Multi-site meta-analysis of image-wide genome-wide associations of morphometry

Neda Jahanshad[1], Gennady Roshchupkin[5], Joshua Faskowitz[1], Derrek P. Hibar[1], Boris A. Gutman[1], Hieab H.H. Adams[3,4], Wiro J. Niessen[5], Meike W. Vernooij[3,4], M. Arfan Ikram[3,4,6], Marcel P. Zwiers[7], Alejandro Arias Vasquez[8], Barbara Franke[8], Alex Ing[9], Sylvane Desrivieres[9], Gunter Schumann[9], Greig I. de Zubicaray[a], Katie L. McMahon[b], Sarah E. Medland[c], Margaret J. Wright[c], Paul M. Thompson[1]

[1]Imaging Genetics Center, Department of Neurology, Keck School of Medicine of USC, Marina del Rey, CA, USA
[3] Department of Epidemiology, [4] Radiology, [5] Medical Informatics, [6] Neurology, Erasmus Medical Center, Rotterdam, NL
[7] Donders Institute, [8] Department of Psychiatry, Human Genetics Radboud UMC, Nijmegen, NL
[9] King's College London, London UK
[b]Center for Advanced Imaging, University of Queensland, Brisbane, Australia
[c]QIMR Berghofer Medical Research Institute, Brisbane, Australia

Abstract (150/150). Large-scale distributed analyses of over 30,000 MRI scans recently detected common genetic variants associated with the volumes of subcortical brain structures. Scaling up these efforts, still greater computational challenges arise in screening the genome for statistical associations at each voxel in the brain, localizing effects using "image-wide genome-wide" testing (voxelwise GWAS, vGWAS). Here we benefit from distributed computations at multiple sites to meta-analyze genome-wide image-wide data, allowing private genomic data to stay at the site where it was collected. Site-specific tensor-based morphometry (TBM) is performed with a custom template for each site, using a multi channel registration. A single vGWAS testing 10^7 variants against 2 million voxels can yield hundreds of TB of summary statistics, which would need to be transferred and pooled for meta-analysis. We propose a 2-step method, which reduces data transfer for each site to a subset of SNPs and voxels guaranteed to contain all significant hits.

Keywords: Neuroimaging genetics, GWAS, meta-analysis, Big Data, multiple comparisons correction, multi-site.

1 Introduction

Imaging genetics is an emerging field in which variations in the human genome are related to brain differences, in an attempt to discover specific genetic variants that affect brain development, connectivity, and risk for disease. Genome-wide association

studies (GWAS) test for statistical associations between brain measures and over a million single nucleotide polymorphisms (SNPs), or base-pair variants, in the genome[1]. To simplify the screening effort, studies often focus on one or a handful of measures extracted from brain scans, such as the overall volume of the hippocampus [1]; a recent study of over 30,000 brain MRI scans identified 8 genetic loci that were consistently associated with intracranial and subcortical structural volumes, in 50 cohorts worldwide [2]. Given these recent successes with simple volumetric measures, there is great interest in screening the image space more deeply. Each image contains many more features, e.g., at individual voxels, allowing better localization of gene effects and their patterns in the brain. Testing effects of $\sim 10^6$ genetic variants at $\sim 10^6$ voxels requires around 10^{12} statistical tests, but recent volumetric associations achieved significance levels of $p<10^{-23}$, suggesting that effects would survive even brute-force Bonferroni corrections for multiple testing. So far, successful GWAS for single traits have required samples of $1x10^4$ to $1x10^6$ individuals to discover and independently replicate statistically significant variants [2]. GWAS tests that screen around a million voxels, are severely underpowered in individual cohorts, due to the large number of tests performed and the stringent statistical criteria needed to establish significance [3]. Initial studies show voxel-wise genome-wide association studies (vGWAS) are feasible, i.e., image-wide genome-wide testing [4-6], but so far these studies have been under-powered to detect true associations in cohorts of ~1,000 individuals. To guard against false positives, it has become standard in genetics to seek replication of results and/or pool data and aggregate evidence across independent cohorts, before associations are considered credible and reproducible. At the same time, privacy requirements governing genomic data, and in some cases also brain scans, may prevent raw data from being transferred and shared. This has led to collaborative efforts using protocol harmonization and meta-analysis to aggregate site-specific statistical results. In addition, as datasets become vast and more numerous, there is some benefit to distributing the computation across sites, and sending the algorithms to the data rather than centralizing all the data. Therefore, approaches are needed to meta-analyze massive amounts of data from a variety of sources, including image-wide statistics.

Several approaches are required to conduct a distributed voxelwise genome-wide search; first, a registration method is needed to map data from multiple cohorts into a single coordinate space. Without such a registration, voxel measures and statistics would not be comparable across cohorts. Second, while dimensionality reduction is commonly applied in neuroimaging studies, there is no strong prior information on which subset of the 10^{12} SNP x voxel tests are more likely to support the strongest association signals. Even if the image features were reduced to several thousands, a GWAS evaluating 10^6-10^7 SNPs yields around 400TB of compressed

[1] At each of these SNP locations, there are two possible nucleotides (or alleles), and each individual has two chromosomes that will carry one variant or the other. Therefore, at each SNP location, and individual will have 0, 1 or 2 copies of the minor allele (which is the term used to refer to the least prevalent variant in the population). When testing for statistical associations using an additive genetic model, each SNP is coded as 0, 1 or 2 in each individual.

data to sort and filter at each site, making data transfer of all sites summary statistics to a centralized site, less than ideal.

While methods for developing optimized single-site vGWAS techniques are also under development [10], our work is applicable to such approaches and focuses on the issues related to the meta-analysis of multi-site vGWAS. Here we develop a multi-site adaptable protocol using freely available and common neuroimaging software for voxelwise volumetric analysis by tensor-based-morphometry (TBM), and we then describe genome-wide association testing in the resulting data. We show it's usefulness in a simple univariate approach to vGWAS, though optimized approaches could benefit similarly. Each site conducts statistical analysis on its own cohort locally, yielding results specific to that cohort. We then show how to map each site's results into a common space based on four large and arguably representative cohorts worldwide; we also address the issue of prioritizing data transfer.

2 Methods

2.1 Harmonizing voxelwise associations for meta-analysis of 7 sites

Data from seven separate cohorts were included in this study, listed in **Table 1.**
Preprocessing: Following the approach in [2], subcortical and cortical segmentations were performed on 3D anatomical brain MRI scans from each site using FreeSurfer. Quality control protocols were implemented to remove poorly segmented images and outliers, using procedures developed and tested by the ENIGMA Consortium [1,2]. To attempt to reproduce previously reported genetic associations with volumetric measures [2], the integer-valued segmentations of these brain structures were included as part of the fidelity term (the image similarity metric) in a multichannel nonlinear registration approach. The resulting tensor-based morphometry (TBM) workflow was implemented across all sites to allow voxel-level inferences about genetic associations with regional brain volumetric differences, determined using TBM.
Site-specific and global minimal deformation templates (MDT): MDTs were constructed using the Advanced Normalization Tools (ANTs; http://stnava.github.io/ANTs/) software package and accompanying scripts. Approximately 30 scans per cohort were used to create each site-specific template. Images were first linearly aligned to a common space consistent with the MNI brain template before SyN [7] non-linear registration was used to obtain deformation fields. To create the multi-channel template, a weight was assigned to each channel, corresponding to the contribution of that channel to the total warp. We set the T1-weighted channel to 1, the cortical ribbon to 0.5 and the subcortical segmentations to 0.2. These parameters resulted in stable MDTs; the resulting MDTs remained robust in cases where poor quality scans were deliberately included.

Four of the cohorts (ADNI-1, the Rotterdam study (RSS), the Queensland study (QTIM) and BIG; see **Figure 1**) - representing two older adult and two younger adult cohorts - were used to create a representative MDT for all the cohorts, again using 3 channels for registration. We chose not to include all sites in the final

template but instead use representative sites with varying imaging parameters and demographics; in practice, new sites will often join an ongoing study and continuously re-establishing a template could be impractical.

All cohort-specific MDTs were then registered to the final MDT in the same manner. These warps were maintained for later pooling of statistical maps to a common space. Two alternate methods for template construction and registration were also evaluated: 1) *Single-channel template and registration.* The same T1-weighed images used for multichannel registration were used, but *without* the added FreeSurfer cortical and subcortical labels. 2) *Registration to MNI.* To eliminate the effect of the specific cohorts in the meta-analysis, it might be suggested that an existing template be used for registrations; to this end, we registered all subjects directly to the MNI atlas. This has the advantage of staying consistent regardless of added subjects or added cohorts; however, use of a single template not drawn from the population may introduce other sources of bias in the maps; this also makes single-site level analysis on the extracted det(Jacobian) maps less practical, and limit the resulting maps to use in multi-site analysis, rather than our proposed approach which will additionally provide a processing stream for site-specific investigations.

2.2 Voxelwise associations on simulated genetic data

This work was motivated by distributed "big data" analysis that can accommodate partially private genomic data. As such, we generated a dataset with simulated genetic effects, including 100,000 data-points per subject of each dataset to represent an additive genetic effect (0,1, or 2 at each "genetic locus") using a 2D multinomial distribution with probabilities set to the minor allele frequency (MAF) and 1-MAF. The MAF was uniformly distributed (as approximated from the publicly available ADNI-2 data) and maintained greater than 0.1 to avoid rare variants (and by definition < 0.5). Files were saved in the .tped format for integration with PLINK software (http://pngu.mgh.harvard.edu/~purcell/plink/).

For each cohort, a univariate GWAS was performed at every individual voxel. To reduce model complexity, covariates including sex and age were removed from the image and vGWAS was run on residual maps using PLINK2. Data was parallelized across 100 processing nodes, and each PLINK run using the --mpheno flag over $N_{voxels}/100$ phenotypes. Each voxel output was 8.5MB in size, generating about 2TB of data for the downsampled image. As this was by a factor of 4 in each dimension, the full-size image would produce about 4^3x2TB of data, or ~128TB.

2.3 Simulating genetic data with predefined volumetric effects

Of the generated SNPs, five were designed to meet certain volumetric summary criteria for each individual cohort. Summary measures were defined extracted and defined according to ENIGMA protocols.

1) SNP with MAF = 0.1 simulated to be marginally (z=1.96) associated with average bilateral thalamic volume (after removing the intracranial volume ICV effect). **2)** Same as 1 but with MAF = 0.3 **3)** SNP with MAF = 0.1 were generated to

have an effect size of z=min(N/10, 5) when regressed with bilateral hippocampal volume (ICV effect removed) such that the significance was related to the cohort size, N, yet was not excessive (|z| < 5). **4)** Same as #3 but with MAF = 0.3 **5)** SNP with MAF=0.3 was set to similarly associate with ICV, a feature intended to be removed from the voxelwise associations, as TBM was performed on images that had been linearly aligned to include scaling after skull stripping, and this effect was not included here.

To enforce these associations, a correlation coefficient was determined from the set Z-statistic (1.96 for **1, 2** above). Using the fact that vectors with mean 0 have corr=cos(*theta*), where *theta* is the angle between them, we centered and orthogonalize the response variable (e.g, HV) with a QR-decomposition and scaled back; as this method does not lead to the integer values 0,1,2 needed, the values were rounded and correlation values were recomputed, and the process was iterated until the final correlation was +/-0.1 of the desired value.

As this work is intended to be a proof-of-concept, we downsampled the images by a factor of 4 for the genome-wide analysis, such that the images contained ~30,000 voxels (this varied slightly by site). The final MDT had 31,725 voxels.

2.4 Reduction of data transfer to eliminate negative exchange

The inverse-variance based aggregate p-value for SNP i at trait (voxel) v, is [8]: $p_{MA-SE}(i,v) = 2\Phi(|-Z_{MA-SE}(i,v)|)$, where Φ represents the normal transformation and

$$Z_{\text{MA-SE}}(i,v) = \frac{\beta(i,v)}{\text{SE}(i,v)} = \frac{\sum_j \beta_j(i,v) \times \text{se}_j^{-2}(i,v) \Big/ \sum_j \text{se}_j^{-2}(i,v)}{\sqrt{1 \Big/ \sum_j \text{se}_j^{-2}(i,v)}}$$

Here $\beta_j(i,v)$ is site j's effect-size and $se_j(i,v)$, its standard error for SNP i and trait v. Statistical significance implies that the p-value is less than a given threshold (p_{cutoff}), or similarly, the magnitude of the statistic, denoted by z, must be greater than a specific threshold ($|Z_{cutoff}|$):

$$|Z_{MA-SE}(i,v)| \geq 0.5\Phi^{-1} p_{cutoff}(i,v) = |Z_{cutoff}(i,v)|.$$

If a SNP (i_k) at a given v_k passes this meta-analytical threshold, then the maximum Z-score for that SNP across all v must also pass the threshold:

$$\underset{v}{\text{argmax}}\left(Z_j(i_k,v)\right) \geq Z_j(i_k,v_k) \equiv \frac{\underset{v}{\text{argmax}}\left(\beta_j(i_k,v)\right)}{\underset{v}{\text{argmax}}\left(se_j^2(i_k,v)\right)} \geq \frac{\beta_j(i_k,v_k)}{se_j^2(i_k,v_k)}$$

Therefore, we can collapse the image localization information to take only the most extreme values for the SNP across the full image. We note that in accordance with the inverse-variance meta-analysis formula, to order to ensure the maximal statistic for each site, we take the Z-statistics to be the most extreme Beta divided by the square of the standard error. If the most extreme +Z-scores for each SNP are taken across all sites, then it can only exceed the significant Z-score at i_k,v_k.

$$Z_{\text{MA-SE}}(i_k) = \frac{\sum_j \underset{v}{\text{argmax}}(\beta_j(i_k,v)) \times \underset{v}{\text{argmin}}(se_j(i_k,v))^{-2} \Big/ \sum_j \underset{v}{\text{argmin}}(se_j(i_k,v))^{-2}}{\sqrt{1 \Big/ \sum_j \underset{v}{\text{argmin}}(se_j(i_k,v))^{-2}}} \geq Z_{\text{MA-SE}}(i_k,v_k)$$

Note positive and negative Z's are possible, so the same must be considered for –Z.

Only SNPs with a meta-analyzed statistic $|Z_{\text{MA-SE}}(i)| \geq Z_{\text{cutoff}}$ will be subjected to full meta-analysis across all voxels.

Voxelwise meta-analysis would in a sense filter for localized false positives; for example, if one SNP shows an effect in one voxel in one cohort, it wouldn't necessarily have a false effect in the same voxel in a different cohort, therefore the meta-analysis would not necessarily show a significant effect. However, when collapsing the image the lack of localization can enhance false positives. Given the large number of voxels present in an image, (1,869,764 in the full resolution image, and 31,725 in the downsampled) using this TBM method, the probability that a SNP will reach the threshold for significance at **any** voxel is high. Therefore, we also demonstrate the effect of this approach in segmented images, and show the effect it has on the reduction in the data required for transfer.

The procedure described here therefore involves the following site-specific steps (for which harmonized scripts would be provided):

1) Creating the cohort-specific template, and defining it's mapping to the overall template (this can be done at the central site, or the template and its mapping to the overall template is sent to the central site)
2) Proposed 3-channel registration of all subjects in cohort to the cohort-specific MDT for TBM analysis
3) Voxelwise GWAS at the site level.
4) Finding the minimum and maximum statistic across all voxels (in the full image or a given parcellation) for each SNP, and sending this information to the central site.

The data transfer then is performed in two steps:

1) Sending these minimum and maximum results (+/-Z)
    a. As data is provided for each SNP, this equates to two full GWAS results to the central site for each parcellation of the image
    b. At this stage -- the central site pools all sites results and determines which SNP set is needed for which parcellation from all individual sites
2) Sending (possibly reduced datafiles) of chosen SNPs for each parcellation of the brain (whole brain or ROIs)
    a. ROIs are delineated on the overall-MDT, and inverse-warps from the cohort MDT are applied to the labels using a nearest neighbor interpolation.

3 Results

3.1 Simulated associations of fixed genetic effects

Figure 2a shows the effect of a single variant with set marginal effects on thalamic volume (SNP 1 above) mapped using multiple possible voxelwise analysis methods. The multichannel approach where the cortical and subcortical volume segmentations were used as added registration channels to help drive registration (MDT creation and intersubject registration) showed visibly greater specificity with thalamic variability. **Figure 2b** shows the effect of the multichannel registration approach when meta-analyzing a fixed SNP effect (SNP 4 above) on a voxelwise level. Voxel level

analysis maintained regional specificity with FDR-significant voxels bilaterally in the hippocampus.

3.2 Data reduction by transferring a reduced set of SNPs

When analyzing the brain in full, and collapsing the image to take the most extreme positive and negative statistic across all voxels for meta-analysis, we find, as may be expected from the multiple comparisons, that this did, though not greatly, reduce the number of SNPs that could potentially survive multiple comparisons correction after meta-analysis.

With a strict Bonferroni correction accounting for all SNPs (100000) and voxels (31,725), 84% of SNPs (or 83954 of the 100000) would be needed, already accounting for an approximate 16% reduction in data transfer.

Dividing the image in to two, the left and right hemisphere, led to an overall reduction to 82495 SNPs across the brain. However, data transfer could further be reduced as only 78041 were found in one hemisphere, therefore for 4454 of the 82495 SNPs, data for ½ the image would not need to be transferred.

Further breaking down the image in to bilateral ROIs, including bilateral subcortical regions such as hippocampus, putamen, etc., as well as cortical parcellations including anterior, posterior temporal/parietal lobes, cingulate gyrus etc., and filling in remaining sections resulted in an even more reduced dataset for transfer. Though a total of 82711 SNPs were still identified, SNPs were contained in small ROIs. Most smaller ROIs were less than 1000 voxels and when their voxels were collapsed held between 0 and a few hundred possibly significant SNPs out of the full 100000. between the ROI with the most amount of potential significant SNPs (51% of the total) for follow up included the bilateral superior frontal gyrus, which also made up about 9% of the image.

Separating the above-mentioned regions into their respective L and R hemispheres (84 total, not listed here for brevity) resulted in a drastically reduced SNP set of only 73% of the total (73250 of the 100000) for the full image. However, once again, as certain SNPs were only significant in certain regions, the transfer of their information from non-significant voxels is not necessary, further reducing the data transfer to less than half of the data generated. The superior frontal gyrus again held the most number of possible SNPs, with 18528 in the left hemisphere and 22389 in the right. Clearly separating this region into L and R already reduced the total number of SNPs by approximately 10% of the total, and 20% of those identified for the region itself. **Figure 3** shows a plot of the number of possibly significant SNPs as a function of ROI size when the ROI is collapsed to the most extreme statistic. Clearly as the number of ROIs increases, the data transfer in the first step also increases, but numbers are orders of magnitude less than if the images were transferred on a voxel by voxel level.

4 Discussion

Here we showed how to extend large-scale meta-analytic genetic association studies to image-wide analyses, by including steps to pool data across templates and make inter-site transfer of data more efficient. A distributed parallel computation can be highly beneficial as cohorts increase in size and add to a study or as new cohorts join an analysis. Due to the high levels of computation, is also not practical to re-run analyses at every site, so our approach makes use of common analyses many cohorts have already performed on structural MRI scans to make volume measurements, leading to a harmonized protocol for voxelwise association studies. While previous voxel-level meta-analyses have been performed, they involve pooling data from published results, which may only highlight the association at specific points in particular populations, and the mapping between regions and the summary statistics are unknown [9]. Using simulated genetic markers, we show our technique can maintain full structure-level volume associations (i.e., effects on hippocampal volume) when mapping out voxel-level associations, not only in a single cohort, but when data is pooled across a diverse set of cohorts.

In our analysis, we generated 100,000 data points to represent genomic markers with varying allelic frequencies. Here, all SNPs were generated independently of the others; therefore we are not evaluating any effect of linkage seen in standard GWAS. An exhaustive search of all SNPs and voxels may be avoidable by using dimensionality reduction methods to both the image and genome. Reduced rank regression, parallel ICA, canonical covariates analysis, and nonlinear machine learning approaches have all been proposed to fit sets of genetic predictors to imaging data. SNP set selection methods, annotation methods, and Bayesian priors have all been proposed to prioritize sets of SNPs in models rather than search them in a bias-free way. However, at the time of writing, reproducible associations have been hard to identify with these methods, while volumetric associations have been reproduced across 50 cohorts using univariate association testing [2]. Hopefully, the inefficiencies in these standard univariate GWAS approaches will be overcome in the future.

Our work here shows however a method applicable to any voxelwise or vertexwise GWAS that is to be extended to a meta analysis. Rather than transferring all data, a multi-stage process can be conducted to break down the image into regions of interest, find the most significant measures for each SNP at those locations, and determine whether they would reach global significance in the meta analysis. Only if so, then more information would be requested. This allows for a significant reduction in data transfer between sites as well as expedited meta-analysis that would not need to read hundreds of TB of text files and store them in memory simultaneously.

References

1. Stein, J.L., et al., *Identification of common variants associated with human hippocampal and intracranial volumes*. **Nature Genetics,** 2012. **44**(5): p. 552-61.
2. Hibar, D.P. et al., *Common genetic variants influence human subcortical brain structures,* **Nature**, Jan. 2015. doi: 10.1038/nature14101.

3. Medland, S.E., et al., *Whole-genome analyses of whole-brain data: working within an expanded search space*. **Nat Neurosci.** 2014 Jun;17(6):791-800.
4. Stein, J.L., et al., *Genome-wide analysis reveals novel genes influencing temporal lobe structure with relevance to neurodegeneration in Alzheimer's disease*. **NeuroImage**. 2010 Jun; 51(2):542-54.
5. Hibar, D.P. et al., *Voxelwise gene-wide association study (vGeneWAS): multivariate gene-based association testing in 731 elderly subjects*. **NeuroImage**. 2011 Jun 15;56(4):1875-91.
6. Vounou, M., et al., *Sparse reduced-rank regression detects genetic associations with voxel-wise longitudinal phenotypes in Alzheimer's disease*. **NeuroImage**. 2012 Mar;60(1):700-16.
7. Avants, B.B., et al., *Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain*. **Med Image Anal.** 2008 Feb;12(1):26-41.
8. Willer, C.J., Y. Li, and G.R. Abecasis, *METAL: fast and efficient meta-analysis of genomewide association scans*. Bioinformatics, 2010. **26**(17): p. 2190-1.
9. Salimi-Khorshidi, G., S. M. Smith, J. R. Keltner, T. D. Wager, and T. E. Nichols, "Meta-analysis of neuroimaging data: a comparison of image-based and coordinate-based pooling of studies," *Neuroimage*, vol. 45, pp. 810-23, Apr 15 2009.
10. Huang M, Nichols T, Huang C, Yu Y, Lu Z, Knickmeyer RC, Feng Q, Zhu H, Alzheimer's Disease Neuroimaging I (2015) FVGWAS: Fast voxelwise genome wide association analysis of large-scale imaging genetic data. Neuroimage.

Table 1 describes the imaging and demographic data from each cohort. Note QTIM and HCP are family based studies, yet only one person per family was included in this study.

	ADNI1	**ADNI2**	**BIG**	**HCP**	**IMAGEN**	**QTIM**	**RSS**
Scanner	GE, Siemens Philips,	GE	Philips	GE	GE, Siemens Philips,	Siemens Brunker	GE
Field Strength	1.5T	3T	1.5 and 3 T	3T	3T	4T	1.5T
Location	US multi-site	US multi-site	Nijmegen, NL	Saint Louis, USA	EU, multisite	Brisbane, AUS	Rotterdam, NL
Voxelsize (mm^3)	1.25×1.25×1.2	1.25×1.25×1.2	1 x1x1	0.7x0.7x0.7	1.1×1.1×1.1	0.9×0.9×0.9	1x1x1.6
N	837	815	62	207	80	590	64
Age	75+/- 6.6 (60-89)	72.8+/- 6.6(48-90)	21.5+/- 1.7 (18-25)	28.7+/- 3.5 (22 – 35)	14 +/-0.4	22.9+/- 2.8 (18-30)	67.49 +/- 11.40

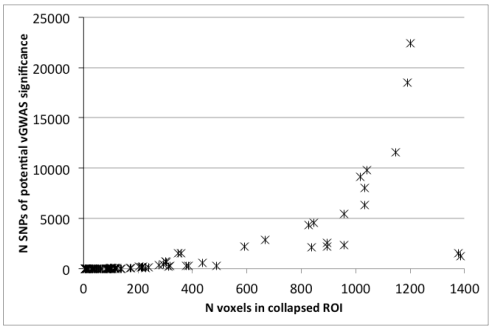

Figure 3 Plot showing the number of potentially significant SNPs meta analyzed in collapse ROIs of a given size. The ~1400 voxel ROIs in the cerebellar region do not follow a similar trend as other brain regions.

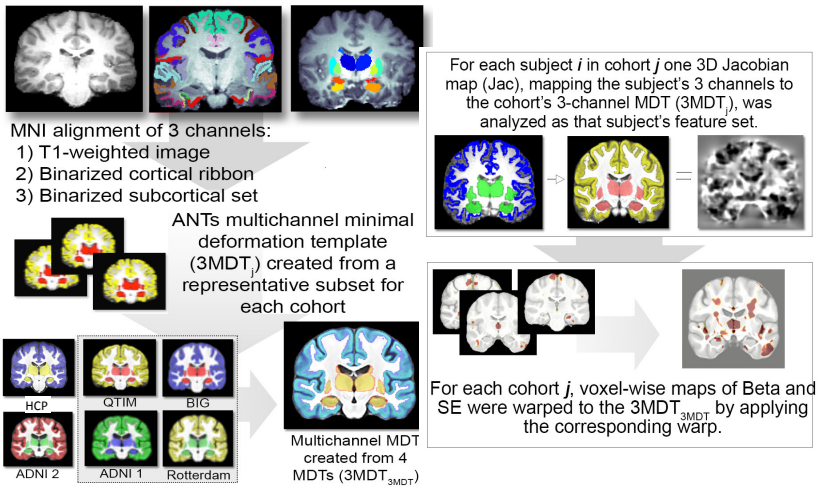

Figure 1 Flow diagram of template creation and registration. T1-weighted images run through common software, Freesurfer, and evaluated to have good quality cortical and subcortical parcellations were used along with the Freesurfer outputs to drive multi-channel registrations to a cohort-specific template. The multiple channels were used to reduce variability between cohorts to create a minimal deformation template (MDT) from 4 datasets. All associations are performed in cohort-specific space and the transformation from cohort to template space was applied to the resulting statistical maps for meta-analysis.

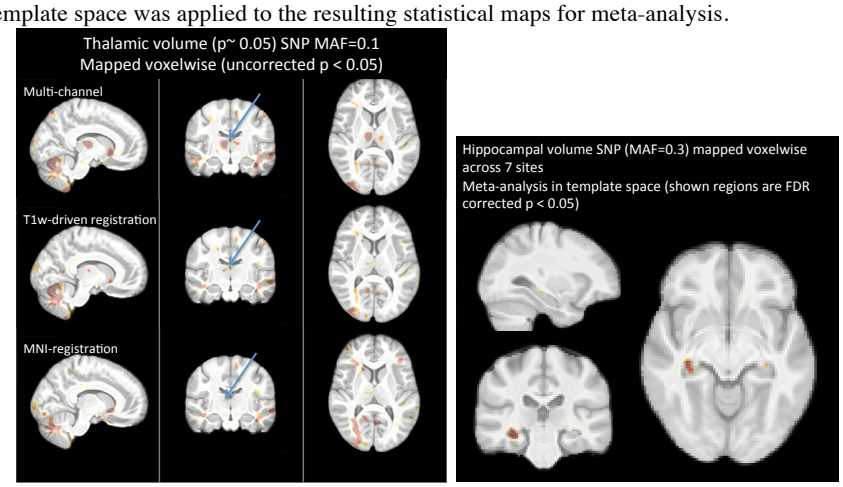

Figure 2: a) A SNP with MAF = 0.1 was simulated to be marginally (z=1.96) associated with average bilateral thalamic volume in a single cohort (after removing ICV). The effect of maintaining specificity to the thalami was compared between multiple templates. No method produced voxelwise significant maps, however, evaluating the uncorrected association results of the methods shows greater thalamic effects in the multi-channel method. b) A SNP with MAF=0.3 was generated for each of 7 cohorts, to have z=min(N/10,5) such that the significance was related to cohort size, yet was not excessive (|z| <5). Beta and SE maps for all cohorts were mapped to template space, and voxelwise meta-analysis revealed associations localized to both hippocampi.