Mining for associations between text and brain activation in a functional neuroimaging database

Finn Årup Nielsen[1,2], Lars Kai Hansen[2], Daniela Balslev[1,3]
[1]Neurobiology Research Unit, Rigshospitalet, Copenhagen University Hospital, Copenhagen, Denmark;
[2]Informatics and Mathematical Modelling, Technical University of Denmark, Lyngby, Denmark
[3]Danish Research Centre for Magnetic Resonance, Copenhagen University Hospital Hvidovre, Denmark

January 4, 2005

Abstract

We describe a method for mining a neuroimaging database for associations between text and brain locations. The objective is to discover association rules between words indicative of cognitive function as described in abstracts of neuroscience papers and sets of reported stereotactic Talairach coordinates. We invoke a simple probabilistic framework in which kernel density estimates are used to model distributions of brain activation foci conditioned on words in a given abstract. The principal associations are found in the joint probability density between words and voxels. We show that the statistically motivated associations are well aligned with general neuroscientific knowledge.

Keywords: Databases; Data Interpretation, Statistical; Information Storage and Retrieval; Magnetic Resonance Imaging; Positron-Emission Tomography; Brain Mapping; Meta-Analysis; Neuroimaging; Data mining

1 Introduction

One of the important goals of functional neuroimaging is to associate a specific cognitive function with a specific brain area or a network of such areas. A neuroimaging study typically investigates a cognitive function by manipulating the subjects' behavior and observes the changes in brain activity that correlate with this manipulation. Brain areas with the highest change in activity are reported as a list with three-dimensional coordinates (locations) of the activation foci in the so-called Talairach space, — a stereotactic space allowing approximate comparison of brains from different humans (Talairach and Tournoux, 1988). The accuracy of such 'brain maps', however, depends critically on several factors. The experiment should be able to control the subject's behavior; the data acquisition protocol should be able to accurately measure the changes; the data analysis should be able to tease out the stimulus-induced changes from other fluctuations and confounds; the registration of the activation locations to the standard volume should be accurate. It is very difficult if not impossible for an experiment to fulfill these criteria, hence, a single neuroimaging study represents a 'noisy' measurement of the given effect. To reduce the noise in a neuroimaging database we may investigate for consensus in the brain-behavior relationships across a multiple relevant studies. Ideally, this would be a search through the published neuroimaging studies, in practice we are limited to mining of rather limited databases.

Neuroimaging databases provide interesting multidimensional multimedia data sets, involving the neuroimages, e.g., represented as raw reconstructed volume data acquired by functional magnetic resonance imaging as in the fMRIDC database (Van Horn et al., 2001), or in the form of brain activation foci sets as in the BrainMap, BrainMap DBJ or Brede databases (Fox and Lancaster, 1994; Fox and Lancaster, 2002; Nielsen, 2003). These databases furthermore contain context information typically derived from the paper in which the neuroimage was published.

We will here describe a method for mining for associations between spatial features and the textual description as available in the abstract. We let the words in the abstract represent cognitive function and the list of locations represent the spatial distribution. We mine across all cognitive functions in the entire brain to find the most characteristic fact. A list of associations allows us to make a novel entrance to the information in a neuroscientific database extending our previous methods based on multivariate and asymmetry measures (Nielsen and Hansen, 2004; Nielsen, 2003).

While the present application is original, modeling of combined image and text features is not new, see for example an application involving web-data (Kolenda et al., 2002). Mining for association rules has been described in, e.g., analysis of supermarket transactions (Agrawal et al., 1993) or in geographical information systems (Koperski and Han, 1995). We have previously discussed neuroscientific database technique in terms of cleaning (Nielsen and Hansen, 2002b) and activation volume similarity measures (Nielsen and Hansen, 2004).

2 Method

We use the data from the Brede database which presently records 121 neuroimaging papers with a total of 2655 locations (Nielsen, 2003).

A simple bag-of-words model is used for representation of the abstract: Terms of consecutive letters are extracted and converted to lower-case. Terms that only appear in a single abstract are omitted. An extensive manually constructed *stop word* list of several thousand words is used to filter out words that are expected to carry limited or no semantic information on cognitive function. This, for example, eliminates commonly used words (e.g., “the”, “and”) and terms referring to methodology (e.g., “positron”, “activation”). Since we are interested in modeling abstract associations between context and location we have in this study also eliminated brain anatomy terms (e.g., “temporal”, “prefrontal”).

The bag-of-word representation is not optimal since it does not model sentence structure and does not represent the immediate context, e.g., “time” may appear as “time-specific memories” or “the raw time series”. The bag-of-words representation is chosen basicly because of its simplicity, and surprising efficiency when identifying global context, see e.g., (Lee and Seung, 1999).

The data in the Brede database is organized in a hierarchy adapted from the BrainMap database (Fox and Lancaster, 1994) in which the top level contains scientific papers with bibliographical information. Each of these might have one or more “experiments” which correspond to a single image or “contrast”. Each experiment in turn includes one or more locations. Not all locations that are reported are in full compliance with the Talairach atlas. Some locations are referenced relative to the so-called MNI space, and are transformed before being added to the Brede database (Brett, 1999). We reconstruct the full brainmap activation volume from the listed activated locations by a process we call “voxelization”. In this step the discrete activation foci are first converted to a 3D probability density function and then sampled in a regular grid in the 3D volume. The probability density estimate is obtained by a Parzen kernel density estimator with a Gaussian kernel (Nielsen and Hansen, 2002b; Turkeltaub et al., 2002).

Let $\mathcal{B}_t$ be the set of papers containing the term t, and $\mathcal{L}_b$ the set of Talairach coordinates $\{\mathbf{x}_l : l \in \mathcal{L}_b\}$ in the bth paper. The joint probability of the Talairach space $\mathbf{x}$ and a term t is

$$p(\mathbf{x}, t) = \sum_{b \in \mathcal{B}_t} \sum_{l \in \mathcal{L}_b} p(\mathbf{x}|l)\, p(l|b)\, p(b|t)\, p(t), \tag{1}$$

with $p(\mathbf{x}|l)$ given by

$$p(\mathbf{x}, t) \propto \sum_{b \in \mathcal{B}_t} \sum_{l \in \mathcal{L}_b} (2\pi\sigma^2)^{-3/2} \exp\left[\frac{(\mathbf{x} - \mathbf{x}_l)^\mathsf{T}(\mathbf{x} - \mathbf{x}_l)}{2\sigma^2}\right]. \tag{2}$$

Here we have let the priors be uniform for simplicity: $p(l|b) = p(b|t) = p(t) \propto 1$. The kernel width is fixed in 2 to $\sigma = 10$ millimeter. Alternatively it can be optimized by leave-one-out cross-validation (Nielsen and Hansen, 2002b). Equation 2 is just one form of a number of possible weighting schemes. An alternatively is to set $p(l|b) = 1/|\mathcal{L}_b|$, where the $|\cdot|$ denotes the number of items in the set. This weighting scheme will deemphasize locations that is in a paper with many locations, and individual locations from papers with few location will have large weight. There does not seem to be a simple scheme for choosing an optimal prior.

When a suitable estimate for $p(\mathbf{x}, t)$ is obtained we next consider foci of high joint probability density over Talairach space and terms. The highest "association" $(t^*, \mathbf{x}^*)$ is the maximum in $p(\mathbf{x}, t)$

$$t^*, \mathbf{x}^* = \arg\max_{t,\mathbf{x}} p(\mathbf{x}, t). \tag{3}$$

The highest association for each term is found. For numerical reasons the continuous 3D Talairach space $\mathbf{x}$ is in the present application sampled on a relatively coarse $4 \times 4 \times 4\text{mm}^3$ regular grid. Hence, the probability density $p(\mathbf{x}, t)$ is converted to a discrete probability normalized to unity and we represent the joint probability by the term $\times$ voxel matrix $p(\mathbf{x}, t) \equiv \mathbf{X}$. After sampling the matrix has 304×63648 elements: The joint probability density $p(\mathbf{x}, t)$ is sampled in 63648 grid points (voxels) and 304 terms remain after stop word removal.

Latent semantic analysis (LSI) is a widely used multivariate technique for analysis of text data (Deerwester et al., 1990). In the typical application of LSI a principal component analysis of a term document database is performed and the components with high variance often carry significant semantic content. The components consist of a 'term histogram', i.e., the vocabulary of the given context, and a vector quantifying how much the context is expressed in the set of documents. An alternative to LSI based on non-negative factor has recently been proposed (Lee and Seung, 1999). By retaining a limited set of components the procedure can be considered a denoising scheme in which correlations not relevant for the major contexts are sifted out. We here generalize the non-negative LSI scheme to analysis of spatially coupled text data. We perform K-component non-negative matrix factorization (NMF) (Lee and Seung, 1999) on the term-voxel matrix $\mathbf{X}(T \times V)$

$$\mathbf{X} = \mathbf{W}\mathbf{H} + \mathbf{U}, \tag{4}$$

where $\mathbf{W}(T \times K)$ and $\mathbf{H}(K \times V)$ are the factor matrices and $\mathbf{U}(T \times V)$ contains the residuals. A new K-ranked term $\times$ voxel matrix $\tilde{\mathbf{X}}$ is formed as the product of the reduced positive factor matrices

$$\tilde{\mathbf{X}} = \mathbf{W}\mathbf{H}. \tag{5}$$

We compare in the following results obtained by the original term-voxel matrix and the NMF denoised matrix.

For the present study a number of specialized Matlab functions have been implemented and these have been included in the Brede neuroinformatics toolbox (Nielsen and Hansen, 2000). The toolbox furthermore includes the stop word list and the Brede database. Both the software and the database is presently available on the web from http://hendrix.imm.dtu.dk/.

3 Results and Discussion

Table 3 lists the top of the automatically generated list with terms and x, y and z Talairach coordinates as well as the identifier (WOBIB number) for the bibliographic entry in the Brede database. Only the most salient coordinate is shown for each term.

The first entry in the list is the term "pain" associated with the location $(0, 16, 28)$. This coordinate is in the anterior cingulate according to the labeling produced by the AAL/MRIcro tool (Rorden and Brett, 2000;

	Term	x	y	z	WOBIB
1	pain	0	16	28	13 56 57 58 60 62 69 72 75 76 ...
2	visual	40	−64	−8	1 2 3 5 7 9 14 15 21 23 ...
3	motor	0	−4	52	1 11 16 23 31 40 44 45 57 66 ...
4	memory	0	−52	20	25 27 29 32 35 37 39 49 51 64 ...
5	perception	−36	0	4	13 23 30 33 36 40 42 43 57 62 ...
6	painful	−40	0	4	13 58 60 69 75 83 101 102 114 117 ...
7	noxious	−36	0	4	57 69 72 76 79 95 100 102 117 118
8	retrieval	0	−52	24	20 32 41 50 73 78 80 85 90 105 ...
9	heat	−36	0	4	57 69 75 79 95 100 101 102 114 117 ...
10	thermal	−48	−4	8	56 60 69 72 75 101 102 118
11	patterns	−40	−4	4	16 22 28 30 35 37 45 49 61 69 ...
12	somatosensory	0	8	36	8 13 48 56 57 58 60 62 72 75 ...
13	warm	−40	0	4	60 69 76 102 118
14	verbal	0	20	28	6 10 12 21 25 51 83 85 110 119
15	attention	0	16	28	3 14 18 24 46 58 60 65 83 120
16	hand	−36	−28	52	11 13 45 48 49 57 60 75 84 94 ...
17	perceived	0	20	20	4 52 56 83 95 118
18	sensation	−36	0	4	56 69 72 76 102 118
19	sensory	4	−12	56	14 16 21 22 23 24 31 48 56 58 ...
20	movements	0	−4	56	1 3 11 23 24 45 46 84 104
21	time	28	−56	48	3 10 25 26 30 32 40 52 63 65 ...
22	spatial	0	20	24	3 6 27 44 46 47 53 59 63 65 ...
23	hot	0	16	28	57 60 76
24	interaction	0	16	28	16 18 51 60 76 82
25	cognitive	0	20	24	6 8 51 58 66 70 74 83 85 109 ...
26	finger	−32	−24	52	11 43 44 48 74 108
27	cold	−36	0	4	56 57 58 69 83 102
28	semantic	−4	−52	32	2 10 20 50 78 113 121
29	handed	40	4	12	11 17 25 95 98 100 113 117 118
30	eminence	0	20	24	60 76
31	thenar	0	20	24	60 76
32	auditory	56	−16	0	9 14 19 20 39 42 74 85
33	neutral	−12	−8	−4	4 31 39 65 66 71 93 97 98 117
34	voluntary	−32	−28	52	11 23 84 108
35	integration	−40	−4	4	7 23 48 53 69

Table 1: Automatically generated list with the top associations between terms and the voxels, shown with x, y and z Talairach coordinates in millimeter and Brede database identifier (WOBIB).

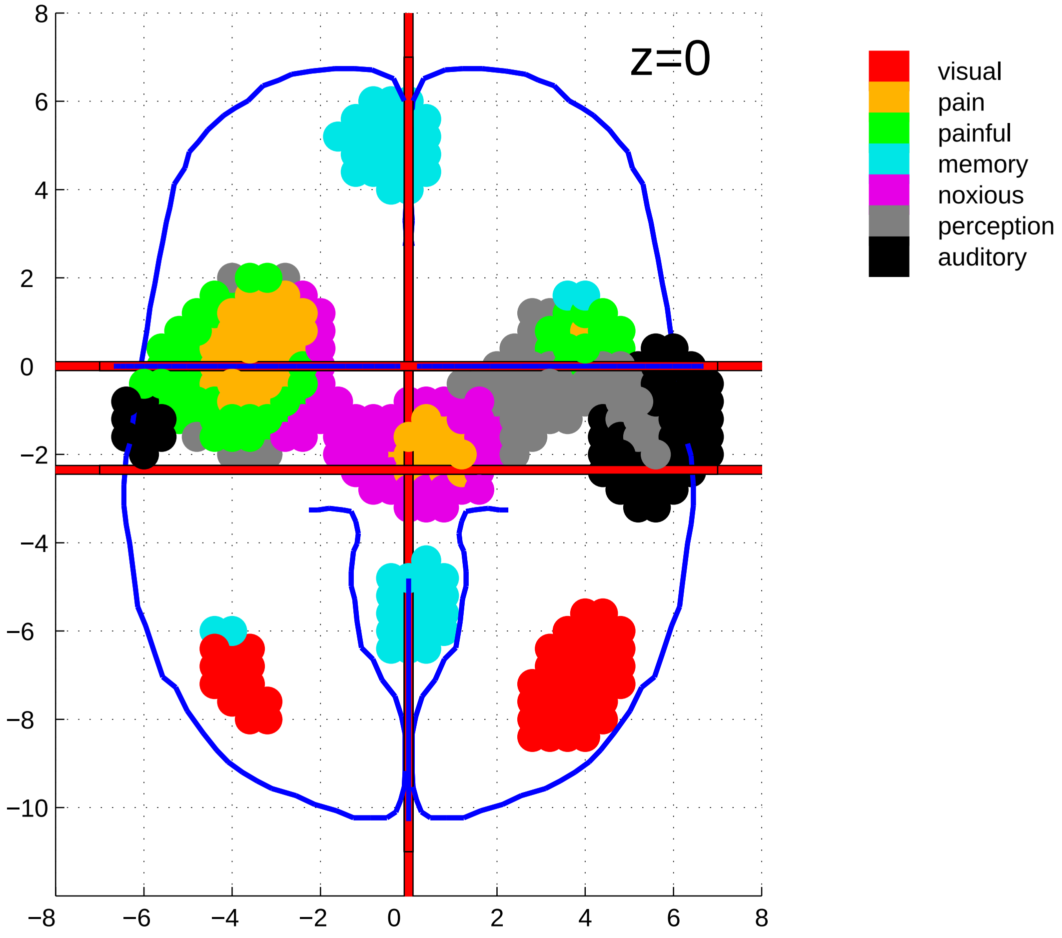

Figure 1: Visualization of the voxel space $\mathbf{x}$ restricted to the $z = 0$ transversal slice. The voxels and terms with the highest associations are plotted together with a blue outline of the brain and red Talairach axes. For each term only the 50 most associated voxels are shown. The labeling on axes are the centimeters from the anterior commissure, — the structure indicating the origo of the Talairach space.

Tzourio-Mazoyer et al., 2002). This is in agreement with an expert review stating "the anterior cingulate cortex (ACC) is the cortical region that is activated in almost every study of elicited pain [...]" (Ingvar, 1999). Figure 1 shows an example where the analysis has been restricted to the $z = 0$ transversal slice (z is the inferior-superior axis). Under this constraint the most important "pain" locations are found in the anterior insula and thalamus — as also found in pain studies (Ingvar, 1999). A number of other pain-related words appear in the list, e.g., painful, noxious and heat. These are associated with voxels in the insular region $(-36, 0, 4)$. Thus the highest association are in alignment with an expert review (Ingvar, 1999). Figure 2 displays the load of the individual studies on the pain components at the voxel with the highest association. 19 papers in total contribute to the association and at the specific voxel with the highest association there are 3 studies that dominates with over 62% of the probability. The reason why the pain topic dominates is that the Brede database presently records many thermal pain studies. Other words are associated with these frequent occurring studies: warm, hot, cold, somatosensory and sensation, and the general term "perception" associates with the same voxel as "noxious".

The second highest is "visual" associated with $(40, -64, -8)$. This is labeled as occipital lobe by the

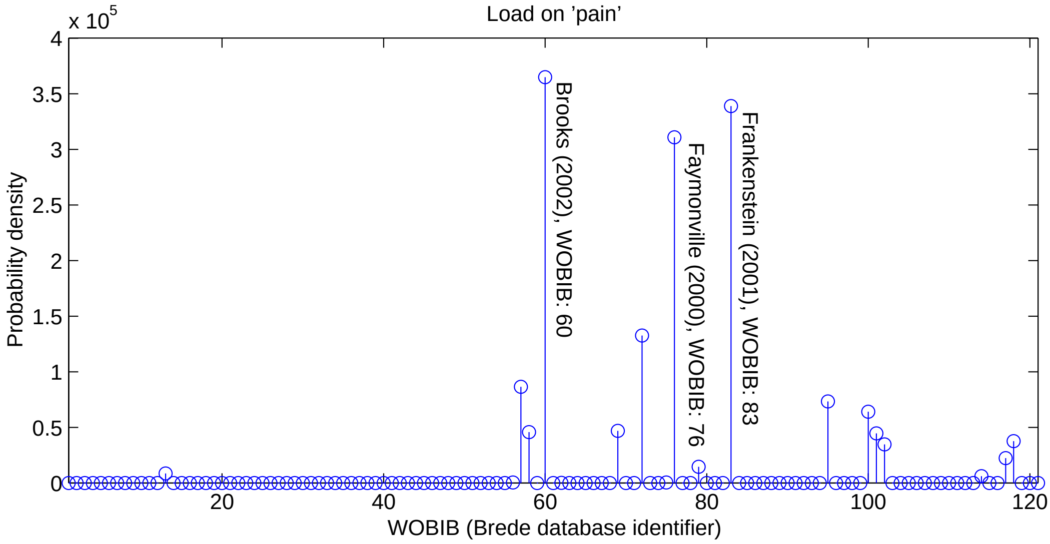

Figure 2: Contributions from the individual studies on the "pain" component at the voxel with the highest association.

Talairach Daemon (Lancaster et al., 1997) and occipital inferior by AAL/MRIcro and is placed laterally compared to the primary visual cortex near the fusiform and temporal gyri. The highest contribution for this association comes primarily from visual object recognition studies such as (Ishai et al., 2000, WOBIB: 28), (Gerlach et al., 1999, WOBIB: 29), (Jordan et al., 2001, WOBIB: 25). Hence, in the context of the present database the term "visual" is associated with visual object recognition areas as this is the dominant visual stimulus type in the relatively limited number of studies included in the database. As more generic visual studies are included the generic term "visual" will presumably associate with earlier visual areas situated medially.

The term "memory" appears as fourth in the list with the Talairach coordinates $(0, -52, 20)$ and the coordinate is labeled "posterior cingulate" by the Talairach Daemon and precuneus by AAL/MRIcro. "Retrieval" is also listed and its major association is close to the "memory" foci. Presently, the Brede database records many studies that activate the posterior cingulate, and episodic memory retrieval is the cognitive function that is presently found to be most associated with this area (Cabeza and Nyberg, 2000), so in the context of the Brede database the association is not surprising.

A number of associations appear for the sensorimotor areas: "hand", "finger" and "motor" near primary and supplementary hand and finger sensorimotor areas. As 32nd in the list "auditory" associates with $(56, -16, 0)$ in the right superior temporal gyrus, i.e., with auditory cortex. Figure 1 shows that "auditory" has high association when the analysis is restricted to $z = 0$.

In table 3 the highest associating terms and Talairach coordinates are shown after the joint matrix has been denoised by the NMF bottleneck (equations 4 and 5). With $K = 20$ components 92.8% of the variance is maintained and it causes only minor re-orderings in the list, e.g., the three top entries are the same. A specific "latent semantic" component corresponding to a column in the right factorized matrix $\mathbf{H}$ is plotted in figure 3. This appears in the posterior inferior part of the brain in both hemispheres and is interpretable as a visual object recognition component.

The textual information (the paper abstract terms) and the brain activation geometry (the locations) form a many-to-many relationship: Each location is not individually associated with specific words. Instead it is a set of words that is associated with sets of location. Where just a few words are associated with a single location as in modeling of the "lobar anatomy" field in the BrainMap database (Nielsen and Hansen, 2002b) the modeling of term t and voxel $\mathbf{x}$ spaces is easier and it is possible to construct atlases and predict

	Term	x	y	z	WOBIB
1	pain	0	12	32	13 56 57 58 60 62 69 72 75 76 ...
2	visual	40	−64	−8	1 2 3 5 7 9 14 15 21 23 ...
3	motor	0	−8	52	1 11 16 23 31 40 44 45 57 66 ...
4	memory	0	−52	20	25 27 29 32 35 37 39 49 51 64 ...
5	perception	−36	0	4	13 23 30 33 36 40 42 43 57 62 ...
6	painful	−40	0	4	13 58 60 69 75 83 101 102 114 117 ...
7	heat	−36	0	4	57 69 75 79 95 100 101 102 114 117 ...
8	retrieval	0	−56	28	20 32 41 50 73 78 80 85 90 105 ...
9	noxious	−36	0	4	57 69 72 76 79 95 100 102 117 118
10	thermal	−44	−4	8	56 60 69 72 75 101 102 118
11	somatosensory	0	4	40	8 13 48 56 57 58 60 62 72 75 ...
12	patterns	−40	−4	4	16 22 28 30 35 37 45 49 61 69 ...
13	warm	−40	0	4	60 69 76 102 118
14	verbal	0	20	24	6 10 12 21 25 51 83 85 110 119
15	sensation	−36	0	4	56 69 72 76 102 118

Table 2: Top of list after the term voxel matrix has been denoised by NMF with $K = 20$ components. This maintains 92.8% of the "variance" as computed by the Frobenius norm of the reconstruction error: $r = 1 - ||\mathbf{X} - \tilde{\mathbf{X}}||_F^2 \,/ ||\mathbf{X}||_F^2$.

labels (Nielsen and Hansen, 2002a). The information on the location level is, however, only anatomical, — not functional. The experiment is an intermediate level between the paper and location levels and this level has functional annotation: In the Brede database there are short texts associated with each experiment and each experiment might be linked to one or more "external component". The external components are functional labels and arranged in a semantic tree. This information is constructed during database entry, i.e., supervised labels. Our present approach is entirely unsupervised and data-driven relying only on the information in the article — not the extra labeling provided during database entry.

Our present method does not distinguish "known" or obvious associations from interesting and novel association, e.g., an example of obvious associations are pain and voxels in the anterior cingulate and insular cortices. For an expert it might appear obvious that pain is associated with these areas. However, what is "expert knowledge" is not always clear, e.g., a text book (Heimer, 1994) written prior to the bulk of functional neuroimaging studies states that "[t]o what extent the perception of pain and temperature requires the cerebral cortex is not clear" and mentions that "part of the anterior cingulate area participates in the perception of pain", but has no reference to the insular area.

4 Conclusion

We have proposed a spatio-textual data mining approach for discovery of association rules between terms and Talairach coordinates in neuroimaging studies. We have illustrated the relative information content between terms and locations by inspecting the principal associations that appear well aligned with neuroscience consensus. Meaningful relations between single terms and single locations appear from the many-to-many relationship between abstracts and sets of Talairach coordinates. The utility of this form of data mining is presently limited by the relatively small databases available. This leads to certain idiosyncrasies induced by 'non-uniform' sampling of neuroimaging papers in the databases, as e.g., in the Brede database where pain

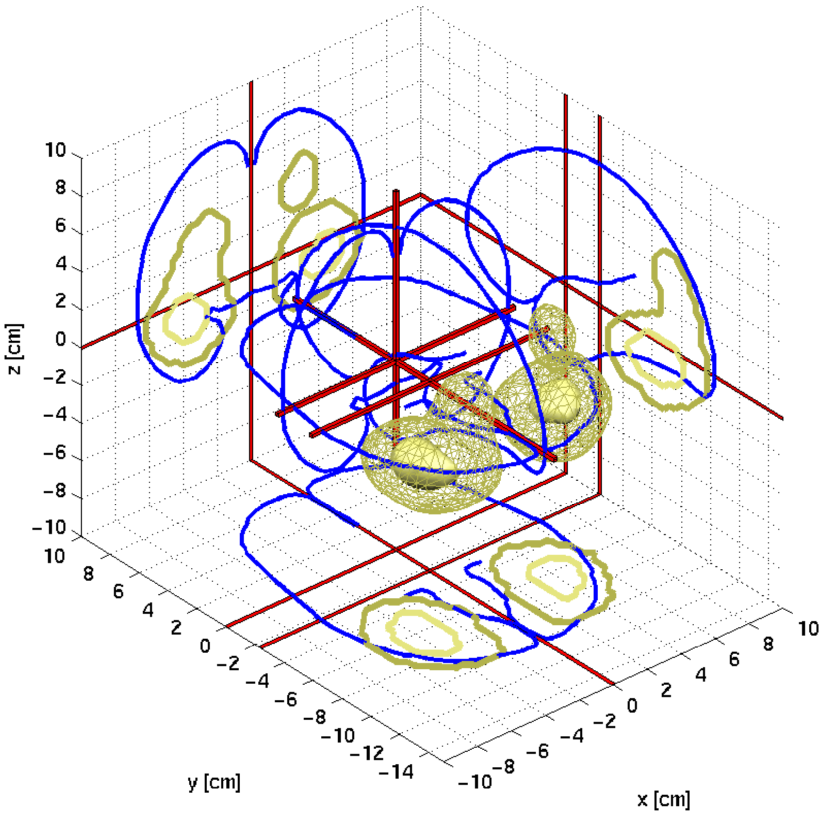

Figure 3: A "latent semantic" component identified by the NMF (a specific row **h** in **H**, c.f., equation 4) plotted in a Corner Cube environment (Rehm et al., 1998). The iso-surfaces are shown with $P = 0.1$ (faces) and $P = 0.5$ (wire-frames) of the top probability inside. The iso-surfaces of the scores in **h** are interpreted as probabilities. The high scoring words from the corresponding row in **W** are "visual", "matching", "category", "perceptual" and "faces" interpretable as visual object recognition. The red lines are Talairach axes and the blue curved lines are the brain outlines.

related studies are relatively abundant. We hope that the intense neuroinformatics activities funded by the NIH's Human Brain Project will increase the size and depth of neuroimaging databases.

Acknowledgments

We thank Tue Lehn-Schiøler for discussions and the four referees that reviewed an earlier version of this paper for the NIPS conference. This work is supported by the European Union project MAPAWAMO, the American Human Brain Project P20 EB02013, and the Villum Kann Rasmussen Foundation by a grant for postdoctoral studies to Finn Årup Nielsen.

References

Agrawal, R., Imielinski, T. and Swami, A. (1993) Mining association rules between sets of items in large databases. In Proceedings of the 1993 ACM SIGMOD International Conference on Management of

Data, Washington, D.C., May 26-28, 1993, Buneman, P. and Jajodia, S., eds. ACM Press, New York, pp. 207–216.

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

Cabeza, R. and Nyberg, L. (2000) Imaging cognition II: An empirical review of 275 PET and fMRI studies. Journal of Cognitive Neuroscience 12, 1–47.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K. and Harshman, R. (1990) Indexing by latent semantic analysis. Journal of the American Society for Information Science 41, 391–407.

Fox, P. T. and Lancaster, J. L. (1994) Neuroscience on the net. Science 266, 994–996.

Fox, P. T. and Lancaster, J. L. (2002) Mapping context and content: the BrainMap model. Nature Reviews Neuroscience 3, 319–321.

Gerlach, C., Law, I., Gade, A. and Paulson, O. B. (1999) Perceptual differentiation and category effects in normal object recognition: a PET study. Brain 122 (Part 11), 2159–2170.

Heimer, L. (1994) The human brain and spinal cord: functional neuroanatomy and dissection guide. Springer-Verlag, New York, 2nd ed.

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

Ishai, A., Ungerleider, L. G., Martin, A. and Haxby, J. V. (2000) The representation of objects in the human occipital and temporal cortex. Journal of Cognitive Neuroscience 12, 35–51.

Jordan, K., Heinze, H. J., Lutz, K., Kanowski, M. and Jancke, L. (2001) Cortical activations during the mental rotation of different visual objects. NeuroImage 13, 143–152.

Kolenda, T., Hansen, L. K., Larsen, J. and Winther, O. (2002) Independent component analysis for understanding multimedia content. In Proceedings of IEEE Workshop on Neural Networks for Signal Processing XII, Bourlard, H., Adali, T., Bengio, S., Larsen, J. and Douglas, S., eds. IEEE Press, pp. 757–766.

Koperski, K. and Han, J. (1995) Discovery of spatial association rules in geographic information databases. In Advances in Spatial Databases. 4th International Symposium SSD '95, Portland, ME, USA, August 6–9, 1995. Proceedings, Egenhofer, M. J. and Herring, J. R., eds. Springer, Heidelberg, Germany, vol. 951 of *Lecture Notes in Computer Science*, pp. 47–66.

Lancaster, J. L., Rainey, L. H., Summerlin, J. L., Freitas, C. S., Fox, P. T., Evans, A. C., Toga, A. W. and Mazziotta, J. C. (1997) Automated labeling of the human brain: A preliminary report on the development and evaluation of a forward-transform method. Human Brain Mapping 5, 238–242.

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

Nielsen, F. Å. (2003) The Brede database: a small database for functional neuroimaging. NeuroImage 19. 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. and Hansen, L. K. (2000) Experiences with Matlab and VRML in functional neuroimaging visualizations. In VDE2000 - Visualization Development Environments, Workshop Proceedings, Princeton, New Jersey, USA, April 27–28, 2000, Klasky, S. and Thorpe, S., eds. Princeton Plasma Physics Laboratory, Princeton, New Jersey, pp. 76–81.

Nielsen, F. Å. and Hansen, L. K. (2002) Automatic anatomical labeling of Talairach coordinates and generation of volumes of interest via the BrainMap database. NeuroImage 16. 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. (2002) Modeling of activation data in the BrainMap$^{\text{TM}}$ database: Detection of outliers. Human Brain Mapping 15, 146–156.

Nielsen, F. Å. and Hansen, L. K. (2004) Finding related functional neuroimaging volumes. Artificial Intelligence in Medicine 30, 141–151.

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, 215–226.

Rorden, C. and Brett, M. (2000) Stereotaxic display of brain lesions. Behavioural Neurology 12, 191–200.

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

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, 765–780.

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, 273–289.

Van Horn, J. D., Grethe, J. S., Kostelec, P., Woodward, J. B., Aslam, J. A., Rus, D., Rockmore, D. and Gazzaniga, M. S. (2001) The functional magnetic resonance imaging data center (fMRIDC): the challenges and rewards of large-scale databasing of neuroimaging studies. Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences 356, 1323–1339.