

ENIGMA-Viewer: Interactive Visualization Strategies for Conveying Effect Sizes in Meta-Analysis

Guohao Zhang
Computer Science and
Electrical Engineering
University of Maryland,
Baltimore County
Baltimore, Maryland
guhaozhang@umbc.edu

Neda Jahanshad
Keck School of Medicine
University of Southern
California
Los Angeles, California
neda.jahanshad@usc.edu

Peter Kochunov
Maryland Psychiatric
Research Center
University of Maryland
Baltimore, Maryland
pkochunov@mprc.
umaryland.edu

Paul Thompson
Keck School of Medicine
University of Southern
California
Los Angeles, California
pthomp@usc.edu

Elliot Hong
Maryland Psychiatric
Research Center
University of Maryland
Baltimore, Maryland
ehong@mprc.umaryland.edu

Jian Chen
Computer Science and
Electrical Engineering
University of Maryland,
Baltimore County
Baltimore, Maryland
jichen@umbc.edu

ABSTRACT

We present ENIGMA-Viewer, an interactive visualization tool for scientists to compare effect sizes and association results on standardized regions of interest from different studies, as performed in large scale consortia studies such as those in the Enhancing Neuro Imaging Genetics through Meta Analysis (ENIGMA) consortium. We report how visualization methods are designed for interactive comparison of brain diffusion tensor imaging (DTI) and cortical thickness statistical test results. ENIGMA-viewer is an open-source software tool, publicly accessible through the NITRC website or online at <http://enigma-viewer.org>.

Categories and Subject Descriptors

H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces—*Web-based interaction*

General Terms

Design

Keywords

3D viewer; neuroimaging; meta-analysis; effect size; interactive visualization; web-based application

1. INTRODUCTION

Large scale harmonization of image processing protocols across different studies around the world and the extraction

of effect sizes across reliably extracted regions of interest, allows for a common framework through which results can be compared, and combined through unbiased meta-analyses as performed in the ENIGMA (Enhancing NeuroImaging Genetics through Meta-Analysis) consortium [42]. offer essential opportunities for brain scientists to produce credible findings through meta-analysis, a method that combines data cohorts collected worldwide to obtain the statistical power otherwise unavailable from a single cohort, in order to find cross-modality data associations that influence brain structures [16]. Cohort studies boost power to detect associations. Seminal accomplishments with promising results in imaging-genomics associations have accelerated scientific discoveries in areas such as schizophrenia [25], bipolar disorders [13], and other neurodegenerative diseases [41].

In meta-analysis, one important task is to interpret effect size, a statistical measure that can be broadly defined as any statistic that quantifies the degree to which sample results diverge from the expectations in the null hypothesis. Computing effect size is important because if effect sizes are stable across studies or even generalizable over some variations in design or analysis, the results are replicable. That is, effect size is a statistical tool for meta-analysis that quantitatively synthesizes effects across different studies. Ranking brain measures in order of their effect sizes for case-control differences can unearth brain measures on the basis of both the stability of the brain volume measures (so-called heritability [24]) and their relevance in the disease being studied [29].

Comparing effect sizes is, however, a multi-variant issue, not only because scientists must choose studies carefully to ensure consistency of protocol use, but also because the variety of cohorts has made it possible to dig more deeply into and disentangle the sources (medication-related geographical or demographics and genetic factors (e.g. [30]) of variations that could explain why brain differences vary across studies and in different phenotypes.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

BCB'16, October 2–5, 2016, Seattle, WA, USA.

Copyright 2016 ACM 978-1-4503-4225-4/16/10 ...\$15.00.

<http://dx.doi.org/10.1145/2975167.2985689>.

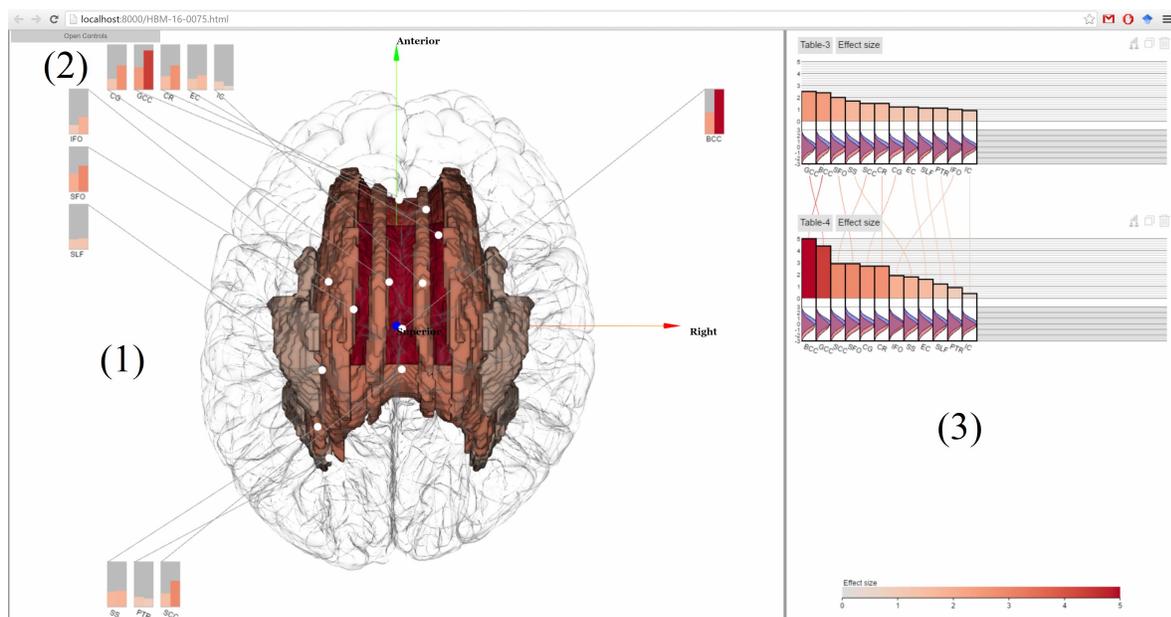


Figure 1: ENIGMA-Viewer Interface: Data are taken from a comparison of normal control and a diseased brain in [23]. (1) **Region-centric view.** Brain white matter regions of interest from the ENIGMA-DTI protocols are shown inside a glass brain and the associated effect sizes of these selected regions belonging to two studies are shown in bar charts (2). (3) **Study-centric view.** The effect sizes are sorted from largest to smallest within each study. The white matter microstructural measure derived from diffusion tensor imaging, fractional anisotropy (FA), distributions across subjects of each of these brain regions are shown. This study cohort includes 243 control subjects and 188 schizophrenia patients.

As new analytical results are produced and lead to increased data dimensionality and size, the bottleneck to human understanding is not limited to data mining and computational approaches, but also to human limited memory capacity. Presenting and interpreting effect sizes and locating regions across studies can be obscured due to the complexity of interpreting the multivariate information space and the problems inherent in presenting rich datasets on a two-dimensional (2D) computer screen. Synthesizing new information for new discoveries and comparison to past results is cognitively demanding. The bandwidth of discovery will be bounded by the characteristics of human perception, and hence the quest for visualization has commenced in the brain sciences, as evidenced by recent reviews and by research on the vital role of visualization in the analysis of multimodal neuroimaging data [28, 31, 43].

Our long-term goals include making analytics results derived from the ENIGMA pipeline accessible to the neuroscience community at large and assisting brain scientists in seeing patterns in massive multimodal computational solutions, as well as encouraging effective communication and collaborative activities through visual means to convey our results to the general public. Here we present ENIGMA-Viewer (Figure 1), an interactive visualization tool to let users explore multimodality brain data to compare effect sizes and associated brain anatomical structures and genomics factors. This work makes several contributions:

- A series of design strategies for spatial and non-spatial parametric data integration in the context of mega-analysis of imaging-genetics.

- Consideration of the brain science domain and tools to aid multivariate comparison studies.
- Priority in integrating different imaging modalities to compare results and locate important information.

2. RELATED WORK

Both neuroscience and visualization scientists have worked extensively on visualizing brain datasets. This section reviews related work in visualization, related data analysis, and multi-modality data visualization.

2.1 Brain Data Visualization

Many brain data visualization tools have addressed important issues in conveying single modality imaging techniques. In tensor field DTI data visualization, Laidlaw et al. designed multivariate tensor field visualization at every voxel using creative artistic rendering [26]. Other powerful techniques have used tensor glyphs to convey tensor shapes [37], or non-photorealistic rendering to resolve complex spatial depth perception [15] as well as validating studies in the large display uses [5] and rendering solutions [32, 46]. Functional brain network visualizations have showed bundling 3D trajectories can support functional network understanding in both 3D [4] or 2D connectivity studies in matrix views [50].

Despite these creative solutions and technical advances, none of the work to our knowledge has exploited features and interpreting results across multiple modalities and multiple datasets, except our own work by Novak et al. [31] and Zhang et al. [49]. A main difference between single

and cohort analyses is that single images become unimportant and statistical results comparing cohorts can lead to valuable understanding of associations between brain regions and diseases. Kehrner et al. have laid out important design challenges in multimodality multi-faceted data visualization in the broad medical imaging areas related to comparative studies as well as possible solutions in the use of multiview visualization to represent multidimensional data [22]. This work follows the multiview solutions to let scientists visually synthesize results from different views.

The work most closely related to ours is Novak et al.’s EnigmaVis. That work lets scientist make quick comparisons among new and existing DTI-GWAS queries through a powerful web interface [31]. This pioneering study is significant because it supports quick hypothesis confirmation through comparison and let brain scientists explore studies and load data in an intuitive visualization environment. Inspired by this study, our work advances the tool development by presenting and combining different imaging and measurements results as well as supporting interactive exploration and comparison between modalities. Our design is very different from that of EnigmaVis in that no prior hypotheses or knowledge of prior studies is required for making a visualization. Brain scientists can freely explore the space and load and compare their data. We believe our solution has excellent potential to support opportunistic discovery and may enable scientists to more easily and interactively investigate broader scientific questions.

2.2 Integrating Spatial and Non-Spatial Data

Our solution to comparative effect sizes is related to spatial and non-spatial data integration to assist data analysis. Our choices of visualization is mostly driven by data types, which is most similar to the design rationale in Keefe et al. [21], where they visualize quantitative parameters using non-spatial data visualization to avoid inaccurate judgement of three-dimensional measurement. Wang and Tao also defines the integration of spatial and non-spatial data visualization [45], as well as Chen, Pyla, and Bowman in three-dimensional interface design [7].

3. SCIENTIFIC BACKGROUND AND DATA SOURCE

In this section we talk about the background that motivates our visualization design, followed by description of the data used in the visualization.

3.1 Introduction to the Goals in ENIGMA DTI-GWAS Data Analysis

The ENIGMA consortium aims to enable image-genetics discoveries by analyzing brain imaging measures and genotypes to examine reproducibility, heritability, and association with diseases [42]. The goal is to address the most fundamental questions in neuroscience by linking brain measures to human well-being. Some of the most intriguing questions include: *what are the effects of aging, degenerative disease and psychiatric illness on the living brain? How do brain measures relate to cognition and behavior? Do brain measures predict our risk for disease, or give prognoses for those who are ill?* [42]

The method is meta-analysis, a quantitative statistical analysis of several separate but similar experiments or studies using pre-agreed covariates in order to test the pooled

data and examine the effectiveness of the results [42]. Subsequently, the p -values and regression coefficients are combined by weighting the results based on the sample size of each contributing cohort. Meta-analysis is not only important for brain white-matter analysis, but it has been the only way to find credible genetic traits of brain disorders with sufficient statistical power to achieve significant effects greater than $p < 10^{-8}$.

Great steps forward have laid foundation for making cohort-comparison possible by addressing challenging technical problems in multiple areas. These include creating common ENIGMA- template [16], harmonization of protocols to synthesize data captured with different protocols [14], generating tract-based spatial statistics skeletonization [39], regions of interests (ROI) extraction [17], and SOLAR statistics [3]. Meta-analyses have also identified the stabilities of brain volume measures (or heritability) in sub-cortical (containing regions associated with human function) and cortical regions across twins, genders, and geolocations [10].

A common workflow in performing meta-analysis is first to follow pre-determined protocols to obtain desirable imaging modalities (here DTI) and genomics in the population under investigation. Tract or voxel-based analyses and associated metrics measures (e.g., fractional anisotropy (FA) or water diffusion and cortical thickness) sensitive to the neuro-degenerations are then derived. Effect sizes in DTI studies are quantitatively compared.

3.2 Data

Brain Imaging Data. The 3D brain imaging dataset labeled total 48 cortical and sub-cortical regions contained in the JHU white matter atlas [44]. The brain volume in this atlas has $182 \times 218 \times 182$ voxels measured at the resolution of $1 \times 1 \times 1$ millimeters. We extract the surface mesh for each cortical region from the atlas. using FreeSurfer [11]. Since the FreeSurfer and the JHU atlas has different numbers of cortical regions, these 70 FreeSurfer regions are transformed to match the JHU atlas using the linear transformation function in FSL [18].

Statistical Analysis Data. The statistical data used in program is from recent studies from ENIGMA group. For sub-cortical analyses on FA values for DTI data, effect sizes are reported as overall Cohen’s d values for case/control effects and Z-scores for quantitative effects from linear regressions of individual subjects. An example data is from the study on the heterochronicity of white matter development and effect of aging in schizophrenia [23]. That study computes effect size values for 12 affected sub-cortical brain regions contained in the JHU atlas.

4. TASK ANALYSIS

The first goal of this study is to characterize the problems being addressed by the brain scientists as visualization tasks.

4.1 Procedure

The task analysis was achieved by working closely with brain scientists, as well as by literature review. Each scientist was interviewed to gather sufficient information on their workflow tasks and goals. Each participant also used our prototype tool of ENIGMA-Viewer and suggested action steps and desirable outcomes. To collect the resulting feedback, we have asked them to answer the following questions: What kinds of questions do you anticipate ex-

Table 1: DTI-GWAS Task List. Our ENIGMA-Viewer is designed for these analytical tasks.

Index	Task	Example questions
1	Viewing distribution, consistency, or inconsistency	What is the most affected brain regions under a certain disorder? What is the overall data distribution of the effect size in all studies? Are the results consistent or very different?
2	Detect trends	What is the distribution of cortical thicknesses and FA values? How does effect size vary among studies?
3	Find association	What are the disorder (brain regions) and genetics correlates to risk? How does brain structural change associated with behavioral risks and changes and vulnerability? How do geographical factors affect disease expression in the brain?
4	Locate extremes	What are the significant brain regions mostly affected by diseases?
5	Find local relationships	What are the differences between studies in terms of their effect size?
6	Compare different disorders	What are the common and different effects in the brain networks between or among multiple disorders?
7	Compare disorders by regions	What are the regions affected or unique to a certain disorder?

ploring using the visualization tool? What would you like to achieve using visualization in general, communication or seeing patterns? Why do the state-of-the-art tools, such as AFNI [9], FSL [40], DtiStudio [19], not address your needs? How would you like the data to be depicted and represented? Should the data be visualized in 2D or 3D? How would you like to interact with and explore the datasets?

4.2 Task List

Neuroscientists are interested in detecting trends and viewing overall data distribution as well as individual regions of interest. The most important tasks are related to (1) comparing similarities and differences in different disorders or in disease and control conditions; (2) comparing effect sizes in meta-analysis to find the truly significant brain regions and associated genetics factors; (3) studying the most important genetics association with these brain regions to establish the DTI-GWAS association; (4) identifying brain regions with high and low heritability. Each of these domain tasks can be abstracted to the fundamental analytic tasks presented by Amar, Eagan, and Stasko [1] and Schulz et al. [38], as listed in Table 1.

5. ENIGMA-VIEWER

This section presents our main contribution, i.e., the design decisions made in the ENIGMA-viewer to address all those users’ tasks.

5.1 Overview of the Design Considerations

The possibilities for encoding and interacting with the data mentioned in Section 2 are vast. Our encodings and layout draw upon existing idioms, and in places the framework suggests that more novelty is required. We investigate visual design options through our experience of working on interface layout, discussion among the team of co-authors and following good design principles. We have also designed interaction techniques such that results from one data type and modality can guide comparative analysis of another in a unified interface level. Data belonging to different types can be visually linked through interaction.

We use *juxtaposition*, which places effect sizes, 3D anatomical regions, and artificial GWAS side-by-side in small-multiples displays similar to that of Chen et al. [6]. We also use *super-*

position which the effect sizes and 3D anatomical regions are overlaid in the same frame of reference, following the comparative visualization classification by Gleicher et al. [12] and Karnick et al. [20].

5.2 Visual Data Fusion

Visual data fusion intermixes different facets of scientific data in a single view using a common frame of reference. In our program, effect sizes in different study cohorts and 3D anatomical regions can be grouped and presented in the space of the 3D glass brain. This visualization addresses scenarios of use in which a brain scientist wants to focus on finding associations of effect sizes in one or more regions of interests. For example, a brain scientist can load new and existing studies and then inspect trends and differences among studies visually. Another example use is to study multiple closely proximate brain regions of cortical and sub-cortical regions. Data from these spatial locations and multiple effect sizes can be discussed together. When the brain scientist’s task is to search for associative relationships between different studies in a common region of interest, this visual fusion would be appropriate to let the user focus simply on one view to obtain all information.

5.3 Focus+Context

Focus+context visualization supports both focused and detailed views as well as context for navigation purpose. Effect sizes of each cohort are displayed in small multiples using bar charts ordered by effect size magnitudes. Since the effect sizes vary across studies, using uniform-scale bar charts would render smaller effect sizes too small to be distinguishable visually. Our solution is to color the magnitudes of the effect sizes. This strategy introduces dual encoding to encode the magnitudes of effect sizes: the bars use length with the most precise magnitude discrimination, while colors encourage pattern finding to locate extreme effect size magnitudes in different cohorts. The diverging color map is perceptually linear and the zero mark appears where the two colors intersect at 0 to represent the least significant effect size. Since effect size is a positive value, the two colors (blue and red) indicate which is larger, the control or the diseased population. When the control is large, the effect size is mapped to the red hue, and otherwise to blue. In this way, users can obtain at a glance the most signifi-

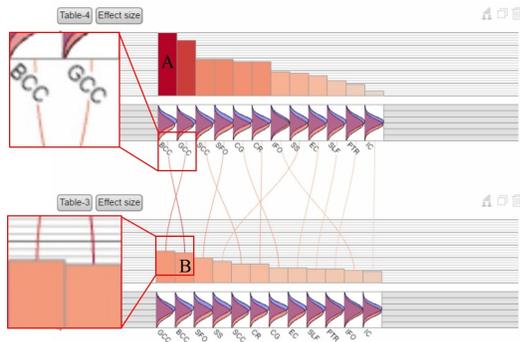


Figure 2: Scented edges to reduce the context switching cost. Here curved lines connect the corresponding regions between two studies (top and bottom ones). The color on each curve varies gradually in the way that the color in the top bar uses the color in the bottom one and vice versa. In this way, a viewer does not need to trace the link to compare study differences.

cant brain regions by searching for the most saturated red or blue regions. We follow best practice and psychophysical principles in using monotonic luminance [35] to show differences in effect sizes. We plot the FA distributions between the patient and control cohorts in order to show the FA differences. Here we can observe that the control cohort has higher average FA values then the diseased ones in all regions.

5.4 Reduce Context Switching Cost

The cost of context switching in visualization [8, 36] is one drawback of the small-multiples display in which bar charts are placed side-by-side. For searching for association between studies, the viewer must constantly switch the viewpoint between studies to look for relevant information in other views. To reduce the cost, our current method is to use “information scent” [33, 48], nuances added to the display to help the user construct visual associations. The edges are “scented” using the color representing the effect size magnitude in the neighboring view so that the user need not visually trace the edge to learn the magnitude in the other view.

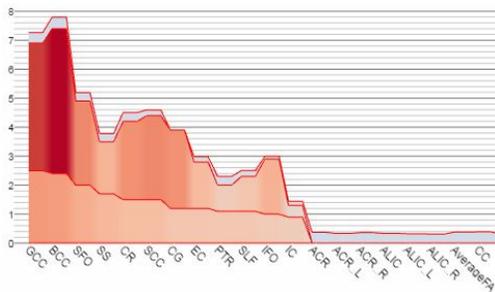


Figure 3: Effect sizes are stacked together to support comparisons among cohorts in different brain regions. Horizontally, the cohorts are ordered by the effect size magnitudes in the bottom cohort. A viewer can easily finds differences between studies.

The second way to reduce the context switching cost and to facilitate comparison of common regions is to use the stacked bar chart (Figure 3). Effect sizes in the same brain region belonging to different cohorts are stacked together and horizontally, the cohorts are ordered by the effect size magnitudes in the bottom cohort. This view facilitates both between and within effect sizes of the same and different brain regions and saves space. It is also easy to find region choice discrepancies between or among studies because some studies many include more regions than others.

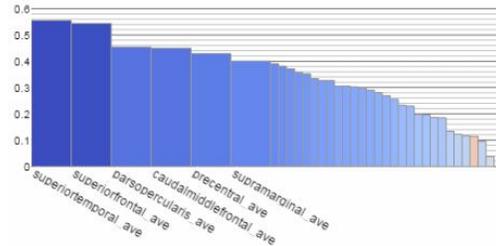


Figure 4: Bars can be re-scaled to make whole dataset visible in one view while leaving regions with large effect size enough screen space. In this example, bars with effect size smaller than 0.4 have their width narrowed to one fifth of normal width.

Our design follows importance-driven interactions. If the screen space is not enough to show all bars for all regions, we keep the important ones, e.g. those with large effect size, unchanged and make less important ones smaller and in context. This scaling mechanism makes the larger effect size regions visually salient. The user can directly interact with the views to rescale the size of the bar charts. Figure 4 shows an example where the bars with the effect sizes lower than 0.4 are toggled to have one fifth of normal bar width and their labels hidden.

5.5 View Reconfiguration

Our tool supports a set of interaction techniques: linking and brushing, zooming, panning, and view reconfiguration. The viewer can manually select interesting brain regions under study in the effect size bar charts and examine the spatial location in the 3D view via brushing [2]. Multiple regions of interest can be selected and visualized and also linked to the artificial Manhattan plot (Figure 5).

Our tools also support drag-and-drop operations to facilitate inter- and intra-study comparison. The bar charts can be dragged and dropped next to other bar chart or to the spatial view. Dragging-and-dropping a bar chart next to other bar charts can be used to rearranged the layout of multiple bar charts, which could make comparison between different studies easier. The user can also drag the bars from the right-side bar chart to the 3D glass brain regions. This action results in the display of a region-centric comparison chart. Brain regions currently being selected will be shown. This design provides a region-specific comparison mechanism.

5.6 Multimodality Visualization

Our visualization supports multimodality visualization in that multiple attributes of brain regions can be visualized together. As can be seen from Figure 6, the Manhattan and

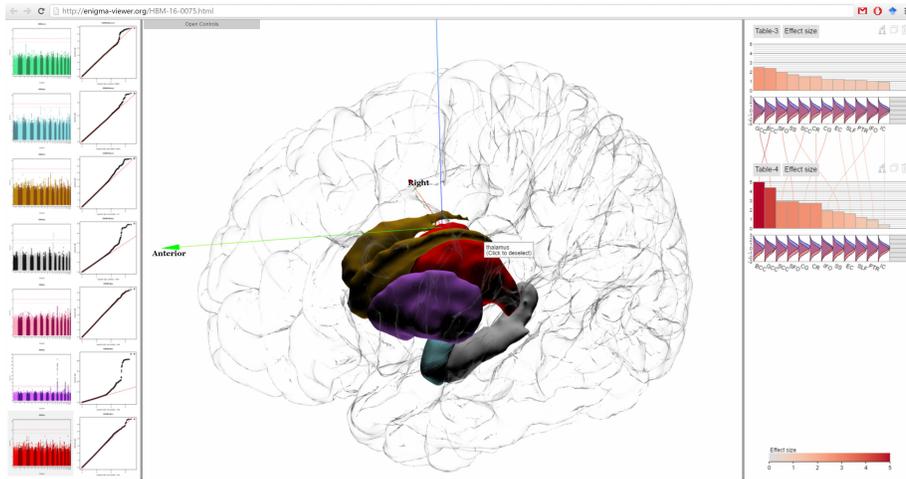


Figure 5: Querying Manhattan plot will highlight 3D regions. The center view displays seven sub-cortical regions and the Manhattan and QQ plots from GWAS analysis. Those regions are colored using the same colors as those in the Manhattan plots.

QQ plots are linked to the 3D brain regions. Figure 7 also shows that the plot modality and the chart modality are both linked to the 3D view. This makes it convenient to visualize multiple attributes of brain regions in the 3D view.

5.7 Implementation

ENIGMA-Viewer is implemented in Google WebGL and JavaScript and can be executed on major web-browsers such as Safari, Firefox, and IE, without requiring any third-party software or add-ons. No account or authorization is required to use ENIGMA-Viewer and users are encouraged to email the developers with all comments and suggestions.

The sub-cortical geometry data is extracted from sub-cortical white matter atlas using marching cubes algorithm [27]. To ensure a fast loading, the mesh is only extracted when a region is selected in the 3D view. We only stored the atlas volume to reduce the data to be loaded to the browser.

6. CASE STUDIES

This section shows two example uses of ENIGMA-Viewer.

6.1 Cortical and Sub-Cortical Effect Sizes

In Figure 6, sub-cortical regions are associated with the Manhattan plot and cortical regions are associated with the effect size displays on the right. Here the user has selected all cortical regions, as shown in the bottom two rows of the effect-size displays. The same regions in all other views are also highlighted. The colors in different views match the associated effect sizes throughout.

6.2 Cortical Effect Sizes

Here only cortical regions are selected (Figure 7). Tooltips show the detailed information about that region when the highlighted regions are moused over.

7. DISCUSSION

This section discusses alternative designs and the logical next steps to improve the usefulness and usability of the visualization design.

7.1 The Interpretation of Effect Size

Visualizing effect sizes only is not enough; neuroscientists must interpret and evaluate effect size for its practical significance and interpret other factors that cause the differences in different studies. Thus, how to interpret effect size is also a crucial question. The common practice here is to use the benchmarks for “small”, “medium”, and “large” effects. However, often this categorization depends on the domain of use and applying existing guidelines directly can be inappropriate.

Because such interpretation often depends on prior effect sizes in the related literature, both the size and nature of the effect should be included in the interpretation to increase the practical significance. Effect size also depends on multiple factors such as the context of the study, the importance of the outcomes, etc. Thus, visualizing the effect sizes and automatic searching and showing related context information to assist effect size interpretation within and between studies is the logical next step in truly helping neuroscientists’ decision -making.

7.2 Limitations of the Current Study

Though we have attempted to use good visual design principles to guide our study, validation is the next step in creating truly meaningful tools for neuroscientists. One cannot assume that understanding has been gained from data represented graphically merely from the fact that the visualization has been presented; still less can one assume that a specific visual analysis can integrate all the capabilities required for multifaceted, spatially complex data analysis. Perception and cognition are complex and evaluation of visualization approaches for information presentation and interpretation is much needed to validate our design.

Our next step is also to improve data processing, handling, sharing, and collaborations using common infrastructure and data format standards. We plan to follow the data-format protocols defined by the ENIGMA consortium. We also plan to design dynamic and interactive queries among variables to support dynamic data analysis and to maximize

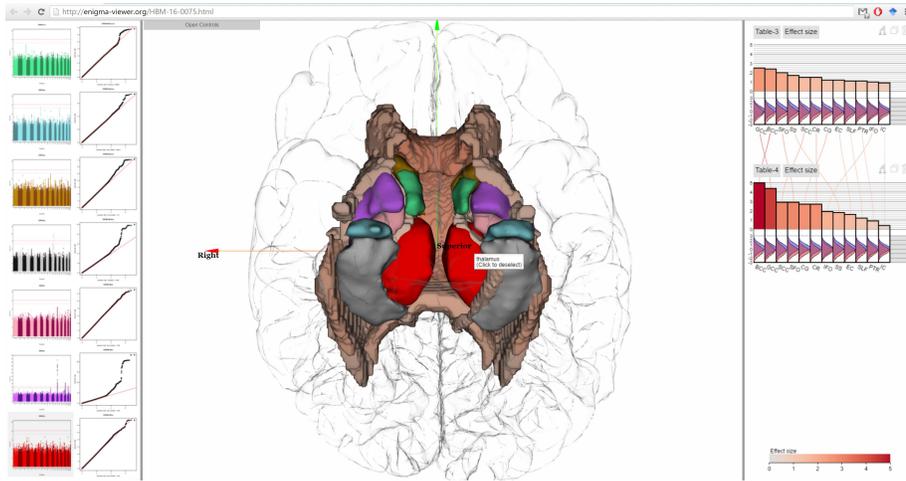


Figure 6: Cortical and sub-cortical regions are highlighted in the glass brain. To the regions in 3D view have the same color as those used in Manhattan plots or bar charts.

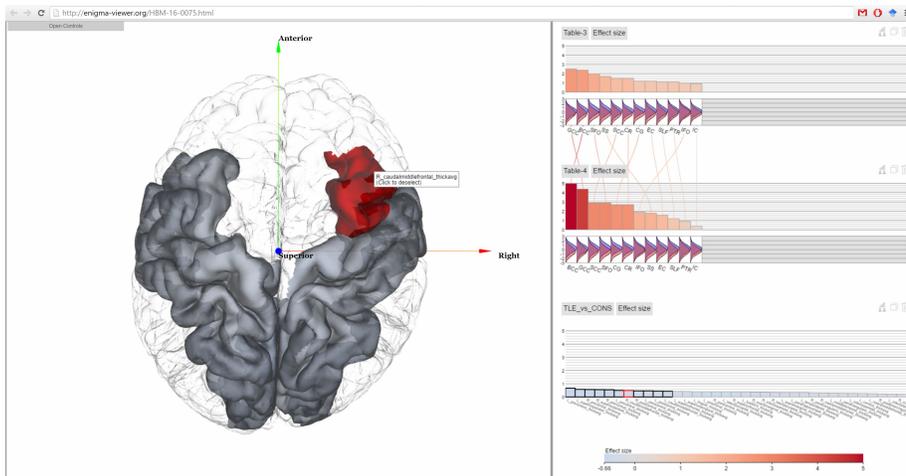


Figure 7: One cortical region is highlighted in the glass brain. Cortical regions can be visualized and interacted with sub-cortical regions together using the same method. The *TLE_vs_CONS* dataset is from a study on 68 cortical regions on temporal lobe epilepsy and contains 339 normal controls and 415 patients [47].

flexibility to cross-link or brush-and-link across displays to find data relationships and compare and filter to remove redundancy. We will integrate the computational solutions and construct the entire workflow so that analysis and visualization can be integrated in a single framework, thus easing computing, data exploration, and human understanding of the massive datasets. The linked multiview visualization also provides a solution for brain scientists to understand how the statistics data provenance [34] in that the factors, such as cohort distributions, that are used to produce the results can be visualized in addition to the final statistics displayed on the 3D brain.

8. CONCLUSIONS

Meta-analytic thinking would considerably facilitate knowledge accumulation in brain science. In principle, meta-analysis does not overemphasize the outcomes of statistical tests in individual studies; instead, it stresses the need to

explicitly design and place studies in the context of prior research. The visualization, reporting, and interpreting of effect sizes are ways to make more explicit comparisons across cohorts in a meta-study. It is particularly beneficial to incorporate prior effect sizes to guide the findings even before a study is conducted. Visualization tools improve accessibility and facilitate quick pattern finding. Our present tool has been used by several teams and has demonstrated the power of visualization to assist reporting and interpreting effect size measures.

9. ACKNOWLEDGMENTS

This work was supported in part by US National Science Foundation IIS-1302755, ABI-1260795, EPS-0903234, and DBI-1062057 and by US National Institutes of Health R01 grants EB015611, NS080655, AG040060, EB008432, MH097268, MH085667, MH089722 and MH094343, and grants U01 AG024904 and P41 EB015922. Any opinions, findings,

and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation and National Institute of Health.

10. REFERENCES

- [1] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization*, pages 111–117, 2005.
- [2] R. A. Becker and W. S. Cleveland. Brushing scatterplots. *Technometrics*, 29(2):127–142, 1987.
- [3] J. Blangero and L. Almasy. SOLAR: sequential oligogenic linkage analysis routines. *Population Genetics Laboratory Technical Report*, 6, 1996.
- [4] J. Böttger, A. Schäfer, G. Lohmann, A. Villringer, and D. S. Margulies. Three-dimensional mean-shift edge bundling for the visualization of functional connectivity in the brain. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):471–480, 2014.
- [5] J. Chen, H. Cai, A. P. Auchus, and D. H. Laidlaw. Effects of stereo and screen size on the legibility of three-dimensional streamtube visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2130–2139, 2012.
- [6] J. Chen, A. S. Forsberg, S. M. Swartz, and D. H. Laidlaw. Interactive multiple scale small multiples. In *IEEE Visualization*, 2007.
- [7] J. Chen, P. S. Pyla, and D. A. Bowman. Testbed evaluation of navigation and text display techniques in an information-rich virtual environment. In *IEEE Proceedings of Virtual Reality*, pages 181–289, 2004.
- [8] G. Convertino, J. Chen, B. Yost, Y.-S. Ryu, and C. North. Exploring context switching and cognition in dual-view coordinated visualizations. In *Proceedings of the International Conference on Coordinated and Multiple Views in Exploratory Visualization*, pages 55–62, 2003.
- [9] R. W. Cox. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, 29(3):162–173, 1996.
- [10] A. den Braber, M. M. Bohlken, R. M. Brouwer, D. van’t Ent, R. Kanai, R. S. Kahn, E. J. de Geus, H. E. H. Pol, and D. I. Boomsma. Heritability of subcortical brain measures: a perspective for future genome-wide association studies. *NeuroImage*, 83:98–102, 2013.
- [11] B. Fischl. FreeSurfer. *Neuroimage*, 62(2):774–781, 2012.
- [12] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, 2011.
- [13] D. Hibar, L. T. Westlye, T. G. van Erp, J. Rasmussen, C. D. Leonardo, J. Faskowitz, U. K. Haukvik, C. B. Hartberg, N. T. Doan, I. Agartz, et al. Subcortical volumetric abnormalities in bipolar disorder. *Molecular Psychiatry*, 2016.
- [14] D. P. Hibar, J. L. Stein, N. Jahanshad, O. Kohannim, X. Hua, A. W. Toga, K. L. McMahon, G. I. de Zubicaray, N. G. Martin, M. J. Wright, et al. Genome-wide interaction analysis reveals replicated epistatic effects on brain structure. *Neurobiology of aging*, 36:S151–S158, 2015.
- [15] T. Isenberg. A survey of illustrative visualization techniques for diffusion-weighted mri tractography. In *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*, pages 235–256. Springer, 2015.
- [16] N. Jahanshad, P. V. Kochunov, E. Sprooten, R. C. Mandl, T. E. Nichols, L. Almasy, J. Blangero, R. M. Brouwer, J. E. Curran, G. I. de Zubicaray, et al. Multi-site genetic analysis of diffusion images and voxelwise heritability analysis: A pilot project of the ENIGMA-DTI working group. *Neuroimage*, 81:455–469, 2013.
- [17] N. Jahanshad, O. Kohannim, A. W. Toga, K. L. McMahon, G. I. De Zubicaray, N. K. Hansell, G. W. Montgomery, N. G. Martin, M. J. Wright, and P. M. Thompson. Diffusion imaging protocol effects on genetic associations. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 944–947, 2012.
- [18] M. Jenkinson, C. F. Beckmann, T. E. Behrens, M. W. Woolrich, and S. M. Smith. FSL. *Neuroimage*, 62(2):782–790, 2012.
- [19] H. Jiang, P. C. van Zijl, J. Kim, G. D. Pearlson, and S. Mori. DTIStudio: resource program for diffusion tensor computation and fiber bundle tracking. *Computer Methods and Programs in Biomedicine*, 81(2):106–116, 2006.
- [20] P. Karnick, D. Cline, S. Jeschke, A. Razdan, and P. Wonka. Route visualization using detail lenses. *IEEE Transactions on Visualization and Computer Graphics*, 16(2):235–247, 2010.
- [21] D. Keefe, M. Ewert, W. Ribarsky, and R. Chang. Interactive coordinated multiple-view visualization of biomechanical motion data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1383–1390, 2009.
- [22] J. Kehrer and H. Hauser. Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 19(3):495–513, 2013.
- [23] P. Kochunov, H. Ganjgahi, A. Winkler, S. Kelly, D. K. Shukla, X. Du, N. Jahanshad, L. Rowland, H. Sampath, B. Patel, et al. Heterochronicity of white matter development and aging explains regional patient control differences in schizophrenia. *Human Brain Mapping*, 2016.
- [24] P. Kochunov, N. Jahanshad, D. Marcus, A. Winkler, E. Sprooten, T. E. Nichols, S. N. Wright, L. E. Hong, B. Patel, T. Behrens, et al. Heritability of fractional anisotropy in human white matter: a comparison of human connectome project and ENIGMA-DTI data. *Neuroimage*, 111:300–311, 2015.
- [25] P. Kochunov, N. Jahanshad, E. Sprooten, T. E. Nichols, R. C. Mandl, L. Almasy, T. Booth, R. M. Brouwer, J. E. Curran, G. I. de Zubicaray, et al. Multi-site study of additive genetic effects on fractional anisotropy of cerebral white matter:

- comparing meta and megaanalytical approaches for data pooling. *NeuroImage*, 95:136–150, 2014.
- [26] D. H. Laidlaw, E. T. Ahrens, D. Kremers, M. J. Avalos, R. E. Jacobs, and C. Readhead. Visualizing diffusion tensor images of the mouse spinal cord. In *IEEE Proceedings of Visualization*, pages 127–134, 1998.
- [27] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *ACM Siggraph Computer Graphics*, volume 21, pages 163–169, 1987.
- [28] D. S. Marcus, M. P. Harms, A. Z. Snyder, M. Jenkinson, J. A. Wilson, M. F. Glasser, D. M. Barch, K. A. Archie, G. C. Burgess, M. Ramaratnam, et al. Human connectome project informatics: quality control, database services, and data visualization. *Neuroimage*, 80:202–219, 2013.
- [29] S. E. Medland, N. Jahanshad, B. M. Neale, and P. M. Thompson. Whole-genome analyses of whole-brain data: working within an expanded search space. *Nature Neuroscience*, 17(6):791, 2014.
- [30] M. Nivard, K. Verweij, C. Minică, J. Treur, E. M. Derks, S. Stringer, K. J. Verweij, N. A. Gillespie, J. M. Vink, C. C. Minică, et al. Connecting the dots, genome-wide association studies in substance use. *Molecular Psychiatry*, 2016.
- [31] N. M. Novak, J. L. Stein, S. E. Medland, D. P. Hibar, P. M. Thompson, and A. W. Toga. EnigmaVis: online interactive visualization of genome-wide association studies of the enhancing neuroimaging genetics through meta-analysis (ENIGMA) consortium. *Twin Research and Human Genetics*, 15(03):414–418, 2012.
- [32] D. Penney, J. Chen, and D. H. Laidlaw. Effects of illumination, texture, and motion on task performance in 3D tensor-field streamtube visualizations. In *IEEE Pacific Visualization Symposium*, pages 97–104, 2012.
- [33] P. Pirolli. The use of proximal information scent to forage for distal content on the world wide web. *Adaptive Perspectives on Human-Technology Interaction: Methods and Models for Cognitive Engineering and Human-Computer Interaction*, pages 247–266, 2006.
- [34] E. D. Ragan, A. Endert, J. Sanyal, and J. Chen. Characterizing provenance in visualization and data analysis: an organizational framework of provenance types and purposes. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):31–40, 2016.
- [35] B. E. Rogowitz and A. D. Kalvin. The “which blair project”: a quick visual method for evaluating perceptual color maps. In *Proceedings of Visualization*, pages 183–556, 2001.
- [36] Y. S. Ryu, B. Yost, G. Convertino, J. Chen, and C. North. Exploring cognitive strategies for integrating multiple-view visualizations. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 47, pages 591–595, 2003.
- [37] T. Schultz and G. L. Kindlmann. Superquadric glyphs for symmetric second-order tensors. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1595–1604, 2010.
- [38] H.-J. Schulz, T. Nocke, M. Heitzler, and H. Schumann. A design space of visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2366–2375, 2013.
- [39] S. M. Smith, M. Jenkinson, H. Johansen-Berg, D. Rueckert, T. E. Nichols, C. E. Mackay, K. E. Watkins, O. Ciccarelli, M. Z. Cader, P. M. Matthews, et al. Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *Neuroimage*, 31(4):1487–1505, 2006.
- [40] S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, et al. Advances in functional and structural mr image analysis and implementation as fsl. *Neuroimage*, 23:S208–S219, 2004.
- [41] J. L. Stein, S. E. Medland, A. A. Vasquez, D. P. Hibar, R. E. Senstad, A. M. Winkler, R. Toro, K. Appel, R. Bartecek, Ø. Bergmann, et al. Identification of common variants associated with human hippocampal and intracranial volumes. *Nature Genetics*, 44(5):552–561, 2012.
- [42] P. M. Thompson, O. A. Andreassen, A. Arias-Vasquez, C. E. Bearden, P. S. Boedhoe, R. M. Brouwer, R. L. Buckner, J. K. Buitelaar, K. B. Bulaeva, D. M. Cannon, et al. Enigma and the individual: Predicting factors that affect the brain in 35 countries worldwide. *NeuroImage*, 2015.
- [43] J. D. Van Horn and A. W. Toga. Human neuroimaging as a “big data” science. *Brain Imaging and Behavior*, 8(2):323–331, 2014.
- [44] S. Wakana, H. Jiang, L. M. Nagae-Poetscher, P. C. Van Zijl, and S. Mori. Fiber tract-based atlas of human white matter anatomy. *Radiology*, 230(1):77–87, 2004.
- [45] C. Wang and J. Tao. Graphs in scientific visualization: A survey. In *Computer Graphics Forum*, 2016.
- [46] C. Weigle and D. Banks. A comparison of the perceptual benefits of linear perspective and physically-based illumination for display of dense 3D streamtubes. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1723–1730, 2008.
- [47] C. Whelan et al. Structural brain alterations in epilepsy: A coordinated case-control analysis of 3,876 individuals via the ENIGMA consortium. *In submission*, 2016.
- [48] W. Willett, J. Heer, and M. Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1129–1136, 2007.
- [49] G. Zhang, P. Kochunov, E. Hong, H. Carr, and J. Chen. Towards visual mega- and meta-analysis of voxel-based measurement in brain cohorts. In *Eurographics Visualization Symposium*, 2016.
- [50] D. Zielasko, B. Weyers, B. Hentschel, and T. W. Kuhlen. Interactive 3D force-directed edge bundling. *Computer Graphics Forum*, 35(3):127–134, 2016.