Published in final edited form as:

IEEE Trans Vis Comput Graph. 2013 December ; 19(12): 2916–2925. doi:10.1109/TVCG.2013.161.

GRACE: A Visual Comparison Framework for Integrated Spatial and Non-Spatial Geriatric Data

Adrian Maries,

Department of Computer Science, University of Pittsburgh

Nathan Mays,

Department of Mathematics, Wheeling Jesuit University

Megan Olson Hunt,

Department of Biostatistics, Graduate School of Public Health, University of Pittsburgh

Kim F. Wong,

Center for Simulation and Modeling, University of Pittsburgh

William Layton,

Department of Mathematics, University of Pittsburgh

Robert Boudreau,

Department of Epidemiology, Graduate School of Public Health, University of Pittsburgh

Caterina Rosano, and

Department of Epidemiology, Graduate School of Public Health, University of Pittsburgh

G. Elisabeta Marai [Member IEEE]

Department of Computer Science, University of Pittsburgh

Adrian Maries: amaries@cs.pitt.edu; Nathan Mays: nmays@wju.edu; Megan Olson Hunt: mjo25@pitt.edu; Kim F. Wong: kimwong@pitt.edu; William Layton: wjl@pitt.edu; Robert Boudreau: boudreau@edc.pitt.edu; Caterina Rosano: rosanoc@edc.pitt.edu; G. Elisabeta Marai: marai@cs.pitt.edu

Abstract

We present the design of a novel framework for the visual integration, comparison, and exploration of correlations in spatial and non-spatial geriatric research data. These data are in general high-dimensional and span both the spatial, volumetric domain – through magnetic resonance imaging volumes – and the non-spatial domain, through variables such as age, gender, or walking speed. The visual analysis framework blends medical imaging, mathematical analysis and interactive visualization techniques, and includes the adaptation of Sparse Partial Least Squares and iterated Tikhonov Regularization algorithms to quantify potential neurology-mobility connections. A linked-view design geared specifically at interactive visual comparison integrates spatial and abstract visual representations to enable the users to effectively generate and refine hypotheses in a large, multidimensional, and fragmented space. In addition to the domain analysis and design description, we demonstrate the usefulness of this approach on two case studies. Last,

we report the lessons learned through the iterative design and evaluation of our approach, in particular those relevant to the design of comparative visualization of spatial and non-spatial data.

Index Terms

Design studies; methodology design; task and requirements analysis; integrating spatial and non-spatial data visualization; visual comparison; high-dimensional data; applications of visualization

1 Introduction

Current imaging techniques make possible the acquisition of both volumetric and non-spatial neurological and functional measurements in longitudinal public health studies of unprecedented scale, involving large numbers of subjects. For example, epidemiologists study the relationship between neurological aging and mobility impairment in hundreds of older individuals. Studying the relationship between brain and mobility can help researchers identify structural changes in the brain associated with neurological aging; for example, balance difficulty and slow gait can indicate brain structural and functional abnormalities, which can be used to predict a greater risk for dementia [18]. Once recognized, these structural changes could be targeted for interventions that can prevent the loss of independence that occurs as physical mobility declines in elderly individuals.

The datasets resulting from these studies are typically high-dimensional and fragmented, and pose manifold problems from an analysis perspective. First, the data span both the spatial, volumetric domain – through medical volume images –, and the non-spatial domain – through variables such as age, gender, or walking speed; integrating these data is a challenge. Second, there is the issue of scale: there are a large number of potential predictors of each functional measurement (such as slowing gait), including numerous brain regions and health-related measures. Third, due to the high interconnectedness of the brain structures there is a high level of collinearity among neurological measurements. Thus, the problem of quantifying the correlations in the data is ill-conditioned and sensitive to the numerical algorithms employed.

Furthermore, considering the scale and diversity in the data and the processing algorithms, the data analysis has a fundamentally exploratory nature. Researchers are interested not only in data acquisition and analysis, but also in developing, adapting or selecting the appropriate numerical analysis method, and in hypothesis generation, testing and refinement. Flexible, scalable and exploratory analysis methodologies for identifying correlations in the data thus require the development of novel infrastructures that blend medical imaging, mathematical and statistical methods, and interactive visualization.

Finally, the integration of spatial and non-spatial data – in the context of high-dimensionality and visual comparison – prompts a combination of visual encodings and techniques. The challenge lies in designing the right combination of appropriate techniques, given the application domain and the analysis needs of the expert users. These challenges are typical of other spatial and non-spatial domains, such as geospatial analysis, bioinformatics, or neuroscience applications in general. While multiple successful examples

of visual tools exist across these domains, we do not have a comprehensive understanding of the parameters that influence successful spatial and non-spatial integrated visualization, nor is it easy to determine how close we are to achieving such an understanding. Domain-specific analyses and design studies can help in this direction.

In this paper we describe the design and development of a novel visual tool for the integration, analysis, and exploration of functional geriatric research data. The tool (Fig. 1) blends medical imaging, mathematical and statistical analysis and visualization in an attempt to support Geriatric Research In Ambulatory and Cognitive Excellence (GRACE). The contributions of this work are as follows: 1) We provide a description of the problem domain and an analysis of the questions typically asked in this problem domain. 2) We adapt two methods for quantifying correlations between brain and functional patterns such as gait: Sparse Partial Least Squares regression and iterated Tikhonov Regularization. 3) We describe the design and implementation of an interactive visual comparison tool based on a mixture of visual encodings. 4) We demonstrate the usefulness of the visual tool in two case studies and 5) Report the lessons learned based on feedback from experts in the epidemiology and statistical domains.

2 Related Work

The issue of combining scientific visualization with information visualization in the context of spatial and non-spatial integration was highlighted in 2003 by Theresa-Marie Rhyne, who wonders if the difference between the two sub-fields really matters [34]. Some of the early work to integrate both spatial and non-spatial visualizations includes [10, 5, 21], which, in addition to 3D views, include mathematical abstractions of the data. More recently, bioinformatics has provided a constant stream of datasets that contain both spatial and non-spatial data to be visualized [37, 49, 26].

Although a large number of brain visualization packages exist, attempting to integrate with such packages functional data visualization, while taking into account the domain requirements was not feasible. For example, one core requirement was to allow researchers to visually explore correlations between functional and neurological variables mapped to regions at different depths inside the structure of the brain. MindSeer [27] and BrainMiner [28] rely on iso-surface rendering of the brain, which leads to occlusion, making difficult the mapping of functional data and the interactive selection of significant regions while maintaining context. Another project [22] does allow the user to see transparent images, but has limited interaction, which it tries to overcome by providing "magic mirrors", three projections that attempt to unambiguously reveal the locations of regions inside the brain. Another important requirement was to enable epidemiologists to see how correlations vary across groups of regions that are functionally similar (for example involved in mobility or visuospatial attention). While 3D Slicer [33], BrainVoyager [9] and a system built for visualization of MS lesions [45] allow interactive direct volume rendering, they do not provide a way to encode and visualize the functional structure of the brain. Finally, LONI [4] provides a brain ontology viewer; however, the viewer is not linked to the rest of the visualization.

Many techniques have been proposed (and explored in this project) for the visualization of hierarchies; among them are traditional node-edge displays such as DOITrees [13] and space-filling techniques: treemaps [16, 40], radial layouts [19] or circle-packing techniques [47]. Dendrograms are a type of node-edge display and have been used for clustering genes [38] and brain fiber tracts [15]. Starplots are a common approach for visualizing of multidimensional data [20, 29].

Visual comparison has drawn significant attention in recent years. Gleicher et al. [7] identify three fundamental visual paradigms: juxtaposition, overlays, and abstraction; the latter category encompasses a large number of possible difference abstractions. Earlier work by Taylor [42] examines the benefits of visualizing multiple fields in a layered approach (similar to Gleicher’s overlay paradigm) by using difference visual channels such as color, transparency, texture etc. He found, however, that visualization began to fail when “more than a few techniques were applied at once”; and also noted on the importance of spatial frequency. While perceptual studies on multilayered representations exist in information visualization [14], little is known about visual comparison in the context of integrated spatial and non-spatial data. In our approach, we pursue a linked-view hybrid paradigm which blends juxtaposition, volume rendering transparency, and abstractions of non-spatial data.

Finally, our work builds upon uncertainty research in the visual perception field. MacEachren et al [23] identify nine possible types of information uncertainty, ranging from accuracy/error and precision to credibility and subjectivity. While most spatial uncertainty visualization techniques focus on encoding strictly accuracy or error [17, 32, 2], our focus is on encoding the credibility of a particular relationship.

3 Methods

Our first contribution is an analysis of the GRACE problem domain, its data, and its tasks. We then introduce our tool framework, which blends medical imaging, statistical analysis, and visual encoding (Fig. 2). We briefly describe the data normalization procedure, which we perform to enable the computation of variable correlations. Next, we describe the adaptation of two numerical algorithms for computing correlations. We then describe the design of the analysis tool; the design is informed by the domain analysis, following the principles outlined by Munzner in [30].

3.1 Data and Task Abstraction

3.1.1 Data Analysis—Neuroimaging datasets may include hundreds of subjects, with more than 50 neurological dimensions (brain volumetric measurements) and multiple functional variables (such as gait speed, balance etc.) for each subject. Voxels in a volume image may be grouped together in “regions of interest”, to facilitate cross-subject comparison. A step further, individual brain regions may be grouped together based on similarities in their function in the brain. For example, some regions are involved in mobility, others in visuospatial attention and others in motor imagery.

Furthermore, there are a large number of potential predictors of functional characteristics, including numerous brain regions and health-related measures. In general, the number of

observations is lower than the number of brain regions of interest. Even when applying a selected number of regions based on *a priori* hypotheses, the number of variables could include more than 30 regions of interest for each hemisphere, which yields a total of over 60 comparisons. The interactions between functional variables and the different types of brain measurements give rise, for each brain region, to 5 to 10 correlational values computed with varying degrees of confidence. Brain regions are likely to be correlated due to the inherent brain interconnectedness. Analysis of neuroimaging correlates of behavioral characteristics typically involve a very high number of comparisons and require conservative methods to correct for false positive results.

Finally, the results of statistical analyses are typically sensitive to scaling, selection, and normalization of variables for statistical interpretations. Multiple algorithms may be used to compute the correlations, and researchers are typically interested in comparing the different algorithmic outputs. A particular algorithm's confidence in the strength of a hypothesized correlation may be encoded as a *p*-value. The regression coefficients computed by a particular algorithm may also be of interest.

To summarize, the domain data consist of both spatial and non-spatial multivariate measurements; to these we add computed data, which may result in several other attributes. The datasets feature high-dimensionality, in the number of subjects considered, in the number of variables, and in the number of possible correlations and comparisons. Finally, spatially-related correlations may be computed with varying degrees of confidence through a variety of algorithms, further emphasizing the need for visual comparison.

Datasets: The data we use in our case studies come from two sources. The first dataset is used to study the relationship between brain and gait in 324 community-dwelling older adults [36]. The gait data was recorded on a 4-meter long instrumented walking surface, the GaitMat II [1], and included gait measurements such as gait speed, stride length, base of support, double support time, and latency. The gait data was further summarized as three variables: bradykinesia, gait disturbance and tremor. Bradykinesia is a measure of slow or hesitant motion during activities such as heel tapping. Gait disturbance is a measure of abnormality in gait (e.g., small amplitude or poverty of movement in general), posture, etc. Tremor is rated as present if detected, either at rest or during action in either foot.

The neurological data for this dataset was acquired through Magnetic Resonance Imaging (MRI) of the subject brains. The resulting volume images were an assessment of the macro- and micro-structure of the brain. The macro-structure was measured as volume of gray matter and of white matter tract hyperintensities (lesions in white matter). Diffusion Tensor MRI (DTI) was used to quantify the microstructure of normal-appearing gray (mean diffusivity) and white (fractional anisotropy) matter.

A second dataset is used to validate several algorithms that are used to compute associations between dependent and independent predictors. The dataset was used to verify known results in epidemiology literature, namely the effect of certain regions in the brain on the subjects' performance on the Modified Mini-Mental State Examination (3MSE) and the Digit Symbol Substitution Test (DSST). The former is a brief, general cognitive test which

evaluates the subject's orientation, concentration, language, praxis and immediate and delayed memory [43]. The latter is a pencil-and-paper test designed to examine the subject's psychomotor performance. The subject is presented with a series of digit-symbol pairs and then required to recall the symbols associated with a sequence of (repeated) digits. The score on the test consists of the number of pairs the subject correctly identifies during a 90-second span [24].

The data was obtained from participants in two on-going longitudinal studies, the Cardiovascular Health Study [6], and the Healthy Brain Project, which is a sub-study of the larger Health, Aging & Body Composition Study, [41]. In both studies, the neurological data was acquired in the same manner as described above and collected from 324 subjects in the former study and 314 in the latter.

3.1.2 Task Analysis—Through repeated on-site interviews and job shadowing, we abstracted the list of tasks most commonly performed when analyzing brain/functional task data shown in Table 1. The second column states the task abstraction, while the third column shows typical instantiations revealed through user interviews.

Analyzing the domain tasks abstracted (Table 1), we find first that many of the tasks involve pairwise comparisons, some between spatial/non-spatial variables, and some between spatial/non-spatial outputs of various correlation algorithms (Tasks T1–T3, T5–T6). Furthermore, the comparisons are almost always comparisons of correlations, normalized over many human subjects, and are thus in a certain sense “2nd order” comparisons. Finally, while the comparisons almost always involve spatial patterns, some domain experts appear to initiate their analysis with non-spatial variables, while others start their analysis with spatial variables.

As anticipated, many of the tasks rely on detection of spatial patterns, suggesting the need for direct spatial manipulation such as rotation and slicing. Most of the questions asked during the analysis are, in fact, either of the *How* or *Where* variety; this observation suggests a need for either linking or overlaying the various types of information. Some of the tasks (e.g., T7) require interactive filtering and conditioning.

Additional user requirements revealed include the ability to explore regions mapped at different depths inside the spatial structure; spatial-context visibility; ease of use, low visual-complexity, flexibility, and cross-platform portability of the resulting tool.

3.2 Spatial Data Normalization

To enable the cross-population analysis of correlations, the MRI volume images for each subject are first divided into anatomically-defined regions of interest (ROI). The ROIs have been previously drawn on a template brain according to the automated anatomical labeling (AAL) neuroanatomical atlas [46]. After skull and scalp stripping, and segmentation of gray matter, white matter, and cerebrospinal fluid, the brain atlas and the brain of each subject were aligned. Intensity normalization was done on each individual's volume image as well as on the brain template. This gives each individual the same orientation and image-intensity distribution as the template and improves the registration accuracy. The registration

procedure used a fully deformable, non-rigid registration algorithm [44] that does not warp or stretch the individual brain and thus minimizes measurement inaccuracies. It also allows for a high degree of spatial deformation compared to other standard registration packages. Each ROI was summarized in terms of its voxel volume, for each imaging modality.

3.3 Correlation Quantification

When studying multiple brain region volumes as predictors of an outcome, the regions are likely to be correlated due to the interconnected nature of the brain. Another problem is that there may be a large number of potential predictors, including the many brain regions and demographic confounders. Additionally, brain imaging techniques are often costly and thus the number of observations tends to be low relative to the number of regions of interest.

Various approaches have been used for the type of problem considered, including PCA, PLS, Sparse PLS (SPLS), machine learning techniques (Mutual Information, Independent Component Analysis, Local Linear Embedding, IsoMap), and Tikhonov regularization (ridge regression.) Of these, SPLS and Tikhonov are the most popular for data selection; both address the risk of over fitting (the number of parameters is larger than the sample size problem), and possible colinearity and its resulting magnification of noise/experimental error. For this reason, in this study, SPLS and Tikhonov regularization were of particular interest to the domain experts.

Sparse Partial Least Squares Regression—Sparse PLS has been introduced to add variable selection to the classic PLS regression [3]. The variable selection is accomplished by imposing sparsity in the middle of the dimension-reducing step. As a result, the approach simultaneously reduces the dimensionality of the data and selects a subset of predictors.

In our adaptation of SPLS, the matrices were initialized with the subject's spatial and abstract measurements: brain regions specified as the volume of gray matter, and numerical equivalents for age, obesity, race, gender and total gray matter volume. Models were fit for many values of the tuning parameters, and the percentage of times a variable was chosen was used to quantify the strength of a particular association. Furthermore, instead of a p -value, we implemented bootstrapping (drawing random samples with replacement) to create a 95% confidence interval for the SPLS parameter estimates. From this, a variable was deemed significant if its confidence interval did not include zero. The SPLS analysis was performed in SAS and R using custom code.

Iterative Tikhonov Regularization—The workhorse and most commonly used method for the accurate and reliable solution of ill-posed problems (in which solutions are either over- or under-determined and small errors in data are magnified greatly in the solution) is Tikhonov Regularization. In this method, a regularization parameter $\alpha > 0$ is selected and an approximation to the solution of the system $\mathbf{Ax} = \mathbf{b}$ is computed by solving

$$(\mathbf{A}^t \mathbf{A} + \alpha \mathbf{I}) \mathbf{x} = \mathbf{A}^t \mathbf{b} \quad (1)$$

where $\mathbf{I}$ is the identity matrix, $\mathbf{b}$ is the known outcome matrix, and $\mathbf{A}$ is the predictor matrix.

In our application of the Tikhonov algorithm, the MRI brain data as well as other correlates to non-spatial quantities (such as age, gender, race) are contained within the matrix A , where the columns are these variables and the rows are the patients involved in the study. The observable variables are contained within the matrix b , where each row corresponds to a patient and the columns correspond to the measured characteristic (e.g., gait speed.)

We extend the Tikhonov algorithm by combining an iterative version of it with an L-curve method [11] to find the optimal regularization parameter α . To provide an estimate of the p -values for the regression coefficients, we take a novel approach to examining the sensitivity of each coefficient to random noise. We perform this step by bootstrapping (random selection with replacement) of scaled residuals ε onto the initially solved regression vector $\mathbf{x}$ to generate a *pseudo* b vector

$$b_{\text{pseudo}} = A\mathbf{x} + \varepsilon \quad (2)$$

from which a new regression vector $\mathbf{x}_{\text{pseudo}}$ can be solved using again the iterated Tikhonov algorithm. After T such bootstrapping processes, we sort the collection of T coefficients for each regression variable and we estimate the p -value for the i -th variable using

$$p(i) = \begin{cases} 2k/T & \text{when } k < T/2 \\ 2(k-T/2)/T & \text{otherwise} \end{cases} \quad (3)$$

where k is the index where the coefficient changes sign. The iterated Tikhonov results are obtained from custom code written in Matlab and in Fortran 90.

While for this particular domain we have adapted two correlation algorithms – Sparse Partial Least Squares and Iterative Tikhonov regression – other correlation algorithms can be swapped for these two, interchangeably.

3.4 Visual Design and Encodings

After data normalization and correlation computation, the core phase of this project was deciding on the overall design for visualization, data graphical representations, visual encodings, and interactions to support the domain data, tasks and requirements summarized earlier.

Several design approaches to visual comparison are known, including juxtaposition (side-by-side views), overlays and explicit encodings of differences [7]. Considering the collaborative nature of the domain, in which users with complementary expertise contribute different viewpoints —such as brain function versus brain structure, our top design pursues a hybrid approach augmented by a linked multiview paradigm. We posit that the linked multi-view approach would help generate insights [35] by allowing users who are trained to start their analysis with either spatial or non-spatial features to explore the data from both viewpoints.

Based on the data and task analysis, the application layout consists of five linked views containing spatial and non-spatial visual representations of the data: two volume rendering panels, two corresponding dendrogram panels and a Kiviat diagram panel (Fig. 1). Kiviat diagrams are also known as Radar Charts, and are a variation of a starplot representation. Additional widgets allow data selection, filtering and manipulation. Correlations are encoded across all panels: the volume renderings and dendrograms show the p-values of correlations between spatial and non-spatial measurements, while the Kiviat diagram can also show the absolute value of coefficients computed for those correlations. Finally, the multiple views are connected through interaction, including dynamic queries and brushing and linking. This flexible top-design can be used to support both pairwise comparison of cohorts or variables, and pairwise comparison of algorithm outputs. Below we describe each component in detail.

3.4.1 Spatial Data Encoding for Comparison

The two volume rendering panels were designed to allow detection of spatial contiguity in brain regions of interests (Tasks T1, T4, T5; and indirectly tasks T2, T6, T7). For these spatial encodings, we chose hardware-accelerated direct volume rendering of the ROI brain template, as opposed to iso-surface rendering, in order to include filtered-out brain regions and thus help maintain spatial context. Filtered-out regions are rendered transparently. Given the richness of the data and to avoid unwanted visual complexity, for this type of comparison tasks we pursued a juxtaposed layout.

Correlations between a particular non-spatial variable and the spatial brain regions of interest are color-mapped to the two volume renderings based on ColorBrewer.org [12] schemes, against the dark background familiar to the geriatric researchers. We selected two sequential multihue color schemes (green-blue for gray matter and orange-red for white matter), to ensure contrast between spatial variable categories. Both schemes were modified slightly to increase the contrast, especially towards the higher end (blue and red). Considering the expert analysis needs with respect to placing brain regions into categories of correlation strengths (rather than inferring exact values), we map colors to five p-value ranges. The ranges can be interactively updated. Following Sanyal et al. [39], high p-values (i.e., low credibility of a correlation) are mapped to low saturation and high value (S and V in the HSV color space). At the request of the domain experts, we also include a rainbow color scheme. We note that alternative “credibility” encodings such as transparency, fuzziness, animation or glyphs would conflict in this particular domain with satisfying the user requirements regarding context visibility and ease of use.

Users can interact with the brain model using either standard direct manipulation or rotation widgets, for stricter control. A slicing widget is also provided to allow exploration of regions which are in the central part of the brain and are obscured by cortical regions. To facilitate spatial comparison, the two volume renderings are synchronized during manipulation. To further preserve context during manipulation, a small axis-aligned coordinate system is shown in the corner of each rendering. Finally, on-demand pop-up labels specify the names of regions of interest.

3.4.2 Non-Spatial Data Encoding for Comparison

Non-spatial information such as brain region function was equally important to the domain experts (Tasks T3, T5, T6). To support comparison of this rich information, we selected a hybrid (juxtaposed and overlay) layout. The dendrogram panels and Kiviat diagram are abstractions of hierarchical information about neurological function, respectively an abstraction of correlation strength.

Two side-by-side dendrogram panels enable the user to compare functional, hierarchical information for different groups of regions, for both white matter and gray matter regions. The dendrogram encoding was selected to represent the functional structure of the brain from a variety of other prototypes, winning as a simple, traditional technique for displaying hierarchies and, consequently, easier to understand and more familiar to the domain experts. The leaves of the dendrogram encoding represent individual regions while their parents correspond to groups of regions associated with particular functions. For example, some regions are involved in mobility, others in visuospatial attention and yet others in motor imagery (Fig. 3.)

Correlations are color-mapped to the dendrograms using the same color scheme as for the volume renderings, further emphasizing the strength of a computed correlation. Similarly, linking and brushing are used to gray out the regions filtered out by the user. Mnemonic labels are provided for groups of regions, while on-demand pop-up labels specify the full name of each region (Fig. 3.) We note that the linear dendrogram layout, as well as the full replica of the tree structure in the two panels were design choices informed by the domain expert preference: prototype circular layouts, respectively reducing or reverting the bottom dendrogram collided with the experts' understanding of a hierarchical structure.

The Kiviat diagram view panel further enables the pairwise comparison of correlation values across regions for different spatial and non-spatial variables (Task T6); and can be used alternatively as an overview. Per the request of the domain experts, the Kiviat diagram can also be used to visualize regression coefficients, when available. The experts emphasized, nevertheless, that the primary information they pursued was related to p-values, and that the regression coefficients were secondary information. This encoding follows an overlay approach: during the prototyping phase it became clear that the Kiviat diagram representation was the user-preferred option for the overview visualization of correlations. Alternative representations evaluated by the domain experts during parallel prototyping were regular star plots, stacked displays, mosaic plots and Chernoff faces, parallel coordinate plots and linked histograms.

In this encoding, the Kiviat diagram axes correspond to brain regions, and the value mapped to each axis represents the significance level. Correlation or regression coefficients are mapped according to their absolute values, per the domain expert request. Radio buttons allow the user to filter out either positive or negative correlation/regression coefficients. The color-mapping for the Kiviat diagram is distinct from the other panels, in order to emphasize the information separation between the regression coefficients (a byproduct of the numerical algorithms) and the rest of the data.

3.4.3 Linked-Views and Interaction

The view panels are linked through color-coding and through interaction, to facilitate the integration of the various types of information (all tasks). Several widgets (Fig. 1) allow data selection (datasets, variables, and algorithms), filtering (data types, correlations and significance values) and data manipulation. Data selection and filtering are facilitated through combo-boxes, check-boxes, and radio-buttons (Fig. 1), and correspond to the T1-T7 tasks outlined by the domain analysis. Direct data manipulation (Tasks T1, T4, T6) is further facilitated by rotation and slicing widgets.

Dynamic queries through the widgets are employed to enable the user to filter the correlations to the desired significance level (Tasks T1-T7). Linking and brushing are used consistently to filter out regions of interest across the multiple views, as well as to remove the significance level mapping for the regions not displayed in the two volume rendering panels. Brushing and linking are also used to visually connect the separate panels, including the dendrograms and the Kiviat diagram. Finally, correlations are encoded using the same color-map across the separate panels.

The interface implementation uses custom code written in C++ with OpenGL and QT, including a hardware-accelerated module for volume rendering [25].

4 Evaluation

We demonstrate the usefulness of our application on two case studies. Feedback on the design was acquired through repeated validation and evaluation sessions with our target users, some of whom are co-authors of this paper: three epidemiologists, two neurologists, two mathematicians, one statistician and two computational scientists; as well as three other research groups located off site. Some of this feedback is already reported in earlier sections. Three of our end-users (co-authors CR, MOH and KFW) provide the following two detailed case studies, while the larger group feedback is reported in Section 4.3.

For the two case studies, a think-aloud protocol and targeted questions were used. The first case study examines neurological-mobility correlations: we compare two different non-spatial gait variables, bradykinesia and gait disturbance. The second case study is a comparison of algorithms used to compute associations in neurological-functional data.

4.1 Brain-Gait Correlation Analysis

The lead GRACE researcher studies the correlations between neurological aging and mobility impairment in the elderly. She had previously discovered an association between specific gait variables and neurological data using a partial least squares algorithm and spreadsheet analysis and was particularly interested in visualization as a means to explore more complex and subtle spatial patterns in the data. In particular, she wanted to know how two gait variables (bradykinesia and gait disturbance) were similar or different in terms of the brain regions correlated with each. The researcher's team is well-versed in state of the art brain visualization software, and was unable to perform this type of analysis using such packages.

Correlations among the data were computed using SPLS and loaded into the visual comparison tool. Once some of the more significant p -value categories were selected, the researcher noticed that the brain regions correlated with bradykinesia were distributed more cortically than those for gait disturbance (volume rendering, Fig. 4). However, she found the fact that only two of the mobility regions and only one motor imagery region associated with gait disturbance were significant at the 5% level suspicious. She decided that this had to be investigated further.

The next step was to switch to correlations between gait variables and microstructure brain measurements (Fig. 1), where the senior researcher immediately noticed that the two mobility regions associated with gait disturbance were neuro-anatomically contiguous. Two significant mobility regions associated with gait disturbance (thalamus and caudate, observed in the lower dendrogram, **Mobility** branch), do not appear to correlate with bradykinesia (upper dendrogram). The Kiviat diagram further emphasizes the coverage difference between the two variables. However, the shape of the gait disturbance Kiviat diagram (purple, two protruding arms) indicates the correlation algorithm is confident about the association. Adding this to the expert's prior knowledge about the position of the two regions in the basal ganglia, the researcher concluded the association could indeed be a biological possibility. The researcher stated that people with severe damage in the basal ganglia (such as a result of Parkinson's disease) also have severe gait disturbances. Even though participants in the study from which the data was collected are "normal" individuals, less severe damage in the same area could be correlated with mild gait disturbance.

When prompted about usage of the individual visual components in her analysis, the researcher reported that she used the dendrogram panels to easily identify the extent of overlap between two mobility measures. The volume renderings further facilitated interpretation of neuroanatomical contiguity, sparing the researcher the mental computation to figure out which regions are located next to others. Being able to interactively manipulate the volume renderings was important, and so was brushing and linking of the complementary visual representations. The Kiviat diagram was thought to have two benefits: when displaying regression coefficients, 1) the Kiviat diagram would give the user an idea about the strength of association, and 2) it would reinforce the idea of overlap. Overall, the researcher was using the combined visual encodings to identify whether specific associations might be spurious or not. The complementary visual encodings helped her decide if the finding was worth pursuing or not. A finding such as the example above would strengthen the belief that the finding is worth pursuing. The lead researcher also stressed that she would often use the visualization as an exploratory tool, where there is large, highly-dimensional data to sift through. This type of exploration was particularly cumbersome in existing brain or statistics visualization packages. This target user has adopted early the prototype as a screening research tool.

4.2 Correlation Algorithm Analysis

Researchers in the School of Public Health are interested in either adopting or developing numerical algorithms capable of capturing complex relationships in the mobility-neurology

data. This second case study focuses on regression algorithm comparison: SPLS and iterative Tikhonov regularization.

Given the exploratory nature of the brain-gait data, which features few *a priori* known associations, two junior researchers in the group decided to test the two algorithms instead on the DSST and 3MSE datasets. In each of these datasets the brain regions were used, along with a few commonly examined demographics to predict two different outcomes, DSST and 3MSE, for two different patient cohorts. Specifically, the researchers aimed to verify that the two algorithms confirmed the following hypotheses: the dorsolateral prefrontal and the lateral superior and inferior parietal regions are predictive of DSST, while the hippocampus, parahippocampus and entorhinal cortex are associated with 3MSE. In addition, the researchers wanted to further explore the datasets for potential associations that might be worth pursuing. For example, if non-hypothesized regions had low *p*-values consistently over different algorithms and/or cohorts, the researchers may study those new regions in the future.

Exploring the dataset, the two researchers noticed that the results generated using the two algorithms generally agree on the hypothesized regions – the regions tend to appear in most outcome/cohort combinations. One thing the visualization tool allowed them to do was zero-in on the one combination where the results disagreed. As emphasized by all panels in Fig. 5, none of the 10 regions which are significant at the 5% level overlap. After further exploration, the researchers agreed that, at least in this specific outcome/dataset combination, the two algorithms were sensitive to different aspects of the data, with SPLS being stricter in variable selection than iterative Tikhonov.

Both researchers were enthusiastic about the tool. They stated that having different ways (abstract and spatial) of seeing the data really helped, and particularly emphasized the value of the interactive, linked volume renderings (“when you rotate the brain, you feel like you can almost touch it”). They also commented that the visualization tool could be very useful for epidemiologists, medical school students and policy-makers of the health care system who are not very familiar with the brain structure and function, as using the tool would improve their understanding of the brain regions locations and functions.

4.3 Domain Expert Feedback

The domain expert feedback from repeat evaluation meetings showed remarkable enthusiasm for the framework and its implementation in a tool; specific comments included “elegant”, “GRACEful”, “scales appropriately”, “simple enough to use that we could do an epidemiologic study”. The most important aspect of the tool was its ability to support the domain tasks and data, and the overall user workflow. The ability to seamlessly integrate spatial and non-spatial information was particularly appreciated, as well as the ability to interactively filter regions based on *p*-values; also the ability to interact directly with the 3D data.

From the individual components of the visual interface, the mapped volume renderings with the *p*-value filtering and dynamic queries were most eagerly embraced by the multiple research groups, followed by the dendrogram representations. The Kiviat diagram was

initially unfamiliar to the experts, and required most training; it took a few rounds of demonstrations until the experts endorsed its enhancing effect on exploratory analyses, and became enthusiastic about its usage. Visual support from the familiar volume renderings and hierarchical representations was essential in this training period.

Some of the researchers were well versed in functional MRI tools such as FSL, SPM, and BrainVoyager, as well as more general packages like Slicer/VTK or AFNI. While they agreed that they would use these particular tools “if under duress” and for explanatory purposes, they also stated that almost without exception the existing tools were appropriate for individual-subject analyses, and did not support their research workflow: “nice looking but difficult to use in a statistical setting” and sometimes simply “not good”. Tools that link machine learning or statistics with spatial data (FSL, AFNI, R) were particularly appreciated, but not adopted for research on account of their lack of integration of tractography, and sometimes spatial data altogether. To complete analyses of a single non-spatial variable correlation with spatial data, the researchers would typically complete the statistical analysis in SPSS, and then use their own mental model of neurological data to interpret the resulting tables. Visualization tools would be used only in post-processing, to generate images for explanatory purposes (publications, training, and presentations.) Complex, flexible, exploratory analyses like the ones reported in the case studies in this paper — which involve comparisons and correlations among more than one non-spatial variable at a time — would be, in the experts’ opinion, nearly impossible in the absence of GRACE. Several of the domain experts have adopted the tool for research purposes.

5 Discussion

The two case studies and the expert feedback indicate that the visual analysis tool is of significant help in the comparative exploration of correlations between high-dimensional spatial and non-spatial geriatric data. While the individual visual encodings are not novel, the exploration of this relatively large design space and the combination of visual encodings in a tool to handle spatial and non-spatial information in geriatrics is novel. The chosen spatial and non-spatial visual encodings have shown complementary strengths: interactive volume rendering was particularly useful for detecting spatial patterns, while the dendrograms and Kiviat diagrams captured both correlation overlap and similarities among correlation algorithms. The use of linked-views and interaction, in the context of comparison tasks, to navigate uncertainty in this hypothesis space is also novel. When connected through interactive filtering, these multiple views on the data provided insight into domain problems, steered the investigation, and allowed for the generation of new hypotheses – a holy grail of visual analysis in general. A measure of the success of our software is the adoption of the tool for research purposes by our epidemiologist co-authors and their research labs. The senior epidemiologist refers to the resulting tool as “*incredibly close to [her original] dream*”.

Spatial and non-spatial visual integration challenges are common across domains, from geospatial applications to bioinformatics to functional neurology applications. We believe the lessons learned in terms of linked-views and successful spatial and non-spatial visual encodings from this specific design study have application to other domains, as well.

The most important lesson that emerged is that users with different backgrounds may employ a spatial and non-spatial visualization application differently. This may be particularly important for the design of collaborative meeting environments, in which users contribute complementary expertise; in the words of a domain expert, “*everyone has holes in their knowledge*”. In our case studies, epidemiologists who are not very familiar with the spatial structure of the brain tended to start their analysis with the dendrogram panels – to see if they can detect a pattern. Only then would they look at the volume rendering panels to see if the spatial layout might give them any additional information. In contrast, neurologists tended to start their analysis directly with the spatial volume rendering, and use the non-spatial abstractions to further strengthen or refute their observations.

This observation supports a multi-view design of visual tools that seek to integrate spatial and non-spatial information; while the information would thus be somewhat fragmented, our evidence points to the benefits of separating the views. In view of the Baldonado et al [48] guidelines, the disadvantages of context-switching between the views were clearly outweighed by the user familiarity with a particular type of analysis (spatial or non-spatial). Similar designs may benefit other spatial and non-spatial application domains. For example, bioinformatics also features complementary user-expertise; this expertise covers 3D spatial protein structure and non-spatial information such as gene sequencing. Simplifying the data available in any one view has further advantages, by allowing a stronger emphasis on similarities and discrepancies in the data. Finally, linking the views allows the user to harness and expand their previous analysis experience.

We credit much of the success of the tool to its explicit “comparative” design, fueled by the domain analysis. The theory of design for visual comparison, in particular in the context of spatial and non-spatial data, is largely unexplored. We note that Globus argued in 1994 that the question “Compared to what?” lies at the heart of quantitative reasoning; he advocated the use of comparative, rather than descriptive visualizations [8]. Munzner cites previous research to conclude that side-by-side visual comparisons outperform both animations and zooming [31]. Gleicher et al. further showed that visual comparisons is typically performed through juxtaposition, superposition, explicit encoding or a combination of any two of them [7]. Nevertheless, there are very few guidelines to visual comparison design. In most instances in our design we favored juxtaposed (side-by-side) layouts, in order to avoid the clutter issues that come with superposition. One exception is the Kiviat diagram panel, where the lack of a physical structure supports superposition. As in [7], we found overall that a hybrid approach supported best the tasks revealed by the domain analysis.

Furthermore, our domain analysis highlights the wide range of tasks that “comparison” actually may cover, from analysis of similarities to detecting patterns. Our experience indicates that linked multi-view designs can successfully address this range of comparison tasks. In both case studies, the juxtaposed spatial views were found very useful in detecting and comparing spatial patterns; the spatial views seem to be more useful the more expert the user is in brain neurology (spatially-oriented). In contrast, the task of comparing subgroups of regions based on the functional structure of the brain seemed to be easily solved by the non-spatial panels. Finally, the combination of Kiviat diagram and volume rendering panels was helpful in comparing different variables in the datasets: the Kiviat diagram panel

worked as an overview, where users could see and compare in detail p -values for all selected regions for both selected variables; the spatial dimension views then strengthened or weakened the expert's belief in a particular hypothesis. The layout of our comparative panels was guided by the user preference. Nevertheless, the question of recommended juxtaposed layout arrangements remains open.

In terms of combining scientific visualization paradigms with information visualization paradigms, we note that while our target users were well versed in using spatial encodings though brain imaging software – in particular surface and volume renderings to present data and findings – they posed initially significant resistance to non-spatial abstractions. It was encouraging to see that persistence helped and, in the longer run, that these visual abstractions were found to be intuitive and rather essential in understanding the spatial data. A similar observation follows the encoding of correlation strength through linked hue and saturation in both the spatial and abstract representations, despite the researchers' being used to qualitative (rainbow) color mappings; however, for presentation and publication of results the target users still prefer qualitative colormaps.

Finally, we note that while current visual perception work tends to examine particular iconic encodings of uncertainty, integrating both data and data uncertainty representations in the same visual encoding is still an open question. We empirically found that saturation and value mapped to volumetric data can encode successfully the credibility of a correlation. A step further, scientific visualization in general would benefit from a better understanding of specific uncertainty components – beyond accuracy – and their relationships to specific application domains, to expert users and to their analysis needs. In particular, limited attention has been directed to the use of interactive techniques such as brushing, linking, dynamic queries and conditioning in the context of uncertainty exploration. Our tools points to the benefits of matching such interaction types to spatial and non-spatial correlation analysis tasks, yet few guidelines exist for interactive visualization in general.

6 Conclusion

In conclusion, we have designed, developed and validated a novel framework for the comparative visual analysis and exploration of spatial and non-spatial correlations in functional geriatric research data. In this process, we have generated an analysis of the application domain – with emphasis on spatial and non-spatial information integration and comparison tasks. Our domain task analysis may be built upon across application domains to extract a consistent taxonomy of tasks typical of spatial and non-spatial integration.

Furthermore, we have created a novel infrastructure which blends medical imaging, mathematical and statistical methods, and interactive visualization. We also adapted and extended two algorithms – Sparse Partial Least Squares and iterated Tikhonov Regularization – to quantify correlation strengths in neurology-functional data.

Our design study points to the benefits of linked-views in the context of spatial and non-spatial integration, in particular when domain users have complementary expertise. The visual comparison tool successfully connects highly-dimensional spatial and non-spatial geriatric information through a novel, interactive, linked-view hybrid design. The separation

into multi-views allows users to gradually build up confidence in the tool; despite initial resistance to some non-spatial representations, the approach may make the learning-curve less steep and facilitate early adoption of a tool. This flexible design supports the information needs of the domain users, while maintaining the user-required low visual-complexity.

As far as we know, this is the first evaluation of the benefits of linked views for visual comparison in the context of spatial and non-spatial information. Our approach allows experts to use any view to form a base, exploratory belief, and the other views to strengthen or refute that initial belief. The lessons reported though this design study may inform the design of visual tools across other domains which seek to integrate spatial and non-spatial information, including geospatial analysis, bioinformatics, and neuroimaging. Evaluation on two geriatry case studies shows that the resulting interface enables the users to effectively generate and refine hypotheses in a large, multidimensional, and fragmented space.

Acknowledgments

This work was supported by NIH R01-AG029232, NSF IIS-0952720, and by a University of Pittsburgh Small Multidisciplinary Grant. Grateful acknowledgments to our Geriatric Psychiatry, Epidemiology, Physical Therapy and Computational Science senior collaborators on this project: Howard Aizenstein (University of Pittsburgh), Levent Yilmaz (University of Pittsburgh), Vijay Venkatraman (University of Pittsburgh), Patrick Sparto (University of Pittsburgh), Anne B. Newman (University of Pittsburgh), Fridtjof Thomas (University of Texas, San Antonio), Kristine Yaffe (University of California, San Francisco), Lenore J. Launer (National Institute of Aging, NIH), and to their respective research groups for many insightful conversations and for their repeated feedback. Thanks to Sriranjani Mandayam and John Yackovich for an early volume rendering prototype, and to the Pitt Vislab for feedback.

References

1. Brach JS, Berthold R, Craik R, VanSwearingen JM, Newman AB. Gait variability in community-dwelling older adults. *Journal of the American Geriatrics Society*. 2001; 49(12):1646–1650. [PubMed: 11843998]
2. Bui, A.; Hsu, W. *Medical Imaging Informatics*. Chapter 4: Medical Data Visualization: Toward Integrated Clinical Workstations.
3. Chun H, Keles S. Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *Journal of the Royal Statistical Society Series B Statistical methodology*. 2010; 72(1):3–25.
4. Dinov ID, Valentino D, Shin BC, Konstantinidis F, Hu G, MacKenzie-Graham A, Lee EF, Shattuck D, Ma J, Schwartz C, et al. Loni visualization environment. *Journal of digital imaging the official journal of the Society for Computer Applications in Radiology*. 2006; 19(2):148–158.
5. Doleisch H, Hauser H. Smooth brushing for focus+context visualization of simulation data in 3d. *Journal of WSCG*. 2002; 10(1):147–154.
6. Fried LP, Borhani NO, Enright P, Furberg CD, Gardin JM, Kronmal RA, Kuller LH, Manolio TA, Mittelmark MB, Newman A. The cardiovascular health study: design and rationale. *Annals of Epidemiology*. 1991; 1(3):263–276. [PubMed: 1669507]
7. Gleicher M, Albers D, Walker R, Jusufi I, Hansen CD, Roberts JC. Visual comparison for information visualization. *Information Visualization*. 2011:1–28.
8. Globus A. Principles of information display for visualization practitioners. *The Review of Economic and Business Studies REBS*. 1994; (2):161.
9. Goebel R, Esposito F, Formisano E. Analysis of functional image analysis contest (fiac) data with brainvoyager qx: From single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. *Human Brain Mapping*. 2006; 27(5):392–401. [PubMed: 16596654]

10. Gresh, DL.; Rogowitz, BE.; Winslow, RL.; Scollan, DF.; Yung, CK. WEAVE: a system for visually linking 3-D and statistical visualizations applied to cardiac simulation and measurement data. *Ieee*; 2000. p. 489-492.
11. Hansen PC, O'Leary DP. The use of the l-curve in the regularization of discrete ill-posed problems. *SIAM Journal on Scientific Computing*. 1993; 14(6):1487–1503.
12. Harrower MA, Brewer CA. Colorbrewer.org: An online tool for selecting color schemes for maps. *The Cartographic Journal*. 2003; 40(1):27–37.
13. Heer J, Card SK. Doitrees revisited: scalable, space-constrained visualization of hierarchical data. *Group*. 2004:421–424.
14. Javed W, McDonnell B, Elmqvist N. Graphical perception of multiple time series. *IEEE Transactions on Visualization and Computer Graphics*. 2010:927–934. [PubMed: 20975129]
15. Jianu R, Demiralp C, Laidlaw DH. Exploring 3d dti fiber tracts with linked 2d representations. *IEEE Transactions on Visualization and Computer Graphics*. 2009; 15(6):1449–1456. [PubMed: 19834220]
16. Johnson, B.; Shneiderman, B. Tree-maps: a space-filling approach to the visualization of hierarchical information structures. *Proceedings of the 2nd conference on Visualization '91, VIS '91*; Los Alamitos, CA, USA. IEEE Computer Society Press; 1991. p. 284-291.
17. Johnson CR, Sanderson AR. A next step: Visualizing errors and uncertainty. *IEEE Computer Graphics and Applications*. 2003; 23(5):6–10.
18. Jorm AF, Jolley D. The incidence of dementia: a meta-analysis. *Neurology*. 1998; 51(3):728–733. [PubMed: 9748017]
19. Keim, D.; Mansmann, F.; Schneidewind, J.; Schreck, T. Monitoring network traffic with radial traffic analyzer. *2006 IEEE Symposium On Visual Analytics And Technology*; 2006. p. 123-128.
20. Keim DA. Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics*. 2002; 8(1):1–8.
21. Kniss J, Kindlmann G, Hansen C. Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics*. 2002; 8(3):270–285.
22. Knig, A.; Doleisch, H.; Grller, E. Multiple Views and Magic Mirrors - Multimodal Volume Visualization of the Human Brain. 1999.
23. MacEachren AM, Robinson A, Hopper S, Gardner S, Murray R, Gahegan M, Hetzler E. Visualizing geospatial information uncertainty: What we know and what we need to know. *Cartography and Geographic Information Science*. Jul; 2005 32(3):139–160.
24. Matarazzo JD, Herman DO. Base rate data for the wais-r: test-retest stability and viq-piq differences. *Journal of Clinical Neuropsychology*. 1984; 6(4):351–366. [PubMed: 6501578]
25. McReynolds, T.; Blythe, D. [accessed 5-December-2011] Programming with opengl: Advanced rendering. 1997. <http://www.opengl.org/resources/Online>
26. Meyer M, Munzner T, DePace A, Pfister H. Multeesum: a tool for comparative spatial and temporal gene expression data. *IEEE Transactions on Visualization and Computer Graphics*. 2010; 16(6):908–917. [PubMed: 20975127]
27. Moore EB, Poliakov AV, Lincoln P, Brinkley JF. Mindseer: a portable and extensible tool for visualization of structural and functional neuroimaging data. *BMC Bioinformatics*. 2007; 8(March):389. [PubMed: 17937818]
28. Mueller, K.; Welsh, T.; Zhu, W.; Meade, J.; Volkow, N. BrainMiner: A Visualization Tool for ROI-Based Discovery of Functional Relationships in the Human Brain. 2000.
29. Müller, F. Hypervariate information visualization. In: Baur, D.; Sedlmair, M.; Wimmer, R.; Chen, Y-X.; Streng, S.; Boring, S.; De Luca, A.; Butz, A., editors. *Trends in Information Visualization*. Munich, Germany: 2010.
30. Munzner T. A nested model for visualization design and validation. *Visualization and Computer Graphics, IEEE Transactions on*. Nov-Dec;2009 15(6):921–928.
31. Munzner, T. Image Rochester NY. 2011. Applying information visualization principles to biological network displays.
32. Pang AT, Wittenbrink CM, Lodha SK. Approaches to uncertainty visualization. *The Visual Computer*. 1997; 13(8):370–390.

33. Pieper, S.; Lorensen, B.; Schroeder, W.; Kikinis, R. The NA-MIC Kit: ITK, VTK, Pipelines, Grids and 3D Slicer as An Open Platform for the Medical Image Computing Community. IEEE; Apr. 2006 p. 698-701. Number
34. Rhyne T. Does the difference between information and scientific visualization really matter? IEEE Computer Graphics and Applications. 2003; 23(3):6–8.
35. Roberts, JC. State of the art: Coordinated & multiple views in exploratory visualization. Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization CMV; 2007; 2007. p. 61-71.
36. Rosano C, Aizenstein H, Brach J, Longenberger A, Studenski S, Newman AB. Gait measures indicate underlying focal gray matter atrophy in the brain of older adults. J Gerontol A Biol Sci Med Sci. 2008; 63(12):1380–1388. [PubMed: 19126852]
37. Rbel, O.; Weber, GH.; Kernen, SVE.; Fowlkes, CC.; Hendriks, CLL. PointCloudXplore: Visual analysis of 3D gene expression data using physical views and parallel coordinates. Citeseer: 2006. p. 203-210.
38. Saldanha AJ. Java treeview – extensible visualization of microarray data. Bioinformatics. 2004; 20(17):32463248.
39. Sanyal J, Zhang S, Bhattacharya G, Amburn P, Moorhead RJ. A user study to compare four uncertainty visualization methods for 1d and 2d datasets. IEEE Transactions on Visualization and Computer Graphics. 2009; 15(6):1209–1218. [PubMed: 19834191]
40. Shneiderman B. Treemaps for space-constrained visualization of hierarchies. New York. 2009; 1992(March 1991):1–15.
41. Simonsick EM, Newman AB, Nevitt MC, Kritchevsky SB, Ferrucci L, Guralnik JM, Harris T. Measuring higher level physical function in well-functioning older adults: expanding familiar approaches in the health abc study. The journals of gerontology Series A Biological sciences and medical sciences. 2001; 56(10):M644–M649.
42. Taylor, R. IEEE Computer Graphics and Applications. 2002. Visualizing multiple fields on the same surface; p. 6-10.
43. Teng EL, Chui HC. The modified mini-mental state (3ms) examination. The Journal of clinical psychiatry. 1987; 48(8):314–318. [PubMed: 3611032]
44. Thirion JP. Image matching as a diffusion process: an analogy with maxwell’s demons. Medical Image Analysis. 1998; 2(3):243–260. [PubMed: 9873902]
45. Tory M, Miller T, Atkins MS. Visualization of Time-Varying MRI Data for MS Lesion Analysis. Spie. 2001; 4319:590–598.
46. Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M. Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain. NeuroImage. 2002; 15(1):273–289. [PubMed: 11771995]
47. Wang, W.; Wang, H.; Dai, G.; Wang, H. Visualization of large hierarchical data by circle packing. Proceedings of the SIGCHI conference on Human Factors in computing systems CHI 06; 2006. p. 517
48. Wang Baldonado, MQ.; Woodruff, A.; Kuchinsky, A. Guidelines for using multiple views in information visualization. Proceedings of the working conference on Advanced visual interfaces AVI 00; 2000. p. 110-119.
49. Weber GH, Bethel EW, Fowlkes CC, Hendriks CLL, Depace AH, Simirenko L, Eisen MB, Biggin MD, Hagen H, Malik J, et al. Pointcloudxplore 2 : Visual exploration of 3d gene expression. Gene. 2008; 2(i)

Fig. 1. GRACE collaborative layout, supporting visual integration and comparison of spatial and non-spatial geriatric data. Following a hybrid visual comparison design, the layout consists of five linked views: Two juxtaposed volume rendering panels (left) encode spatial information; Two linked, juxtaposed dendrogram panels and one overlaid Kiviat diagram panel (center) encode non-spatial information; Additional widgets (right, top, bottom) allow data selection (datasets, variables, and algorithms), filtering (data types, correlations and significance values) and data manipulation. Correlations are encoded across all panels using two sequential multihue color schemes (green-blue for gray matter, and orange-red for white matter). This flexible design successfully supports the comparison of spatial and non-spatial variables, functional attributes, datasets, and algorithms, while interaction enables seamless formation and refinement of hypotheses.

Fig. 2.

GRACE framework. Medical images acquired through a multitude of imaging protocols are first normalized against a region of interest (ROI) anatomical brain template, in order to enable spatial analysis across hundreds of subjects. The resulting ROI-based measurements, as well as the non-spatial measurements are then numerically processed to quantify correlations. The results, together with the original template and additional hierarchical ROI information are then visually encoded and presented to the user. The domain-expert controls the visual interface, and uses this interface to formulate new hypotheses.

Fig. 3.

Dendrogram panel encoding the functional structure of the brain. Leaves represent brain regions, e.g. cerebellum, while inner nodes correspond to groups of regions, i.e., regions associated with Mobility (Mob), Visuospatial Attention (VA), Motor Imagery (MI), Cognitive Processing Speed (CPS), or Memory (Mem). The "Cerebellum (R)" label is the result of a mouse hovering action. The color-map is linked to the one used for volume rendering and encodes the significance level of the correlations.

Fig. 4.

Comparison of two gait variables, bradykinesia (top) vs. gait disturbance (bottom) using macro-structure brain measurements. The brain regions correlated with bradykinesia (upper volume rendering panel) were distributed more cortically than those for gait disturbance (lower volume rendering panel). Also noticed here is that only two of the mobility regions and only one motor imagery region associated with gait disturbance (lower dendrogram) were significant at the 5% level. This prompted the lead GRACE researcher to examine next the correlations between gait variables and microstructure brain measurements.

Fig. 5. Comparison of two algorithms (iterative Tikhonov Regularization vs. SPLS) used to compute the associations between brain regions and DSST. Visual analysis of all visual encodings indicates very limited overlap between the results generated by the two algorithms. To the researchers' surprise, the correlated regions selected by the two algorithms were significantly different.

Table 1

Task Abstractions and Instantiations for Spatial and Non-Spatial Geriatric Analysis

Index	Task Abstraction	Instantiation
T1	For the same non-spatial variable, compare correlations of a particular spatial variable to the correlations of another spatial variable	<i>Compare correlations of microstructure to correlations of macrostructure for the same gait variable.</i>
T2	For the same spatial variable, compare correlations of a particular non-spatial variable to the correlations of another non-spatial variable	<i>Compare correlations of one gait variable to correlations of another gait variable</i>
T3	Compare correlations of a subgroup of spatial locations, for a particular non-spatial, hierarchical attribute	<i>Compare correlations of subgroups of brain regions, according to the functions they perform (e.g., in Mobility)</i>
T4	Detect spatial patterns for a particular non-spatial variable	<i>Visually see brain spatial patterns for tremor</i>
T5	Detect non-spatial variables that vary in the same way, followed by T1	<i>Visually see spatial patterns for gait variables that vary in the same way</i>
T6	Compare correlations computed through a particular algorithm to correlations computed through a different algorithm	<i>Visually see whether SPLS and Tikhonov yield similar correlations</i>
T7	Dynamically query correlations based on confidence level, followed by T1–T6	<i>Form and refine gait-brain hypotheses by filtering correlations to a desired confidence level</i>