

Workflows for visualizing data mining results

Finn Årup Nielsen

Center for Integrated Molecular Brain Imaging
DTU Informatics, Technical University of Denmark
Neurobiology Research Unit, Copenhagen University Hospital, Rigshospitalet

April 1, 2009

Abstract

I review methods and tools for visualizations in neuroimaging as well as point to some of the problems with interpreting such visualizations. I furthermore describe Internet-based visualization and information visualization. I focus on workflows associated with the Brede Toolbox for display of data mining results. A few neuroinformatics databases now exist that record results from neuroimaging studies in the form of brain coordinates in stereotaxic space. The Brede Toolbox was originally developed to extract, analyze and visualize data from one of them — the BrainMap database. Since then the Brede Toolbox has expanded and now includes its own database with coordinates along with ontologies for brain regions and functions: The Brede Database. With Brede Toolbox and Database combined we setup automated workflows for extraction of data, mass meta-analytic data mining and visualizations. Most of the Web presence of the Brede Database is established by a single script executing a workflow involving these steps together with a final generation of Web pages with embedded visualizations and links to interactive three-dimensional models in the Virtual Reality Modeling Language.

1 Introduction

In a narrow sense neuroimaging workflows involve neuroimaging image processing and analysis. In a more broader sense the workflow in a neuroimaging study involves a number of other processes: gathering information, designing the experiment, brain scanning, interpretation of the study, relating it to other studies and communicating the study. Data mining in neuroimaging may not only be applied as the standard neuroimaging analysis but also set to work on other components in workflow, and visualization of the data mining results may help the individual researcher in understanding his or her data

as well as in communication with other researchers.

A number of tools exists for visualizing neuroimaging data mining results when the result is a volumetric neuroimage. There are, however, also visualization tools for other aspects of the neuroimaging process, and one example is our Brede Toolbox (Nielsen and Hansen, 2000a). Starting out as a program for handling and visualization of data from the BrainMap database of Fox et al. (1994) the Brede Toolbox now includes its own database of results from neuroimaging—the Brede Database (Nielsen, 2003)—as well as analysis and visualization functions for a range of tasks. Below I will describe some of the tools for neuroimaging visualization, and as our experience has been mostly with the Brede Database and Toolbox I will take a particular focus on workflows and visualizations with these tools. We have setup an automated workflow involving a few non-interactive batch scripts that construct practically the entire Web presence of the Brede Database with static Web pages and visualizations. Furthermore, automated workflows using the ontologies of the Brede Database can perform mass meta-analysis across brain functions or brain regions (Nielsen, 2005; Nielsen et al., 2006a).

2 Visualization tools

The abundance of tools for visualization as well as for other aspects of the neuroimaging process has spawned an interest in generating overviews for these tools, and now there exist several Web-based directories: Neuroscience Database Gateway (NDG) (Gardner and Shepherd, 2004), Neuroscience Information Framework (NIF) (Gardner et al., 2008), Neuroimaging Informatics Tools and Resources Clearinghouse (NITRC) (Buccigrossi et al., 2008), *I Do Imaging* and Internet Analysis Tools Registry (IATR), see also (Dinov et al., 2008). Some of these have an API so that HTML or XML for a specific tool can be re-

quested. The NIF resource may be downloaded as an XML file. NITRC, IATR and I Do Imaging have Web 2.0 components with user-provided tool ratings and NITRC has an associated wiki. Since 2001 I have updated the *Bibliography on Neuroinformatics* which also lists numerous tools. Recently I began the Brede Wiki with structured information about neuroscience including neuroimaging visualization tools. Anyone can ‘micro-publish’ relevant information, and the structured content allows for off-wiki database queries (Nielsen, 2009).

Among the visualization tools are, e.g., MRICron (Rorden and Brett, 2000), MindSeer (Moore et al., 2007), iiV (Lee et al., 2008), DataViewer3D (Gouws et al., 2009), MIPAV (Bazin et al., 2007). Several workflow management systems exists (Fissell, 2007) and they often have associated visualization programs for the neuroimaging data. Presently NITRC has 35 tools listed under visualization, and IATR returns 48 on a query on visualization.

3 Visualizing neuroimages

In the published literature the most used visualization is the display of summary statistics. For example, in the *NeuroImage* issue volume 21 number 1 of May 2006, practically all articles display functional neuroimaging data mining results as slices with a thresholded and color-coded summary statistics image (e.g., the t -value map from statistical parametric mapping) on top of a grey-scale background of an anatomical magnetic resonance image. There are a few other types of visualization present: Projection of the activations on the cortical surface visualized in three dimensions (3D), either with the method implemented in the SPM tool (Frackowiak et al., 2003) or with the more elaborate approach involving cortical surface extraction, that may be further turned into two dimensional (2D) flat-maps where the sulcal and gyral curvature is indicated with a gray-scale, for a recent example see Van Essen and Dierker (2007). This type of visualization is also used for results on cortical thickness analysis. Other types are ‘glass-brain’ visualizations with maximum intensity projection on three orthogonal planes. The few EEG studies in this *NeuroImage* issue use their specialized visualization, and a meta-analysis plot activation coordinates from the literature in 2D and 3D as well as plot showing the ranges of x , y and z coordinates. Other forms of visualization that appear in the literature are, e.g., volume rendering and corner cube visualization. Corner cube visualization combines a 3D rendering with 2D tri-planar plots, where, e.g., thresholded summary statistics appear

as 3D objects with their projection on surrounding 2D ‘walls’ (Rehm et al., 1998). Other more specialized renderings are associated with diffusion tensor imaging and tractography and visualization for estimated EEG dipoles that require glyphs with directionality. Results reported with respect to brain regions—not volumes or coordinates—may be displayed by color-coding a segmented brain.

The typical visualization displays areas that survive a threshold and usually color-codes the area above threshold according to t -value or F -value. There are, however, other elements from the data mining result that may be visualized. What they are can be revealed by what has been termed in the ‘only’ statistical formula by Sackett (2001):

$$\text{Confidence} = \frac{\text{Signal}}{\text{Noise}} \times \sqrt{\text{Sample size}} \quad (1)$$

Thresholding at a certain t -value will, e.g., be based on the ‘confidence’ part of the equation, and that is depended on the signal, the noise and the sample size. In certain situations a researcher is actually interested in the ‘signal’ or the signal to noise ratio. It is possible to plot estimates of these terms. In the General Linear Model the so-called contrast image relates to the signal, while the residual is the noise. The correlation coefficient is a kind of signal to noise ratio. That the issue of what is plotted is not entirely unimportant is seen with the Reimold-effect (Reimold et al., 2005): The combination of spatially varying signal and noise patterns together with spatial smoothing as a step in the image processing may conspire to make it seem as if the ‘activations’ in the t -map area moved to areas of no signal, e.g., white matter. Reimold et al. (2005) suggest viewing the contrast image instead of the t -map.

When viewing the thresholded t -maps or correlation visualizations one should remember that the thresholded voxels are selected in a multiple comparison procedure and the values are no longer unbiased estimates, — unless corrected for. If a scatter plot is made with values of a thresholded voxel across time or subject as a function of a design variable of interest, then the scatter plot and its associated correlation may give the naive reader the impression of a stronger association than actual. This will also be the case for, e.g., a bar plot of percent signal changes for a peak voxel.

The biased correlation has recently been referred to as ‘voodoo correlations’ by Vul et al. (2009) or less colorful ‘biased *post hoc* correlations’ by Lieberman et al. (2009). Methods exist that delivers an unbiased scatter plot. In machine learning it is common to split the data into two sets to get an unbiased estimate of the model performance (Mørch

et al., 1997). This may also be done in more standard neuroimaging analysis with one part of the data acting as a hypothesis-generating data set to determine the region above threshold and the other part acting as the hypothesis-testing data set that gives unbiased estimates within the thresholded region. This procedure is not widespread, but have been applied in a few PET studies, see, e.g., Fiez et al. (1996) and Law et al. (1998). Though these two studies do not have unbiased statistical values for other reasons: the split does not occur across subjects or subjects are excluded in the hypothesis-testing group if they do not survive the threshold.

Interpretation of a visualization *with* a threshold should be done with caution. Such a visualization may hide an area that did not survive the threshold, and should, say, an activation appear in the right hemisphere and not in the left hemisphere, one is (usually) not allowed to conclude that there is a statistical significant difference between the left and right hemispheres: If the threshold is selected on, say, $t = 4$ the right activation may be 4.1 while the left 3.9, and given variation in the data the 0.2 difference may not be significant, see also the discussion by Jernigan et al. (2003). The phenomenon is not limited to neuroimaging and statisticians has discussed it in a paper with the telling title *The Difference Between “Significant” and “Not Significant” is not Itself Statistically Significant* (Gelman and Stern, 2006), which also lists the standard issues with thresholding, e.g., the arbitrariness with the threshold on, e.g., 0.05.

There is a range of non-spatial visualization for neuroimaging data. Some of these pertains to quality assurance, e.g., Keator et al. (2008) report the computation of values such as center-of-mass, mean intensity, per-slice spikiness and have displays of ‘mean of running difference’ and ‘per-slice variation’, and Luo and Nichols (2003) present plots with several different statistical measures to characterize signal and noise, see also Duff et al., this issue. Image processing programs could generate these kinds of visualization as part of a workflow.

Large-scale studies may process individual subjects as they are scanned, and plots can render the evolution of the statistical tests on the hypothesis of interest as more and more subjects are added, see Figure 1 and Frokjaer et al. (2008). In the Brede Toolbox bootstrap resampling gives an indication of the variation of the estimate as more subjects are added. If the scanning of subjects is stopped based on the statistical test it should be noted that multiple dependent comparisons have been performed and the statistical threshold should be corrected. The statistical area of *sequential analysis* deals with this problem.

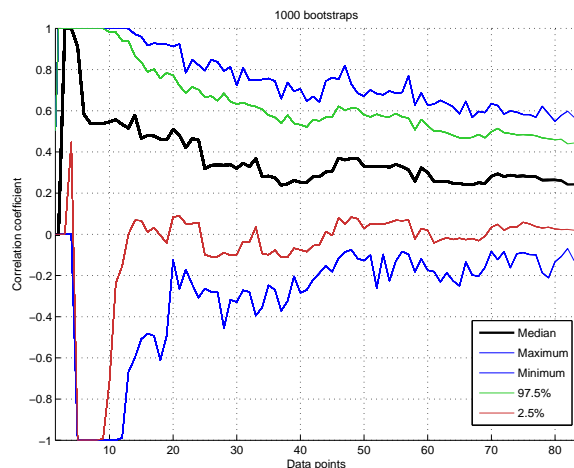


Figure 1: Plot of the evolution of the partial correlation coefficient as the number of subjects included in the analysis is increased from 1 to 83. The lines indicate maximum, 97.5% fraction, median, 2.5% fraction and minimum correlation among 1000 bootstrap samples. Original data from Frokjaer et al. (2008).

Temporal signals, such as the fMRI time courses, are often visualized with line plots. Duann et al. (2002) suggest another way of revealing the time course of fMRI data with the ‘BOLD image’ that plots multiple time-series as an image with the amplitude color-coded. In the Brede Toolbox we also use ‘image’ plots to visualize the data and results in partial correlation analysis, see Figure 2. Kherif et al. (2003) describe one example of visualization with multidimensional scaling (MDS) plots of spatial and temporal similarities between subjects. Similarities are computed as model-based RV coefficients between all pairs of subjects. The resulting MDS shows subjects as individual points in the 2D plots and makes it relatively easy to visually spot temporal or spatial outliers. The ‘receptor fingerprints’ of Zilles and Palomero-Gallagher (2001) display brain region results of multiple dimensions in a polar plot.

4 Visualizing predictors

Standard neuroimaging data mining results in a volume, — often termed a statistical parametric map. In this mode of analysis design variables can be regarded as predictors and the neuroimaging data as response variables. Another mode of data mining uses the neuroimage data to predict the experimental variables (Lautrup et al., 1995). Stan-

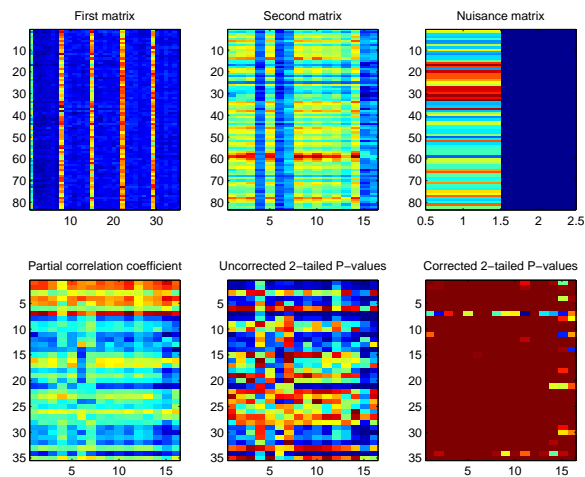


Figure 2: Partial correlation analysis data and results presented in the Brede Toolbox. Multiple personality trait scores, multiple brain regional values from positron emission tomography and nuisance variables (sex and age) for 83 subjects are displayed in the upper row. The lower row presents the results of partial correlation analysis: the correlation matrix and matrices of corrected and uncorrected P -values. Data from Frokjaer et al. (2008)

standard statistics, such as multiple regression, cannot do statistical inference in this case as the predictors form a massively ill-posed problem, and the problem is often recast as a machine learning task. As such the constructed machine learning algorithms can classify patients from healthy controls (Lautrup et al., 1995) or classify fMRI and PET scans with brain activation induced by tasks and stimuli (Mørch et al., 1995, 1997). The approach usually splits the data into a training and a test set. First the model parameters of the machine learning algorithm are estimated on the training set, and then the classification or regression performance of the model is evaluated on the test set. The primary result is not an image that may be visualized, but ‘just’ the prediction model and its performance. However with further computation, images can be formed that tell how much voxels are involved in the prediction. For linear models a weight vector is usually associated with the prediction, and this can be visualized, for a recent example see, e.g., (Mourão-Miranda et al., 2007). For non-linear models, such as multilayer sigmoidal neural networks and non-linear support vector machines, the visualization is not straightforward, but some approaches have been suggested. The *saliency map* technique, inspired from the method of *Optimal Brain Damage* of Le Cun et al. (1990), esti-

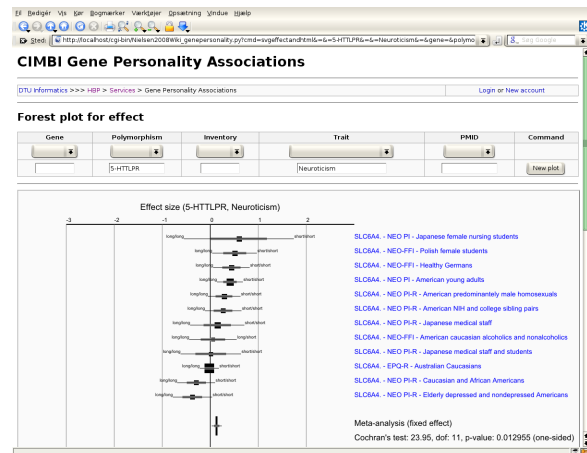


Figure 3: Meta-analytic forest plot as a Web service with studies on personality genetics. Components in the Scalable Vector Graphics image file are hyperlinked and the content may be controlled interactively through a HTML form.

mates how much the elimination of a voxel affects the prediction performance. A closed-form approximation for the effect exists for a trained multilayer neural network and the resulting map can be visualized in 3D brain space (Mørch et al., 1995). The *sensitivity map* asks a related question: What is the effect of a small perturbation of a voxel (Kjems et al., 2002). Both the saliency and sensitivity map can also be applied with linear models. Although linear weight vectors, saliency and sensitivity maps may be compared to results from standard statistical analysis, e.g., by Receiver Operating Characteristics curves and histogram equalization (Lange et al., 1999; Hansen et al., 2001), the values of their individual elements are not directly ‘statistically’ interpretable. For gaining a further statistical interpretation of the magnitude of the individual elements of a linear weight vector, a saliency or sensitivity map these multidimensional results may be compared to results from resampling. One approach, the NPAIRS framework of Strother et al. (2002), uses split-half resampling. Such methods lead to statistically justifiable thresholds for use in visualization of neuroimaging predictors.

5 Meta-analytic visualization

Many meta-analyses use so-called *forest plots* and *funnel plots*, where scatter plots with whiskers display effect sizes and estimators of their variations in two dimensions (Lewis and Clark, 2001), see Figure 3. These meta-analyses typically investigate

a single variable—continuous or dichotomous—and its relation to another variable, e.g., a personality trait and its association with a genetic polymorphism. In neuroimaging meta-analysis we have a quite different situation: The neuroimage result contains not just one variable but many variables, i.e., voxels. One would need thousands of standard meta-analysis plots to capture the result across studies. Another much more fundamental problem stems from the fact that neuroimaging researchers typically only report the positive results, e.g., areas with activation to a given task, — not signal changes for brain regions that did not survive the statistical threshold selected. Meta-analysts usually regard the discarding of negative results as a heresy, referring to it as the *file drawer problem* or with the term *publication bias*. All the standard statistical meta-analysis technique require that also negative results are reported, — at least to some extent (Hedges and Olkin, 1985). So we may ask if it is at all possible to make appropriate analyses and visualizations across studies in neuroimaging?

One simple visualization simply plots the positive results—the reported coordinates—in stereotaxic space. The program associated with the original BrainMap database displayed coordinates in 2D tri-planar plot (Fox et al., 1994). This type of visualization is maintained in a newer version of the database with the program Sleuth (Laird et al., 2005). WebCaret may display coordinates in 3D as colored spheres together with an inflated cortical surface (Van Essen and Dierker, 2007), see Figure 4. The Brede Toolbox can generate 3D visualizations in the corner cube style of Rehm et al. (1998). Plotting points in 3D is not straightforward, — simple ‘zero’ dimensional graphics do not give an important perception of depth, therefore we use 3D glyphs of different color and shape. To help the viewer in spatial localizing the coordinates we can add components in a configurable workflow such as AC/PC axes, stalks for the glyphs, glyph shadows on the tri-planar walls, contour and cerebral cortex outlines from the atlas of Talairach and Tournoux (1988). Figure 5 shows two visualizations of this kind with Figure 5(a) displaying all coordinates in the Brede Database from papers authored by Edward T. Bullmore and Figure 5(b) displaying cingulate coordinates colored according results from a text mining of the associated abstracts (Nielsen et al., 2005, 2006a). The batch script setup for the Brede Database will automatically generate a plot like Figure 5(a) for each author mentioned in the author ontology. Sometimes these simple plots reveal interesting features: The Bullmore coordinates appear somewhat limited to the middle of the inferior-superior axis perhaps reflecting a restricted

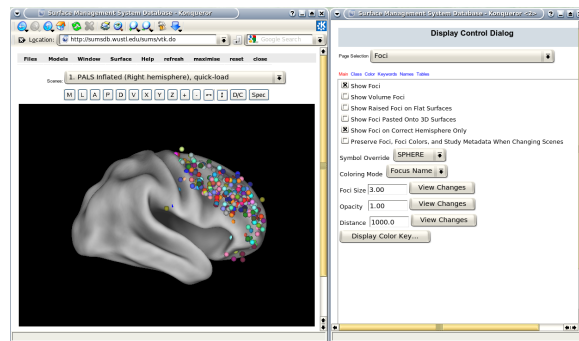


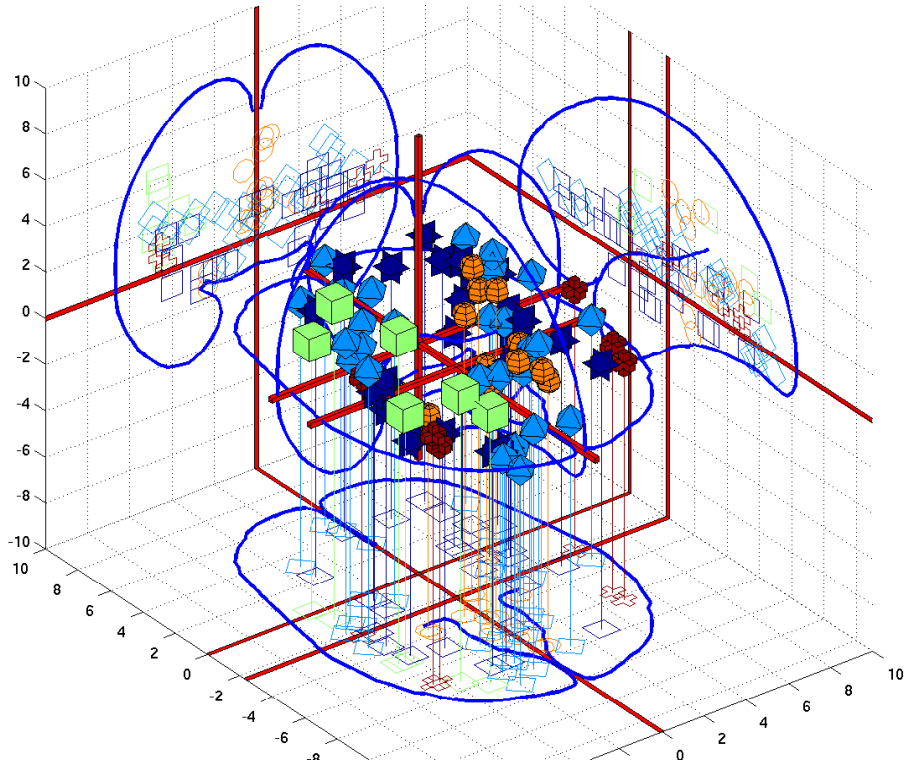
Figure 4: WebCaret server-side display of returned coordinates from the Surface Management System Database (SumsDB) with a query on ‘middle frontal gyrus’. The right window offers some control over the rendering and the buttons in the left window may rotate the cortical surface. SumsDB allows the query on a neuroanatomical label to be invoked from another program or Web site by simple Web linking, and the Brede Wiki automatically constructs such links.

field of view selected for some of the studies. The elaborate and automated workflow for generating a plot like Figure 5(b) involves


1. Select a brain region and from the Brede Database brain region ontology get all naming variation of the brain regions and its sub-areas. With these names extract coordinates from papers recorded in the database, model their spatial distribution and include extra non-matched coordinates that lies within the region.
2. Get abstracts from the Brede Database that—for the brain region in question—have one or more coordinates and perform text mining, which results in clusters of themes, such as ‘pain’ and ‘memory’ and documents belonging to these clusters.
3. Perform statistical tests on the spatial distribution of the coordinates grouped according to the text mining clusters to determine if the text mining has discovered functions that are segregated in the region.

The procedure is done for all brain regions in the Brede Database brain region ontology and Figure 5(b) shows one of the regions that listed high after sorting brain regions according to statistical significance in the spatial distribution test.

Data mining directly with the coordinates has been termed coordinate-based meta-analysis




(a) Coordinates from the five studies in the Brede Database authored by Edward T. Bullmore. The 3D glyphs have type and color according to paper.



(b) Cingulate coordinates colored according to the clustering results after a text mining of abstracts in the Brede Database. Dark magenta glyphs are from the 'memory' cluster while the light yellow are from the 'pain' cluster. From Nielsen et al. (2006a).

Figure 5: Two examples of coordinates in a 3D corner cube visualization.



(a) Gaussian mixture model of the three main behavioral domains in BrainMap: Perception (red wireframe), cognition (green surface) and motion ('M'-textured surface). From Nielsen and Hansen (1999).

(b) Kernel density modeling of auditory (red wireframe) and vision (green) studies. From Nielsen and Hansen (2000b).

Figure 6: VRML visualizations with coordinate-based meta-analysis of data from BrainMap with isosurfaces in conditional probability densities.

(CBMA) and several methods exists (Wager et al., 2009), see also Laird et al., this issue. For the most part they involve a form of estimation of a conditional probability density $p(\mathbf{v}|c)$ in stereotaxic space \mathbf{v} . The conditioning, c , may be, e.g., for a specific brain function or a specific anatomical label. Once the probability density is estimated it can be converted to a volume by sampling the probability density in voxels and visualized in the same way as standard neuroimages, or the density can be used to color-code the cortical surfaces in a 3D visualization, see Wager et al. (2009).

Fox et al. (1997) introduced the method to model the probability density: A single confined area—the primary motor area for the mouth—were examined so only a model with mean and standard deviation was devised, i.e., a simple Gaussian model. As more complex brain functions are distributed in brain space more flexible models are needed. Our first effort in modeling the probability density was by Gaussian mixture models (Nielsen and Hansen, 1999)

$$p(\mathbf{v}|c) = \sum_k^{K_c} p(\mathbf{v}|k) P(k|c), \quad (2)$$

where each $p(\mathbf{v}|k)$ estimates a 3D Gaussian probability density. Figure 6(a) shows the isosurfaces in model of this type where the parameters have

been fitted to data from the BrainMap database. Here each ellipsoids corresponds to a single Gaussian $p(\mathbf{v}|k)$ and c corresponds to three different labels of 'behavioral domain' from the BrainMap database that are associated with each coordinate. Although the Gaussian mixture model may generalize, the ellipsoids do not look neuroanatomical plausible and call for yet more flexible models. Figure 6(b) is generated with kernel density estimation using a Gaussian kernel (Nielsen and Hansen, 2000b). Such models seems to generate probabilities that are somewhat more neuroanatomical plausible than the Gaussian mixture model.

The isosurfaces in the probability densities in both subplots of Figure 6 has been set for display purpose. More statistically grounded values can be obtained with the methods by Turkeltaub et al. (2002); Nielsen (2005); Costafreda et al. (2009). The methods for probability density estimation of coordinates are not limited to activations but may be applied to any kind of coordinates in stereotaxic space from 'deactivations', cortical stimulations, lesions or structural changes, e.g., obtained with voxel-based morphometry.

When a probability density estimate is constructed for a set of coordinates and it is converted to a voxel-volume, then the volumes across multiple sets of coordinates may be aggregated into

a single data matrix \mathbf{X} (sets \times voxels). This data matrix may then be decomposed with multivariate analysis in a number of ways, e.g., with singular value decomposition for principal component analysis, $\mathbf{ULV}' = \mathbf{X}$, where the left factorization matrix \mathbf{U} (sets \times components) contains loading over sets of coordinates for each principal component and the right factorization matrix \mathbf{V} (voxel \times components) contains loadings over voxels. Other types of decomposition for this matrix is independent component analysis ($\mathbf{MS} = \mathbf{X}$, with \mathbf{M} the mixing matrix and \mathbf{S} the source matrix), non-negative matrix factorization ($\mathbf{WH} = \mathbf{X}$) and K-means clustering ($\mathbf{CA} = \mathbf{X}$, with \mathbf{C} a centroid matrix and \mathbf{A} an assignment matrix). The right decomposition matrices, \mathbf{V} , \mathbf{S} , \mathbf{H} and \mathbf{A} all contain vectors that each represents a volume. As part of the workflow for presenting the information in the Brede Database on the Web the decompositions work on data matrices formed from sets of papers and sets of experiments, and corner cube visualizations are automatically constructed with isosurfaces in the volumes contained in the right decomposition matrices. Figure 7 shows such a visualization for a component from non-negative matrix factorization, i.e., a row in the \mathbf{H} matrix. Such visualizations may be useful for navigating among the studies in the database, and to a certain extent they reveal spatial distributions of the ‘cognitive components’ of the brain. Together with the visualization on the Web page is listed the experiments that have high association with the component, i.e., experiments associated with large elements in a column of the left matrix \mathbf{W} . For the component in Figure 7 they are experiments described as, e.g., ‘Visual object decision’, ‘Buildings visual objects’, ‘Color perception during free viewing’ and ‘Passively viewed scenes’.

Before putting too much trust in visualizations and analysis across studies one needs to remember that the study results may have arisen in quite different ways. In standard meta-analysis the only variations between studies that are usually modeled is the number of subjects and the standard deviation of the data in the individual studies. In neuroimaging meta-analysis and visualization these variables are not usually modeled, for exceptions see Fox et al. (1997) and Laird et al., this issue. Besides there are several other variables that neither are considered: The varying thresholds applied, e.g., corrected and uncorrected P -values (Nielsen et al., 2006b), the difference in field of view between studies, the reporting style of coordinates (e.g., ‘extent threshold’, ‘number of maxima per cluster’) as well as the variation from the different pre-processing and analysis choices that have been made. Furthermore, the different CBMA model




Figure 7: Corner cube visualization on the Web page of the Brede Database with results from a non-negative matrix factorization of experiments in the database.

may produce different results on the same material. Salimi-Khorshidi et al. (2009) compared different CBMA models and their application of a threshold makes a ‘blob’ appear and disappear depending on the type of CBMA.

6 Internet-based visualization

Quite a few tools exist for interactive neuroimaging visualization across the Internet. Often these tools are based on a client-server model with the client implementing the visualization and graphical user interface in Java. Among these tools are JIV that renders multiple volume data by orthogonal slice views implemented as a Java applet (Cocosco and Evans, 2001). iiV implements a similar functionality (Lee et al., 2008), and MindSeer can also render in 3D remotely (Moore et al., 2007). NeuroTerrain implements 3D visualization and has demonstrated its use in connection with a Mouse atlas (Gustafson et al., 2007). The Talairach Applet renders a digital representation of the Talairach Atlas and combines it with neuroanatomical labeling of coordinates via the Talairach Daemon described by Lancaster et al. (2000). Also in connection with the BrainMap database the Java client-program Sleuth plots 3D points in orthogonal 2D slices based on user query to the BrainMap server (Laird et al., 2005).

The *Internet Brain Volume Database* (IBVD) records published values for brain region volumes across variables such as gender and diagnosis (Kennedy et al., 2003). Since the neuroimaging data analysis arrives at one single value—the brain

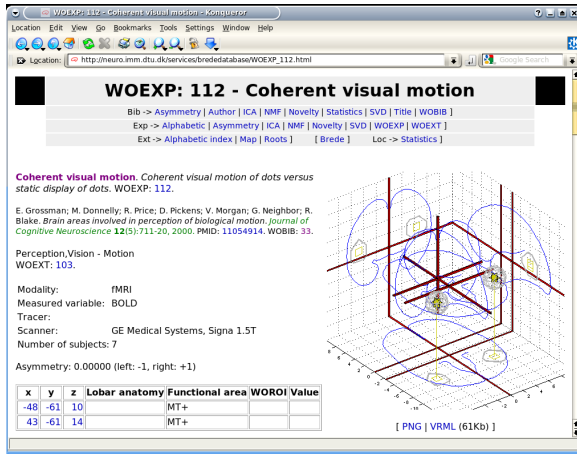


Figure 8: Screenshot of the Web page for an experiment in Brede Database with a corner cube visualization of the coordinates in a experiment together with a wireframe indicating an isosurface of the kernel density estimate with the coordinates. An interactive rendering is provided with the link to a generated corner cube visualization in a VRML file.

volume in cubic centimeters—the visualization of the data is relatively simple compared to other neuroinformatics visualizations: From Web-based user queries IBVD generates on-the-fly PNG image-files with the brain volumes from the different studies plotted as a function of age with color-coding and the variability indicated. Interactive visualization systems for neuroimages with server-side 3D rendering have been described by Poliakov et al. (2005) and a public system is available with the WebCaret Web service, see Figure 4.

With the Brede Toolbox we construct 3D visualizations browsable on the Web by using the Virtual Reality Modeling Language (VRML) (ISO/IEC, 1997; Nielsen and Hansen, 2000a), see the VRML examples in Figure 6. When defined in the middle of 1990s VRML held great promise to get widespread use for 3D interactive and hyperlinked visualizations, but since then it has had limited growth: VRML lacks good browser implementations and there has been erratic adoption of a scripting language. Nevertheless, it is one of the few means for Web distribution of 3D content in free standardized form. An alternative format is the *Universal 3D File Format* (U3D) that can be embedded in newer versions of the PDF format. Apart from the Brede Toolbox *ImageSurfer* described by Feng et al. (2007) implements VRML export.

For the Web presentation of the Brede Database we generate 3D corner cube visualizations of the

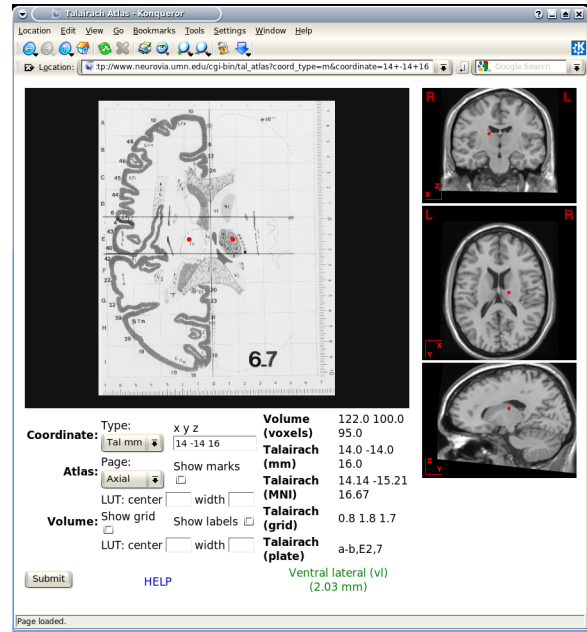


Figure 9: The Web-based INC Interactive Talairach Atlas queried with a coordinate from the Brede Wiki.

coordinates in the database with an offline Matlab batch script, — both as image files embedded on the Web page as well as VRML files, see Figure 8. Matlab is not well suited to work as a Web script, and for the interactive Web scripts associated with the Brede Database there are presently no visualization implemented. The INC *Interactive Talairach Atlas* renders 2D orthogonal slices from the Talairach and the MNI single subject atlases. This Web service can merge a user-given coordinate with the visualization, and as such we use it for visualization of individual coordinates from the Brede Database and the Brede Wiki.

Besides Java, VRML and standard image files such as PNG the *Scalable Vector Graphics* (SVG) format may prove useful for Internet-based visualizations, see Figure 3 for an example. These files may contain hyperlinks and JavaScript. However, Web browsers do not yet consistently implement the standard.

7 Information visualization

Data mining results from neuroimaging analysis are not the only type of information for visualization. Information about the background, design, scanning, analysis procedure, and interpretation surrounds the data mining results of a typical neuroimaging study. In scientific articles the body text

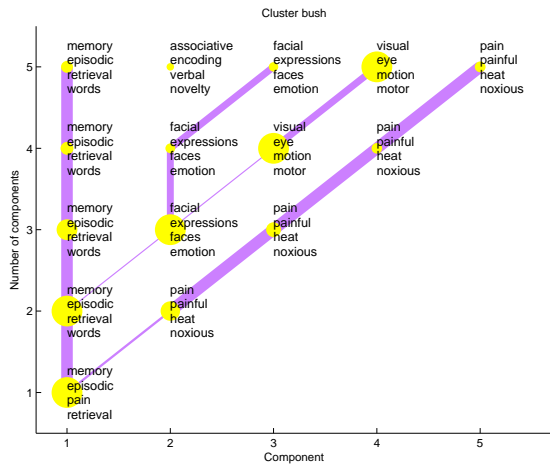


Figure 10: A so-called ‘cluster bush’ visualization of the text mining results of the abstracts in the Brede Database. Each yellow dot is a cluster of articles and words in the article. The four words with highest load on each cluster are listed.

mostly carries this ‘context’ information, though sometimes authors also use tables to describe, e.g., subject information. Authors rarely apply visualizations for this kind of information except in situations with explanation of the experimental design and scanning. The experimental design has a natural temporal evolution and as such the visualization often displays the design as a function of time. Users of the behavioral experiment software from Psychology Software Tools is familiar with the graphical programming environment of *E-Prime* which has this kind of visualization as an integral part of the development of the experiment. Other parts of the neuroimaging study may be visualized with what is usually referred to as *information visualization*.

In a demonstration visualization we employed a torus topology for the an entire neuroimaging study process constructing 3D icons for ‘funding’, the experimental design, authors, experimental subjects etc. (Nielsen and Hansen, 1997). The usefulness of such a visualization depends on how effective it conveys information compared to standard text, and if the visualization format requires specialized and limited distributed programs for rendering and interaction the impact may be small. Manual creation of these visualizations is infeasible, — the visualization should be constructed automatically from description of the study, e.g., the so-called ‘provenance’ (Fissell, 2007). In related visualizations some workflow management systems display




Figure 11: Corner cube visualization of a ‘functional’ atlas generated from a combined data mining of text and brain coordinates.

the processing flow graphically (Dinov et al., 2008).

When neuroimaging studies get reported in articles the relationships between the articles can be turned in to visualizations. Many types of visualizations exist and many relationships may be revealed: Between terms, concepts, citations to and from articles as well as between authors, cited authors and cited journals. The visualizations are of course not limited to articles only in neuroimaging, see, e.g., Card et al. (1999); Chen (1999). For an example in neuroscience Naud et al. (2007) use a spherical embedding algorithm to display a bipartite graph in 3D space with two spheres. One of their illustrations visualized the relationship between poster sessions in the Society for Neuroscience 2006 meeting together with words from the abstracts in the sessions. Another example of text mining result visualization is what we termed a ‘cluster bush’, that describe the clusters in a hierarchical multivariate analysis (Nielsen et al., 2005): Clusters are indicated with dots and thick lines indicate a large similarity between two clusters. Given a set of abstracts the automated workflow for generating a plot like Figure 10 involves the conversion of the texts to a bag-of-words matrix, the exclusion of a large number of words (stop words), hierarchical non-negative matrix factorization and lastly the ‘cluster bush’ visualization all implemented with the functions of the Brede Toolbox.

Coordinate-based meta-analysis and text mining can be combined to form visualizations, see Figure 11 and Nielsen et al. (2004). The workflow for constructing the visualization in the figure involves the setup of a matrix describing the words in




Figure 12: Visualization of data mining result of journal co-citation analysis with singular value decomposition on citation data from NeuroImage.

the abstract of papers and the construction of another matrix from kernel density estimation with the coordinates in each paper. After non-negative matrix factorization each individual factor may be rendered in 3D and associated with words from the abstract, e.g., the blue area in Figure 11 in the occipital lobe is associated with words such as ‘visual’ and ‘eye’.

Based on a corpus of articles published between 1997 and 2000 in the journal *NeuroImage* we could plot cited authors and cited journals in 2D. The data mining with visualization would for example reveal a dichotomy between PET and fMRI (Nielsen, 2002), see Figure 12. Here the workflow involves specialized algorithms that extract citations and the use of matrix computations, particularly singular value decomposition, for multidimensional scaling-like projection of the data onto 2D. For the Brede Database we automatically construct what we have termed ‘bullseye plots’ to display the network of coauthors for each recorded author. Figure 13 shows a larger bullseye plot on coauthors

in the NeuroImage corpus. Authors near the center, such as Friston and Dolan, have high network degrees, which here corresponds to the number of authored articles (Nielsen, 2002).

The well-tested and widely used GraphViz package provides spatial graph layout for a given network (Gansner and North, 2000). At one point the *PubGene* Web service used GraphViz in a large-scale application for displaying relations between genes based on literature in PubMed (Jenssen et al., 2001). GraphViz layouts graphs for the Web presentation of the Brede Database. These graphs display the brain function and brain region ontologies, e.g., indicating that ‘vision’ has ‘perception’ as taxonomic parent or that the cingulate area is a parent for the posterior cingulate, see Figure 14. Our workflow with the Brede Toolbox involves extraction of the ontology from Brede Database XML files, construction of a file with the graph that GraphViz reads, invoking GraphViz for generation of an image file, and then finally construction of the Web page with the image file embedded. GraphViz




Figure 13: Coauthor bullseye plot (target diagram) with data from NeuroImage 1997–2000. A line between two authors indicates that they co-wrote a paper. The concentric circles indicate the number of articles written by the author in the corpus.

can construct HTML image maps so the nodes in the graph image are associated with clickable hyperlinks. On the final Web page a reader may navigate the brain region and brain function ontologies by clicking on the nodes in the graph. The Brede Toolbox can also use GraphViz for layout of other types of data that can be described as a network, e.g., from structural equation modeling of regional neuroimaging data. A number of journal Web sites use plots called *Citation map* in the style of GraphViz for visualizing in- and out-going citations of each article, see, e.g., *BMJ* and *The Journal of Neuroscience* Web sites.

Another type of graph visualization within neuroimaging is the interactive graph visualization with a hyperbolic browser that features in tools from the Laboratory of Neuro Imaging (LONI): LOVE and iTools (Dinov et al., 2006, 2008). *ISI Web of Knowledge* provides a Java applet to render

their citation information with a similar topology.

8 Conclusion

Numerous visualization tools exist for display of neuroimages, and several Web directories record these tools. There are different types of results from neuroimaging analysis that can be visualized, and one should be aware of the problems that a naive reader faces when viewing such results.

With the Brede Toolbox we are able to build a workflow with extraction of data from the Brede Database, automated data mining and visualizations. The automated procedures generate publicly accessible Web pages with interactive visualizations. An advantage of the automated procedure is that little human intervention is required to update the visualizations as new data is added to the

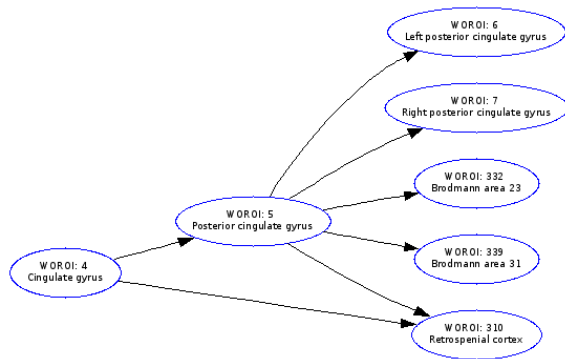


Figure 14: For presenting the Brede Database brain region ontology on the Web the workflow with the Brede Toolbox invokes the GraphViz program which generates hyperlinked plot of the brain region hierarchy, here for the 'posterior cingulate gyrus'.

database. The visualizations can display not only spatial neuroimages, but for example also results from text mining, and visualization can take place across the Internet with data originating on one server and displayed on another.

Acknowledgment

Many thanks to Kristoffer Hougaard Madsen, Daniela Balslev and Lars Kai Hansen for comments on an earlier version of the manuscript. This work was supported by the Lundbeck Foundation through the Center for Integrated Molecular Brain Imaging.

References

- Bazin, P.-L., Cuzzocreo, J. L., Yassa, M. A., Gandler, W., McAuliffe, M. J., Bassett, S. S., and Pham, D. L. (2007). Volumetric neuroimage analysis extension for the MIPAV software package. *Journal of Neuroscience Methods*, 165(1):111–121.
- Buccigrossi, R., Ellisman, M., Grethe, J., Haselgrove, C., Kennedy, D., Martone, M., Preuss, N., Sullivan, M., and Wagner, K. (2008). The neuroimaging informatics tools and resources clearinghouse (NITRC). In *14th Annual Meeting of the Organization for Human Brain Mapping*, pages 319 T–AM. Organization for Human Brain Mapping.
- Card, S. K., MacKinlay, J. D., and Shneiderman, B., editors (1999). *Readings in Information Visualization. Using Vision to think*. The Morgan Kaufmann Series in Interactive Technologies. Morgan Kaufmann Publishers, San Francisco, California, USA.
- Chen, C. (1999). Visualising semantic spaces and author co-citation networks in digital libraries. *Information Processing and Management*, 35(3):401–420.
- Cocosco, C. A. and Evans, A. C. (2001). Java internet viewer: A WWW tool for remote 3D medical image data visualization and comparison. In *Proceedings of the 4th International Conference on Medical Image Computing and Computer-Assisted Intervention*, volume 2208 of *Lecture Notes In Computer Science*, pages 1415–1416, London, UK. Springer.
- Costafreda, S., David, A. S., and Brammer, M. J. (2009). A parametric approach to voxel-based meta-analysis. *NeuroImage*, doi:10.1016/j.neuroimage.2009.01.031.
- Dinov, I. D., Rubin, D., Lorensen, W., Dugan, J., Ma, J., Murphy, S., Kirschner, B., Bug, W., Sherman, M., Floratos, A., Kennedy, D., Jagadish, H. V., Schmidt, J., Athey, B., Califano, A., Musen, M., Altman, R., Kikinis, R., Kohane, I., Delp, S., Parker, D. S., and Toga, A. W. (2008). *iTools*: a framework for classification, categorization and integration of computational biology resources. *PLoS ONE*, 3(5):e2265.
- Dinov, I. D., Valentino, D., Shin, B. C., Konstantinidis, F., Hu, G., MacKenzie-Graham, A., Lee, E.-F., Shattuck, D., Ma, J., Schwartz, C., and Toga, A. W. (2006). LONI visualization environment. *Journal of Digital Imaging*, 19(2):148–158.
- Duann, J.-R., Jung, T.-P., Kuo, W.-J., Yeh, T.-C., Makeig, S., Hsieh, J.-C., and Sejnowski, T. (2002). Single-trial variability in event-related BOLD signals. *NeuroImage*, 15(4):823–835.
- Feng, D., Marchburn, D., Jen, D., II, R. J. W. R. M. T., and Burette, A. (2007). Stepping into the third dimension. *The Journal of Neuroscience*, 27(47):12757–12760.
- Fiez, J. A., Raife, E. A., Balota, D. A., Schwarz, J. P., Raichle, M. E., and Petersen, S. E. (1996). A positron emission tomography study of the short-term maintenance of verbal information. *The Journal of Neuroscience*, 16(2):808–822.
- Fissell, K. (2007). Workflow-based approaches to neuroimaging analysis. *Methods in Molecular Biology*, 401:235–266.

- 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.
- 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.
- Frackowiak, R. S. J., Friston, K. J., Dolan, R., and Price, C. (2003). *Human Brain Function*. Academic Press, second edition.
- Frokjaer, V., Mortensen, E. L., Nielsen, F. Å., Haugbøl, S., Pinborg, L. H., Adams, K. H., Svarer, C., Hasselbalch, S. G., Holm, S., Paulson, O. B., and Knudsen, G. M. (2008). Frontolimbic serotonin 2A receptor binding in healthy subjects is associated with personality risk factors for affective disorder. *Biological Psychiatry*, 63(6):569–576.
- Gansner, E. R. and North, S. C. (2000). An open graph visualization system and its applications to software engineering. *Software — Practice and Experience*, 30(11):1203–1234.
- Gardner, D., Akil, H., Ascoli, G., Bowden, D. M., Bug, W., duncan E. Donohoe, Goldberg, D. H., Grafstein, B., Grethe, J. S., Gupta, A., Halavi, M., Kennedy, D. N., Marengo, L., Martone, M. E., Miller, P. L., Müller, H.-M., Robert, A., Shepherd, G. M., Sternberg, P. W., Van Essen, D. C., and Williams, R. W. (2008). The Neuroscience Information Framework: A data and knowledge environment for neuroscience. *Neuroinformatics*, 6(3):149–160.
- Gardner, D. and Shepherd, G. M. (2004). A gateway to the future of neuroinformatics. *Neuroinformatics*, 2:271–274.
- Gelman, A. and Stern, H. (2006). The difference between “significant” and “not significant” is not itself statistically significant. *The American Statistician*, 60(4):1–4.
- Gouws, A., Woods, W., Millman, R., Morland, A., and Green, G. (2009). DataViewer3D: An open-source, cross-platform multi-modal neuroimaging data visualization tool. *Frontiers in Neuroinformatics*, 3:9.
- Gustafson, C., Bug, W. J., and Nissanov, J. (2007). NeuroTerrain — a client-server system for browsing 3D biomedical image data. *BMC Bioinformatics*, 8:40.
- Hansen, L. K., Nielsen, F. Å., Strother, S. C., and Lange, N. (2001). Consensus inference in neuroimaging. *NeuroImage*, 13(6):1212–1218.
- Hedges, L. V. and Olkin, I. (1985). *Statistical Methods for Meta-Analysis*. Academic Press, Orlando, Florida.
- ISO/IEC (1997). *Information technology — Computer graphics and image processing — The Virtual Reality Modeling Language (VRML) — Part 1: Functional specification and UTF-8 encoding*. International Organization for Standardization/International Electrotechnical Commission. International Standard ISO/IEC 14772-1:1997.
- 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.
- Jernigan, T. L., Gamst, A. C., Fennema-Notestine, C., and Ostergaard, A. L. (2003). More “mapping” in brain mapping: Statistical comparison of effects. *Human Brain Mapping*, 19(2):90–95.
- Keator, D. B., Grethe, J. S., Marcus, D., Ozyurt, B., Gadde, S., Murphy, S., Pieper, S., Greve, D., Notestine, R., Bockholt, H. J., Papadopoulos, P., BIRN Function, BIRN Morphometry, and BIRN-Coordinating (2008). A national human neuroimaging collaboratory enabled by the biomedical informatics research network (BIRN). *IEEE Transactions on Information Technology in Biomedicine*, 12(2):162–172.
- Kennedy, D. N., Haselgrove, C., and McInerney, S. (2003). MRI-based morphometric analysis of typical and atypical brain development. *Mental Retardation and Developmental Disabilities Research Reviews*, 9(3):155–160.
- Kherif, F., Poline, J.-B., Mériaux, S., Benali, H., Flandin, G., and Brett, M. (2003). Group analysis in functional neuroimaging: selecting subjects using similarity measures. *NeuroImage*, 20:2197–2208.
- Kjems, U., Hansen, L. K., Anderson, J., Frutiger, S., Muley, S., Sidtis, J., Rottenberg, D., and Strother, S. C. (2002). The quantitative evaluation of functional neuroimaging experiments: mutual information learning curves. *NeuroImage*, 15(4):772–786.

- Laird, A. R., Lancaster, J. L., and Fox, P. T. (2005). BrainMap: the social evolution of a human brain mapping database. *Neuroinformatics*, 3(1):65–78.
- 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.
- Lange, N., Strother, S. C., Anderson, J. R., Nielsen, F. Å., Holmes, A. P., Kolenda, T., Savoy, R., and Hansen, L. K. (1999). Plurality and resemblance in fMRI data analysis. *NeuroImage*, 10(3):282–303.
- Lautrup, B., Hansen, L. K., Law, I., Mørch, N. J. S., Svarer, C., and Strother, S. C. (1995). Massive weight sharing: A cure for extremely ill-posed problems. In Herman, H. J., Wolf, D. E., and Pöppel, E. P., editors, *Proceedings of the Workshop on Supercomputing in Brain Research: From Tomography to Neural Networks*, pages 137–148, Singapore. World Scientific.
- Law, I., Svarer, C., Rostrup, E., and Paulson, O. B. (1998). Parieto-occipital cortex activation during self-generated eye movements in the dark. *Brain*, 121(11):2189–2200.
- Le Cun, Y., Denker, J. S., and Solla, S. A. (1990). Optimal brain damage. In Touretzky, D. S., editor, *Advances in Neural Information Processing Systems: Proceedings of the 1989 Conference*, pages 598–605, San Mateo, CA. Morgan-Kaufmann. NIPS-2.
- Lee, J. T., Munch, K. R., Carlis, J. V., and Pardo, J. V. (2008). Internet image viewer (iiV). *BMC Medical Imaging*, 8:10.
- Lewis, S. and Clark, M. (2001). Forest plots: trying to see the wood and the trees. *BMJ*, 322(7300):1479–1480.
- Liebermann, M. D., Berkman, E. T., and Wager, T. D. (2009). Correlations in social neuroscience aren't voodoo: A reply to Vul et al. *Perspectives on Psychological Science*. In press.
- Luo, W.-L. and Nichols, T. E. (2003). Diagnosis and exploration of massively univariate neuroimaging models. *NeuroImage*, 19(3):1014–1032.
- Moore, E. B., Poliakov, A. V., Lincoln, P., and Brinkley, J. F. (2007). MindSeer: a portable and extensible tool for visualization of structural and functional neuroimaging data. *BMC Bioinformatics*, 8:389.
- Mørch, N. J. S., Hansen, L. K., Strother, S. C., Svarer, C., Rottenberg, D. A., Lautrup, B., Savoy, R., and Paulson, O. B. (1997). Nonlinear versus linear models in functional neuroimaging: Learning curves and generalization crossover. In Duncan, J. and Gindi, G., editors, *Information Processing in Medical Imaging*, volume 1230 of *Lecture Notes in Computer Science*, pages 259–270, Berlin/Heidelberg. Springer.
- Mørch, N. J. S., Kjems, U., Hansen, L. K., Svarer, C., Law, I., Lautrup, B., Strother, S. C., and Rehm, K. (1995). Visualization of neural networks using saliency maps. In *Proceedings of 1995 IEEE International Conference on Neural Networks*, volume 4, pages 2085–2090.
- Mourão-Miranda, J., Friston, K. J., and Brammer, M. (2007). Dynamic discrimination analysis: A spatial-temporal SVM. *NeuroImage*, 36:88–89.
- Naud, A., Usui, S., Ueda, N., and Taniguchi, T. (2007). Visualization of documents and concepts in neuroinformatics with the 3D-SE viewer. *Frontiers in Neuroinformatics*, 1:7.
- Nielsen, F. Å. (2002). *Neuroinformatics in Functional Neuroimaging*. PhD thesis, Informatics and Mathematical Modelling, Technical University of Denmark, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby. IMM-PHD-2001-87.
- Nielsen, F. Å. (2003). The Brede database: a small database for functional neuroimaging. *NeuroImage*, 19(2). 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.
- Nielsen, F. Å. (2009). Brede Wiki: Neuroscience data structured in a wiki. Submitted to SemWiki2009.
- 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.

- Nielsen, F. Å., Balslev, D., and Hansen, L. K. (2006a). Data mining a functional neuroimaging database for functional segregation in brain regions. In Olsen, S. I., editor, *Den 15. Danske Konference i Mønstergenkendelse og Billedanalyse*, Copenhagen, Denmark. The Department of Computer Science, University of Copenhagen.
- Nielsen, F. Å., Christensen, M. S., Madsen, K. H., Lund, T. E., and Hansen, L. K. (2006b). fMRI neuroinformatics. *IEEE Engineering in Medicine and Biology Magazine*, 25(2):112–119.
- Nielsen, F. Å. and Hansen, L. K. (1997). Interactive information visualization in neuroimaging. In Ebert, D. S. and Nicholas, C. K., editors, *Proceedings of the 1997 workshop on New paradigms in information visualization and manipulation*, pages 62–65, New York, NY, USA. ACM.
- Nielsen, F. Å. and Hansen, L. K. (1999). Modeling of BrainMap data. Online at <http://isp.imm.dtu.dk/publications/1999-/nielsen.nips99.ps.gz>.
- 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.
- 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.
- Nielsen, F. Å., Hansen, L. K., and Balslev, D. (2004). Mining for associations between text and brain activation in a functional neuroimaging database. *Neuroinformatics*, 2(4):369–380.
- Poliakov, A. V., Albright, E., Hinshaw, K. P., Corina, D. P., Ojemann, G., Martin, R. F., and Brinkley, J. F. (2005). Server-based approach to web visualization of integrated three-dimensional brain imaging data. *Journal of the American Medical Informatics Association*, 12:140–151.
- Rehm, K., Lakshminarayan, K., Frutiger, S. A., Schaper, K. A., Summers, 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.
- Reimold, M., Slifstein, M., Heinz, A., Mueller-Schauenburg, W., and Bares, R. (2005). Effect of spatial smoothing on *t*-maps: arguments for going back from *t*-maps to masked contrast images. *Journal of Cerebral Blood Flow & Metabolism*, 26(6):751–759.
- Rorden, C. and Brett, M. (2000). Stereotaxic display of brain lesions. *Behavioural Neurology*, 12(4):191–200.
- Sackett, D. L. (2001). Why randomized controlled trials fail but needn't: 2. failure to employ physiological statistics, or the only formula a clinician-trialist is ever likely to need (or understand). *Canadian Medical Association Journal*, 165(9):1226–1237.
- Salimi-Khorshidi, G., Smith, S. M., Keltner, J. R., Wager, T. D., and Nichols, T. E. (2009). Meta-analysis of neuroimaging data: A comparison of image-based and coordinate-based pooling of studies. *NeuroImage*, 45:810–823.
- 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.
- 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(3 part 1):765–780.
- Van Essen, D. C. and Dierker, D. L. (2007). Surface-based and probabilistic atlases of primate cerebral cortex. *Neuron*, 56:209–225.
- Vul, E., Harris, C., Winkielman, P., and Pashler, H. (2009). Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspectives on Psychological Science*. In press. The paper formerly known as “Voodoo Correlations in Social Neuroscience”.
- Wager, T. D., Lindquist, M. A., Nichols, T. E., Kober, H., and Snellenberg, J. X. V. (2009). Evaluating the consistency and specificity of neuroimaging data using meta-analysis. *NeuroImage*, 1(1, supplement 1):S210–S221.

Zilles, K. and Palomero-Gallagher, N. (2001). Cyto-, myelo-, and receptor architectonics of the human parietal cortex. *NeuroImage*, 14(1, part 2):S8-S20.