NeuroVault.org: A web-based repository for collecting and sharing unthresholded statistical maps of the human brain

Krzysztof J. Gorgolewski[1,6], Gael Varoquaux[2], Gabriel Rivera[3], Yannick Schwarz[2], Satrajit S. Ghosh[4], Camille Maumet[5a], Thomas E. Nichols[5b], Russell A. Poldrack[6], Jean-Baptiste Poline[7], Tal Yarkoni[8], Daniel S. Margulies[1]

1 Max Planck Research Group for Neuroanatomy and Connectivity, Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

2 INRIA Parietal, Neurospin bat 145, CEA, Saclay, 91191, Gif sur Yvette, France

3 InfoCortex UG, Frankfurt am Main, Germany

4 McGovern Institute for Brain Research, Massachusetts Institute of Technology, MA, USA

5a Warwick Manufacturing Group, University of Warwick, Coventry CV4 7AL, UK

5b Department of Statistics, University of Warwick, Coventry CV4 7AL, UK

6 Department of Psychology, Stanford University, Stanford, CA, USA

7 Helen Wills Neuroscience Institute, University of California at Berkeley, CA, USA

8 Department of Psychology, University of Texas at Austin, TX, USA

Contact: krzysztof.gorgolewski@gmail.com

Running title: NeuroVault.org - a database of statistical maps

Abstract

Here we present NeuroVault — a web based repository that allows researchers to store, share, visualise, and decode statistical maps of the human brain. NeuroVault is easy to use and employs modern web technologies to provide informative visualisation of data without the need to install additional software. In addition it leverages the power of the Neurosynth database to provide cognitive decoding of deposited maps. NeuroVault is also a resource for researchers interested in conducting meta- and coactivation analyses. All of the data is exposed through a public REST API enabling other services and tools to take advantage of it.

Introduction

Noninvasive neuroimaging tools such as MRI and PET have enabled unprecedented insight into the localization of various functions in the human brain. As the number of studies using such techniques continues to grow exponentially, the challenge of assessing, summarizing, and condensing their content poses ever-greater difficulty. Even though a single study can take years to conduct, cost hundreds of thousands of dollars, require the effort of dozens of highly trained scientists and is possible only due to volunteered time of many more participants, the output usually is reduced to an academic article, and the original data are rarely shared (Poline et al., 2012). Unfortunately due to the historical legacy of how knowledge is disseminated, the reporting to results in research articles consists mostly of subjective interpretation of data with very little machine readable information. The introduction of common stereotaxic spaces (such as Talairach and MNI305) has provided a framework for reporting of activation locations that has become a community standard in human neuroscience, and enabled meta-analyses, but it still leaves a number of unresolved issues. Peak coordinates are not able to fully describe the 3D shape and extent of a suprathreshold volume on a statistical map. Many papers use figures (2D

or 3D) to present these statistical maps, but authors must make decisions limiting which aspects of the 3D data cube to show. To fully explore the data one would need to be able to interrogate it in an interactive fashion. Furthermore, published figures are not machine readable which forces researchers interested in comparing their own results with published literature to reconstruct regions of interest (ROIs) based on figures from previous work by hand using spheres placed at activation locations.

Additionally it is difficult to put one's results in the context of other studies. The overwhelming number of brain imaging results published each year makes manual comparison both unfeasible and prone to bias. There are attempts to automatically aggregate knowledge across large sets of neuroimaging studies. For example Neurosynth (Yarkoni et al., 2011) is a meta-analysis database that collects coordinates of activation foci from published papers and generates topic maps based on the spatial distribution of those coordinates. Such maps can aid in interpretation of new results. However, comparing new result to a set of topic maps has so far not been implemented in a user friendly fashion.

Furthermore making meta-analytic inferences using only peak coordinates (or statistically thresholded maps) is problematic. It is easy to imagine a subthreshold effect that is consistent across many studies. Such effect would not be picked up by existing meta-analysis methods (Laird et al., 2005; Yarkoni et al., 2011) because it will never be reported in the tables of peak coordinates. Considering how underpowered most human neuroscience studies are, this situation is not that unlikely. Additionally discarding information that is below threshold is akin to not publishing null results which creates a publication bias skewing our perception of accumulated knowledge.

Using full unthresholded statistical maps, instead of solely peak coordinates, would provide a significant advance in meta-analytic power. Coordinate-based meta-analysis (CBMA) methods show only modest overlap with image-based meta-analysis (IBMA; meta-analysis based on unthresholded statistical maps) methods and are less powerful (Salimi-Khorshidi et al., 2009). However, IBMA methods struggle with the source of data. Peak coordinates are easier to obtain because coordinate tables are part of the paper they are easy to share, whereas very few papers provide unthresholded statistical maps (usually by an ad hoc means such as the author's web site).

NeuroVault.org is an attempt to solve these problems. It is a web-based repository that makes it easy to deposit and share statistical maps. It provides attractive visualization and cognitive decoding of the maps, which can improve collaborations and readability of the results. At the same time it also provides an API for methods researchers to download the data, perform powerful analyses, or build new tools.

Results

Platform

One of the key features of NeuroVault is the ease of uploading and sharing of statistical maps. Figure 1 presents a schematic overview of the platform. After logging in, users can upload a broad range of neuroimaging images and associated metadata. These data are then immediately accessible (subject to user-controlled privacy settings) via both an interactive HTML-based interface, and a comprehensive REST API that facilitates programmatic interoperability with other resources. In the following sections, we discuss different aspects of the platform.

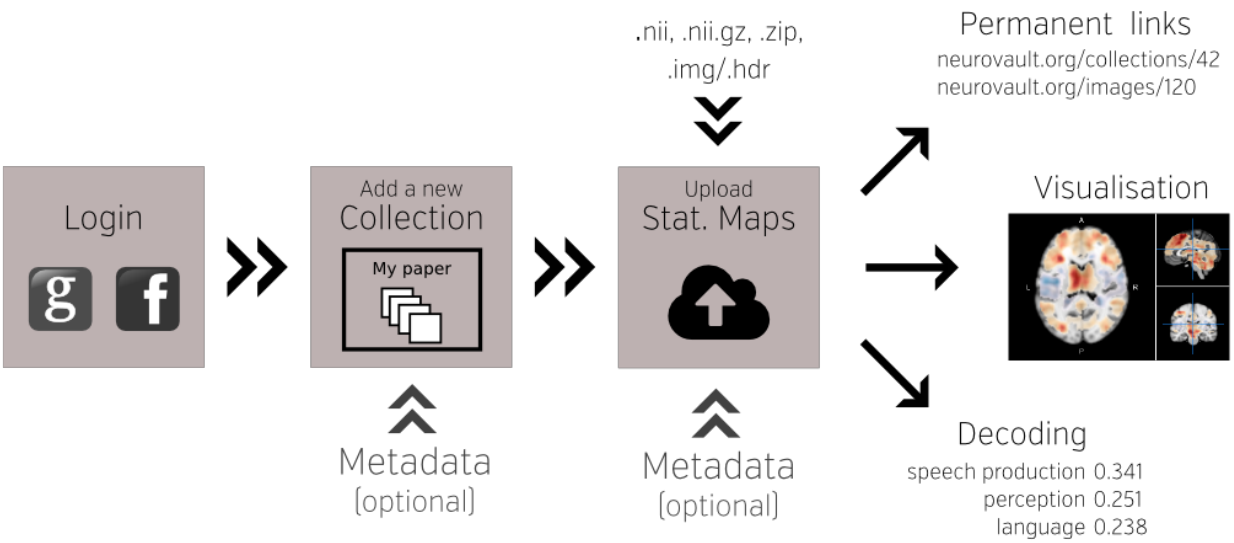

Figure 1. Schematic overview of the NeuroVault platform. To begin working with NeuroVault users are asked to create an account or log in using their Facebook or Google account. After login, the user creates a collection (representing a paper or a study). At this stage users can provide a DOI pointing to a paper associated with the collection and/or fill a number of fields describing the study (see Supplementary Materials for details). This additional information is, however optional. After the collection is created users can add images. This can be done one by one or in bulk by uploading whole folders. Again there is an option to add more metadata describing the images. The process of creating a collection and uploading statistical maps to NeuroVault.org takes only 5-10 minutes. When the maps are uploaded users can start benefiting from permanent link to their results, interactive web-based visualisation, and real-time image decoding.

Image upload. The NeuroVault upload process emphasizes speed and ease of use. Users can rely on existing social media accounts (Google or Facebook) to log in, and can upload individual images or entire folders (see Figure 1). Users are free to arrange their maps into collections or to group them with tags. Each collection and statistical image in NeuroVault gets a permanent link (URL) that can be shared with other researchers or included in papers or other forms of publication (blogs, tweets etc.). Users can specify whether each collection is public or private.

The latter have a unique obfuscated URL that is not discoverable on the NeuroVault website, and thus are accessible only by whomever the owner decides to share their URL with. This gives users freedom to decide who can access their data and can facilitate a scenario in which a collection is shared privately during the pre-publication peer review process, and then made public upon acceptance of a manuscript. Using a third party (such as NeuroVault) to share data that are part of the peer review process eliminates concerns about the reviewers' anonymity. Even though we opted to minimize the required amount of metadata for collections and statistical maps (to streamline the process) we give users an option to provide more information if the user wants to maximise the usability of their maps (see Supplementary Tables metadata_collections and metadata_images). Most importantly we provide ability to link a collection to a paper via a DOI which can promote the associated paper and facilitate meta-analysis.

Data types. NeuroVault is able to handle a plethora of different types of brain maps as long as they are represented as 3D NIFTI files in MNI space. This includes Z or T maps derived from task based, resting state fMRI and PET experiments as well as statistics derived from analyses of structural data (e.g., Voxel Based Morphometry; VBM). In addition results from electroencephalography (EEG) and magnetoencephalography (MEG) experiments can also be used with NeuroVault as long as they are converted to NIFTI volumes through source localization (Phillips et al., 2002). NeuroVault can also handle mask files (for describing ROIs), label maps (a result of parcellation studies), posterior probability maps (coming from Bayesian methods; (Woolrich et al., 2004)), weight maps (coming from multivariate pattern analysis methods; (Haxby, 2012)), and group-level lesion maps (from clinical studies).

User interface. NeuroVault is designed to provide intuitive, interactive visualization of uploaded images. Each image is assigned its own unique URL with an embedded JavaScript 2D/3D viewer. In contrast to traditional, static figures in published articles, users can dynamically interact with images — adjusting statistical thresholds, selecting different color maps, and loading additional brain volumes into the viewer for comparison. Using two embedded open-source JavaScript viewers (Papaya - https://github.com/rii-mango/Papaya and pycortex - https://github.com/gallantlab/pycortex), users can interrogate the data both in the volumetric space as well as on the surface (see Figure 2). Both viewers work inside modern web browsers and do not require any additional software to be installed. In addition to the visual representation of the volume, each page also displays any metadata associated with that image (e.g., experimental contrast, statistic type, etc).

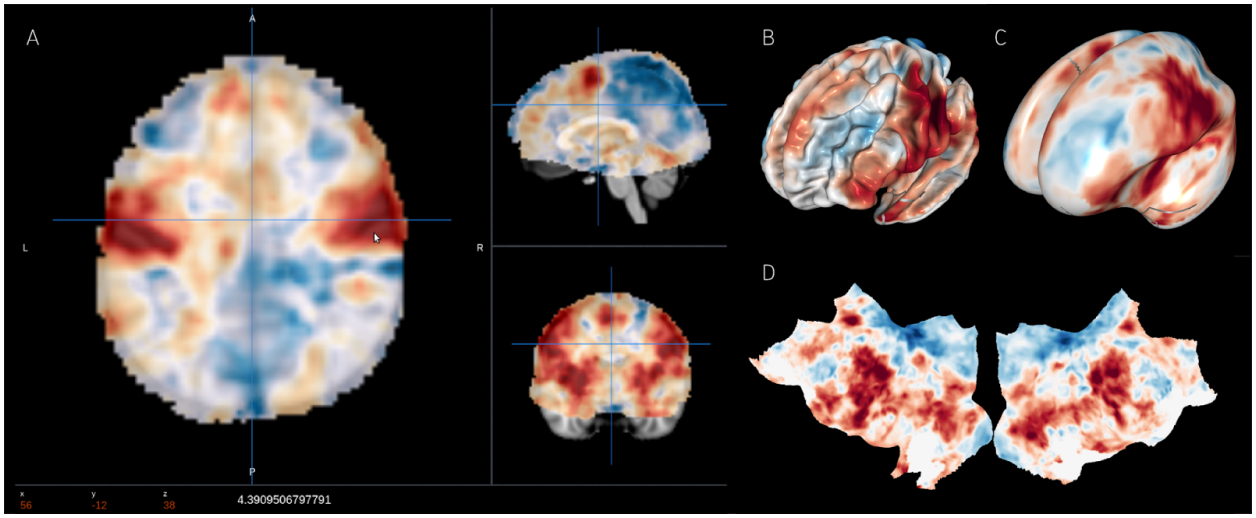

Figure 2. Visualisation options available in NeuroVault. User can choose to interactively interrogate the images using 2D volumetric view (A), 3D fiducial view (B), 3D inflated view (C) or a flattened cortical surface map (D).

Interoperability. A major design goal of NeuroVault is to directly interoperate with other existing web-based neuroimaging resources, ensuring that users can take advantage of a broad range of

computational tools and resources without additional effort. There are two components to this. First, in cases where other relevant resources implemented a public API, NeuroVault can provide a direct interface to those resources. For example, at the push of a single button, each map deposited in NeuroVault can be near-instantly “decoded” using Neurosynth (see Figure 3). In the space of 1 - 2 seconds, the uploaded image is analysed for its spatial correlation with a subset of the concept-based meta-analysis maps in the Neurosynth database. The user is then presented with a ranked, interactive list of maximally similar concepts, providing a quantitative, interactive way of interpreting individual statistical images that is informed by a broader literature of nearly 10,000 studies. Second, NeuroVault exposes its own public REST API, which provides fully open programmatic access to all public image collections and enables direct retrieval of images and associated metadata. This feature allows other researchers to leverage NeuroVault data in a broad range of desktop and web applications. To maximize the impact of data stored in NeuroVault the access to the API is unrestricted, does not require any terms of use agreements, and the data itself is distributed under the CC0 license (http://creativecommons.org/publicdomain/zero/1.0).

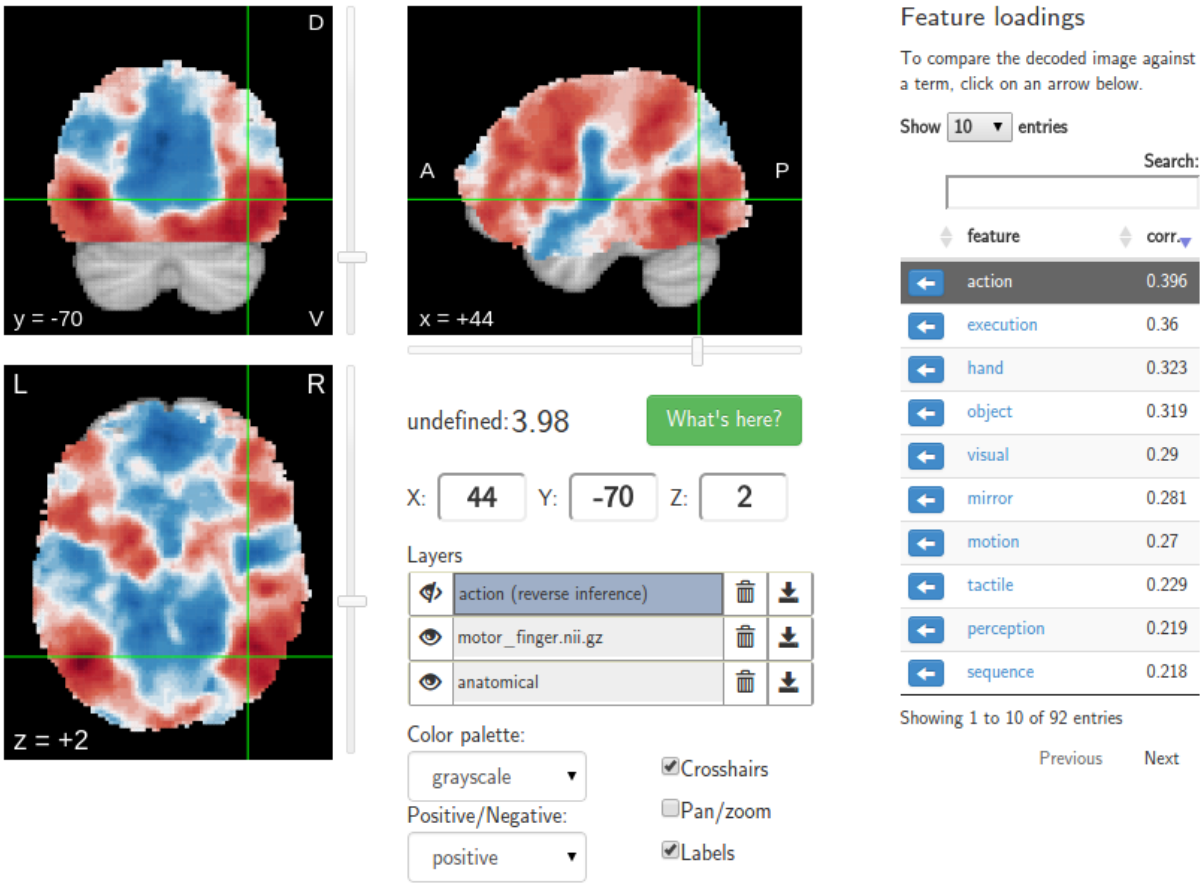

Figure 3. Results of the Neurosynth decoding of a statistical map obtained through NeuroVault API. Users are able to interactively compare their maps with Neurosynth topic maps.

Accessibility. Another advantage of depositing statistical maps in NeuroVault is the increase in longevity and impact of one's research outputs. By providing a free, publicly accessible, centralized repository of whole-brain images, NeuroVault has the potential to increase the flow of data between different researchers and lab groups. Maps deposited in NeuroVault can be used by other researchers to create detailed region of interests for hypothesis driven studies or to compare results of replications. However, one of the most interesting reuse of statistical maps from previous studies is image-based meta-analysis. Even though collection of rich metadata and annotating studies with terms from existing cognitive and experimental ontologies is not the

main purpose of NeuroVault, researchers conducting meta-analyses can use manual (BrainMap), crowd-sourced (BrainSpell) or automatic (Neurosynth) procedures to extract metadata from the papers associated with statistical maps through a DOI. Below we present a proof of concept meta-analysis based on NeuroVault data collected to date. It gives a taste of the potential this platform provides for aggregating knowledge about the human brain.

Meta-analysis using the NeuroVault data

At the time of submitting this publication there were 67 collections (29 of them associated to a publication, 10 private) that consisted of 406 images. We have selected 284 maps to perform proof of concept analyses. The goal of the analyses is to determine whether similar results can be obtained using a limited set of unthresholded maps as compared with large coordinate-based databases. The analyses focus on three aspect: (i) spatial distribution of activations (ii) term-based maps (iii) Independent Component Analysis (ICA) coactivation across maps. Code for this analysis is available at https://github.com/NeuroVault/neurovault_analysis.

Spatial distribution of activations

Figure 4 (middle) shows the spatial distribution of activations for the maps in NeuroVault, depicting the number of times a voxel appeared in a statistical map with a Z- or a T-Statistic greater than 3. Such data reduction simulates analysis conducted on peak coordinates. The distribution is strikingly non-uniform throughout gray matter. In particular, the regions most represented are the frontal part of the insula and dorsal anterior cingulate cortex that form a well-known cingulate-insulate control network associated with salience processing (Seeley et al., 2007). The other structures highlighted in Figure 4 are the inferior parietal sulcus — regions sometimes called the "task-positive network" (Fox et al., 2005) — as well as the occipital lobe, encompassing the visual cortex. The presence of the latter likely reflects the fact that the

majority of experiments rely on visual stimuli. Interestingly, the networks that are most prominent on this map are largely related to attention and executive control.

For comparison we also show a similar map computed from a coordinate-based meta-analysis database, NeuroSynth, that collects coordinates of activation foci from the literature (Figure 4, top). It displays a similar density of activation, with visible attentional networks. It is worth noting that other studies have also reported similar activation density maps (Nelson et al., 2010). However, the visual cortex is much less present in the Neurosynth map comparing to the one based on thresholded NeuroVault data. This could potentially be explained by the fact that researchers tend to report coordinates of high-level contrasts that cancel out low-level effects of stimuli, whereas the NeuroVault database (in its current state) contains a variety of contrasts, including task-versus-baseline maps.

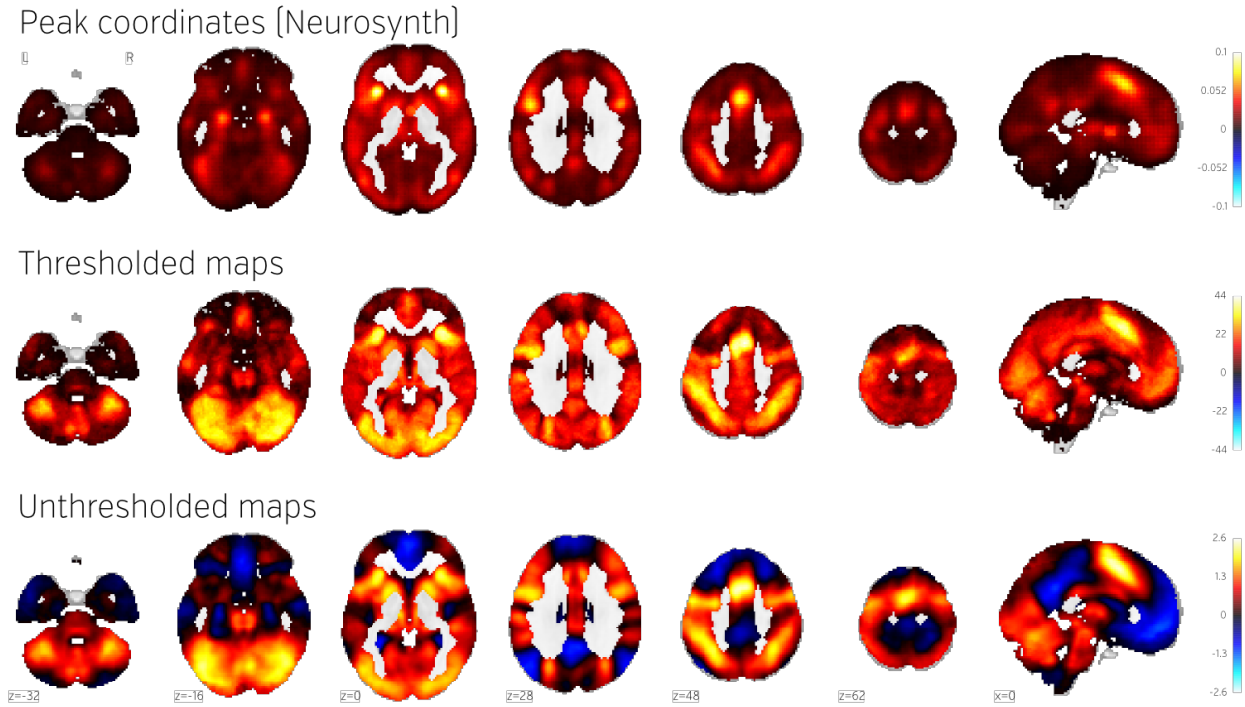

Figure 4. Comparison of frequency of activation across human brain studies obtained using different methods. In top row a probability map obtained from coordinate-based meta-analysis using

NeuroSynth. Middle row a frequency map of voxels of T or Z value > 3 across statistical maps deposited in NeuroVault. Bottom row mean of all T and Z maps (also deposited in NeuroVault).

From a global perspective, to understand better the spatial distribution of detections, it is also interesting to look at the average activation across the entire database. In Figure 4 (bottom) we give the overall average of the statistical maps.

Unlike a simple count of statistically significant detections, as in a coordinate-based meta-analysis, this analysis also captures the dominant sign of the activation, accumulating power in regions that may not cross threshold in analyses from individual studies (note that doing a principled statistical inference, e.g. computing a p-value or a posterior from this heterogeneous collection of maps require methodological developments outside of the scope of this article). For example, the average unthresholded maps clearly show regions that respond, on average, by deactivation, rather than activation. These span the default-mode network, which was historically discovered in a similar analysis, through observation of consistent decreases in activity across a variety of tasks (Shulman et al., 1997).

Conducting term-based meta-analysis of activations

The accumulation of results of brain mapping experiments can be used to run a meta-analysis, pooling results across different experiments on the same topic to define a map related to a given concept that is independent of the specific experimental details (Costafreda, 2011; Somorjai, 2001).

One challenge to running a meta-analysis on data hosted on NeuroVault is that, unlike BrainMap or NeuroSynth, the images do not come with explicit labels describing their content. We use a simple heuristic to assign labels to images: for each image, we look at its metadata (name,

description, and contrast details if available). We assign a label to an image if its metadata contains at least one word from an associated list (see Table 1).

Label	Associated words
language	semantic, linguistic, language, word, words, reading, verb, voice
audio	audio, auditory, audition, listening
motor	motor, button, hand
visual	face, imagery, scrambled, checkerboard, color, visual, visually

Table 1. Labels and associated words.

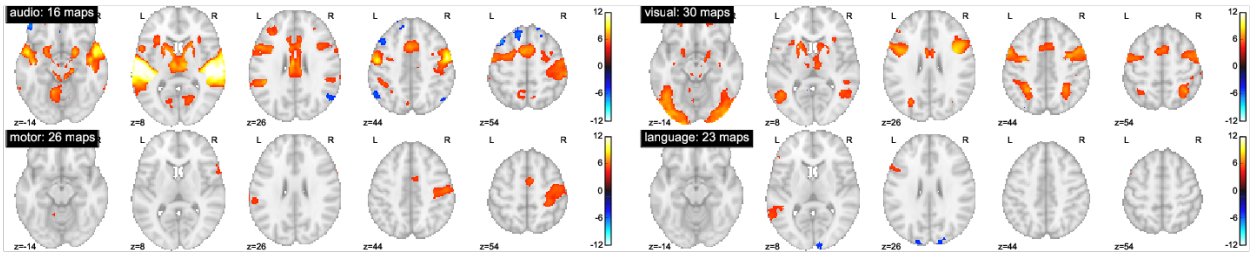

Figure 5. Term-based map for entries of the database: T-map contrasting maps containing a given term with the other maps of the database. Maps were thresholded at T > 5, which corresponds to an family wise error control of p<0.05 with Bonferroni correction.

Figure 5 shows for each label a statistical ‘3rd level’ analysis of all the entries of the database that contains the corresponding terms. As in (Yarkoni et al., 2011), we have found that to retrieve functionally-selective maps, it is important to define a proper baseline: asking whether a voxel is significantly correlated to a term is bound to give imprecise maps as with large-enough sample sizes the full brain becomes significantly correlated (Thyreau et al., 2012). Thus we formulate the problem as follows: is the activation of a given voxel different for the terms of interest than for the rest of the database. We can see that this very rough meta-analysis does capture some meaningful information. The activations related to "auditory" terms highlights predominantly the

auditory cortex; the “visual” terms activate not only the visual system, but also the dorsal-attentional system related to visuo-spatial attention. The corresponding map for the language terms is on the low-end of statistical power but highlights Broca’s and Wernicke’s areas. Such an analysis framework opens the door to principled reversed inference (Poldrack, 2006). Indeed, with a comprehensible-enough database it would be possible to quantify functional specificity of a given activation by testing which terms are significantly associated with the activation of a voxels, rather than which voxels are activated given a term (Yarkoni et al., 2011).

Co-activation maps with independent component analysis

Another approach to uncover large-scale brain organization from a meta-analysis database is to study patterns of coactivation: mapping regions that are often activated together. In this regard, ICA is an interesting data-driven method that can unmix throughout the database different cognitive components that are present entangled in each individual experiment. Here we apply this approach to the NeuroVault database, using a similar strategy as (Smith et al., 2009) applied to the BrainMap database. The cognitive content of the ICA maps is reflected by the unmixing matrix, that gives the loadings of each component on the original maps. On the BrainMap database, these maps have manually-assigned cognitive labels. In NeuroVault, we use the decoding functionally of NeuroSynth to assign such labels.

We use the following strategy to retrieve the cognitive content that these networks capture. First we use the decoding functionality of NeuroSynth to associate cognitive weights for each term in NeuroSynth with each statistical map in Neurovault. The unmixing matrix estimated by ICA is used to transform these map-level loadings to loadings for each network. On Figure 6, we

represent 20 co-activation networks extracted via such an analysis on the 284 maps from NeuroVault, each labeled with four terms most present in the corresponding loadings.

Note that this analysis reproduces the (Smith et al., 2009) analysis, but unlike the original analysis it does not rely on manually labeled and curated data but uses fully automated extraction of information.

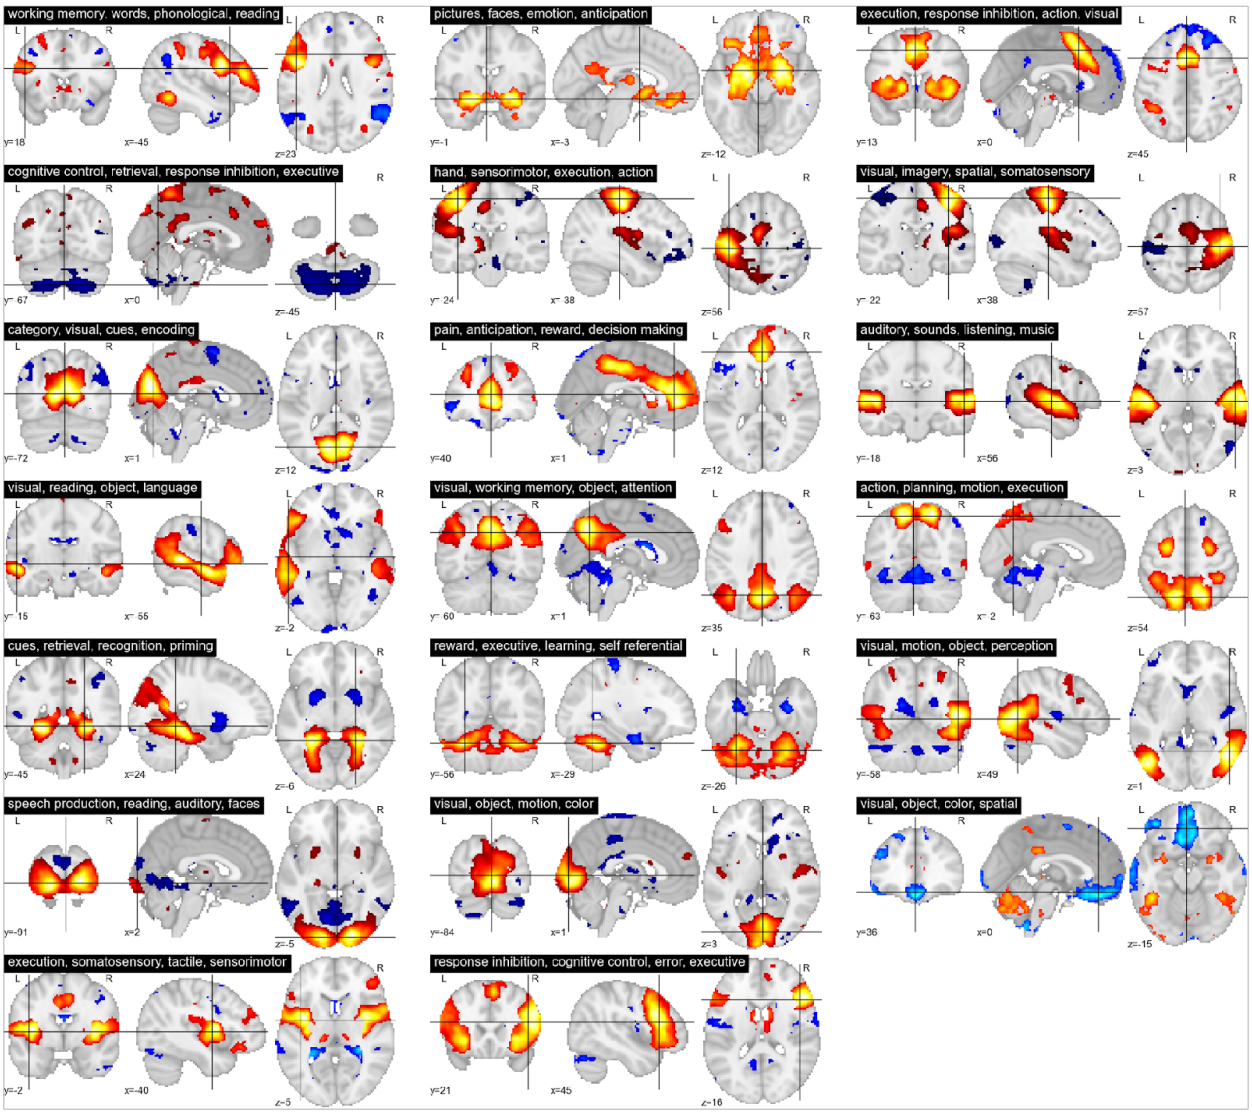

Figure 6. ICA networks extracted from the NeuroVault database. The labels of the networks are automatically computed from the terms decoded on the NeuroVault maps. Note that here we report all the components, unlike in most ICA analyses.

Discussion

We present NeuroVault, a web based platform that allows researchers to store, share, visualise and decode maps of the human brain. This new resource can improve how human brain mapping experiments are presented, disseminated and reused. Due to its web-based implementation NeuroVault does not require any additional software to be installed and thus is very easy to use.

One of the biggest challenges of data sharing platforms is sustainability. Users contributing their data trust that they will be available over an extended period of time. We cannot make any certain claims about the future, but we designed the service in a way that increases its robustness. First of all NeuroVault is an open source project (the code is available at https://github.com/NeuroVault/NeuroVault) that is dependent only on free and open source components (web servers, content management systems, databases etc.). This means that at any given point if need arises any individual with minimum web admin experience can set up NeuroVault to run on a new server. Software is not, however the most important part of the project. To preserve the data we are performing daily offsite backups which are later copied to other locations. The procedure of restoring the service from scratch using the freely available code combined with the backups has been heavily tested. The last component of the service reliability is hardware. It is worth noting that statistical maps take considerably less space than other types of data such as raw fMRI datasets. A 500 Gb hard drive (available for 50$) can store almost 500,000 statistical maps. Furthermore the cost of server maintenance and the uplink to

the Internet can easily be leveraged by existing academic institutions' infrastructures. In short we argue that even though no one is able to guarantee long term availability of NeuroVault, due to the nature of its design and data it is dealing it is easy and cheap to maintain or host at a new location, given there is enough interest and the service will prove to be useful to the scientific community.

NeuroVault is not only a helpful tool for researchers who wants to share, visualise and decode their maps, it is also a resource for researchers wanting to perform meta- and coactivation analyses. Thanks to the public REST API and the CC0 licensing of the data there are no restrictions in terms of how and by whom the data can be used. We hope that this will accelerate progress in the field of human brain imaging. There are many services that could benefit from interaction with NeuroVault. Neurosynth and BrainMap can boost the power of their meta-analyses by working with unthresholded maps stored in NeuroVault instead of peak coordinates extracted from papers. In our analyses we have showed promising results (replication of Neurosynth frequency map, DMN deactivation and ICA topic maps similar to (Smith et al., 2009)) even with the initial, limited in numbers and heterogenous set of maps. Those results were obtained using only a few hundred maps but they nonetheless replicate results from much bigger (but coordinate based) databases (BrainMap and NeuroSynth cover respectively 2500 and 9000 papers). We are confident that increased amount of data will lead to discovering new organizational principles of brain function.

The sharing of neuroimaging data can potentially raise ethical issues related to subject confidentiality (Brakewood & Poldrack, 2013), but sharing of derived data, such as group contrast maps, greatly reduces this concern. Nonetheless, mechanisms are necessary to ensure that uploaded files will not contain personal information, or information permitting

re-identification of subjects. The individual uploading data must take responsibility for this, by agreeing with the data agreement and certifying that he or she has both institutional right and ethical permission to share these derived data. As NeuroVault is mainly focused on group data analyses, this is little chance that personal information will be included and lead to ethical issues, but the platform allows single subject analysis to be uploaded which necessitates safeguards.

To keep the amount of effort needed to create a new collection as minimal as possible, the addition of annotated metadata is optional in NeuroVault. Nevertheless, at the users' discretion, a rich set of metadata can be manually included and stored with the statistical maps. We envision that, in the future, more and more machine-readable information will be shared and these metadata will be populated automatically increasing the potential re-use of the datasets hosted at NeuroVault. Current efforts, including Web pages such as BrainSpell.org, can aid the process of annotating papers (and their corresponding maps) through crowdsourcing. Ideally, machine-readable metadata would be made available directly by the software packages used to generate the statistical maps. The NeuroImaging Data Model (NIDM) (Keator et al., 2013) is an example of metadata standards that could be used to decrease metadata loss between the analysis and the upload of the statistical maps into NeuroVault. Detailed annotation can enable more interesting analyses. For example one can imagine a study trying to train a classifier predicting from the activation maps who performed the study, what software was used, what scanner was used etc. NeuroVault is part of a growing ecosystem of web enabled tools that enables researchers to aggregate and analyse results from many human brain mapping experiments.

It is also worth pointing out that NeuroVault is not only supporting task based fMRI results. Results from resting state fMRI, PET, VBM, DWI, and most interestingly source reconstructed

EEG/MEG experiments can be used with the platform as long as they are NIFTI files in MNI space. We plan to expand this to FreeSurfer surfaces, CIFTI files, and connectomes in the near future. So far aggregating results across modalities has been difficult, and we hope that this platform can to some extent improve this situation.

We are also in the process of integration with the Resource Identification Initiative through The Neuroscience Information Framework (NIF, http://neuinfo.org/). This interdisciplinary project assigns identifiers to resources and tools used in research that are then included in publications and later indexed by Google Scholar and PubMed. This way PubMed LinkOut service (http://www.ncbi.nlm.nih.gov/projects/linkout/) will automatically include links to statstical maps on PubMed web pages describing relevant papers. Assigning resource identifiers to statistical maps would also allow their creators to track how they are used and obtain academically acknowledge credit (even in case when the maps come from unpublished studies). In addition we aim to capitalize on the NIDM standard for representing data analysis results. This will allow tight integration with existing neuroimaging software (such as FSL, SPM, or AFNI). Providing a single click solution for uploading maps to NeuroVault that would be available within analysis software could greatly improve usability of the service.

Conclusion

In this work we have described NeuroVault — a web based repository that allows researchers to store, share, visualise, and decode unthresholded statistical maps of the human brain. This project not only helps individual researchers to disseminate their results and put them in the context of existing literature, but it also enables aggregation of data across studies. Through our analyses we have shown that with only a few hundred statistical maps we can achieve results comparable to those obtained with thousands of sets of coordinates. NeuroVault is free and

unencumbered by data use agreements. The data is available and the database queryable via the web interface and REST API. This opens the door to developing novel method to draw inferences from a meta-analytic database.

Acknowledgements

This work was partially funded by the National Institutes of Health (NIH), R01MH096906 [TY], NSF OCI-1131441 [RP], International Neuroinformatics Coordinating Facility (INCF) and the Max Planck Society [KJG, DSM].

We thank the INCF Neuroimaging Data Sharing task force members for their input during several discussions.

References

Costafreda, S. G. (2011). Meta-Analysis, Mega-Analysis, and Task Analysis in fMRI Research. *Philos. Psychiatr. Psychol.* 18, 275–277.

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., and Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc. Natl. Acad. Sci. U. S. A.* 102, 9673–9678.

Haxby, J. V. (2012). Multivariate pattern analysis of fMRI: the early beginnings. *Neuroimage* 62, 852–855.

Keator, D. B., Helmer, K., Steffener, J., Turner, J. A., Van Erp, T. G. M., Gadde, S., Ashish, N., Burns, G. A., and Nichols, B. N. (2013). Towards structured sharing of raw and derived neuroimaging data across existing resources. *Neuroimage* 82, 647–661.

Laird, A. R., Lancaster, J. J., and Fox, P. T. (2005). BrainMap. *Neuroinformatics* 3, 65–77.

Nelson, S. M., Dosenbach, N. U. F., Cohen, A. L., Wheeler, M. E., Schlaggar, B. L., and Petersen, S. E. (2010). Role of the anterior insula in task-level control and focal attention. *Brain Struct. Funct.* 214, 669–680.

Phillips, C., Rugg, M. D., and Friston, K. J. (2002). Anatomically informed basis functions for EEG source localization: combining functional and anatomical constraints. *Neuroimage* 16, 678–695.

Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends*

Cogn. Sci. 10, 59–63.

Poline, J.-B., Breeze, J. L., Ghosh, S., Gorgolewski, K., Halchenko, Y. O., Hanke, M., Haselgrove, C., Helmer, K. G., Keator, D. B., Marcus, D. S., et al. (2012). Data sharing in neuroimaging research. *Front. Neuroinform.* 6, 9.

Salimi-Khorshidi, G., Smith, S. M., Keltner, J. R., Wager, T. D., and Nichols, T. E. (2009). Meta-analysis of neuroimaging data: a comparison of image-based and coordinate-based pooling of studies. *Neuroimage* 45, 810–823.

Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., Reiss, A. L., and Greicius, M. D. (2007). Dissociable intrinsic connectivity networks for salience processing and executive control. *J. Neurosci.* 27, 2349–2356.

Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., and Petersen, S. E. (1997). Common Blood Flow Changes across Visual Tasks: II. Decreases in Cerebral Cortex. *J. Cogn. Neurosci.* 9, 648–663.

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R., et al. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. U. S. A.* 106, 13040–13045.

Somorjai, R. L. (2001). Meta brain, meta experiment - simultaneous analysis of fMRI data across multiple subjects and/or experiments. *Neuroimage* 13, 253.

Thyreau, B., Schwartz, Y., Thirion, B., Frouin, V., Loth, E., Vollstädt-Klein, S., Paus, T., Artiges, E., Conrod, P. J., Schumann, G., et al. (2012). Very large fMRI study using the IMAGEN database: Sensitivity–specificity and population effect modeling in relation to the underlying anatomy. *Neuroimage* 61, 295–303.

Woolrich, M. W., Behrens, T. E. J., Beckmann, C. F., Jenkinson, M., and Smith, S. M. (2004). Multilevel linear modelling for FMRI group analysis using Bayesian inference. *Neuroimage* 21, 1732–1747.

Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., and Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nat. Methods* 8, 665–670.