

Histopathological image QTL discovery of thyroid autoimmune disease variants

Authors: Joseph D. Barry^{1,2*}, Maud Fagny^{1,2}, Joseph N. Paulson^{1,2}, Hugo J. W. L. Aerts³, John Platig^{1,2}, and John Quackenbush^{1,2}

Affiliations:

¹ Center for Cancer Computational Biology and Department of Biostatistics and Computational Biology, Dana-Farber Cancer Institute, 450 Brookline Ave., Boston, MA 02215

² Department of Biostatistics, Harvard T.H. Chan School of Public Health, 667 Huntington Ave., Boston, MA 02115

³ Department of Radiology, Dana-Farber Cancer Institute, 450 Brookline Ave., Boston, MA 02215

*Correspondence to: jbarry@jimmy.harvard.edu

Abstract: Genotype-to-phenotype association studies typically use macroscopic physiological measurements or molecular readouts as quantitative traits. These approaches have been successful for the identification of disease risk loci, and variants that affect gene transcription and regulation. However, there are comparatively few suitable quantitative traits available between cell and tissue length scales, a limitation that hinders our ability to identify variants affecting phenotype at many clinically informative levels. We overcome these limitations by showing that unbiased quantitative image features, automatically extracted from histopathological imaging data, can be used successfully for image Quantitative Trait Loci (iQTL) mapping and disease variant discovery. Using thyroid pathology images, clinical metadata, and genomics data from the Genotype and Tissue Expression project (GTEx), we establish and validate a quantitative imaging biomarker for thyroid autoimmune disease. A total of 80,791 candidate variants in 769 coding genes highly associated with lymphocyte invasion in our analysis were selected for iQTL profiling, and tested for genotype-phenotype associations with our quantitative imaging biomarker. Significant associations were found with variants in Histone Deacetylase 9 (HDAC9), a gene proposed to be an epigenetic switch in T-cell mediated autoimmunity, but not previously associated with thyroid autoimmune disease. We validated our results using an independent dataset of 1,213 hypothyroidism cases and 3,789 controls from the Electronic Medical Records and Genomics network (eMERGE).

One Sentence Summary: We use a histopathological image QTL analysis to identify genomic variants associated with thyroid autoimmune disease.

Introduction

Modern genomics has had tremendous success in relating genetic variants to a variety of molecular phenotypes. Variants associated with gene expression, so-called expression Quantitative Trait Loci (eQTLs), are enriched for regions of active chromatin, influence gene regulation, and are involved in processes that contribute to disease (1,2). At the macroscopic level, genome-wide association studies (GWAS) that use physiological measurements of organ health have also provided insight into the genetic component of human disease. However, given the immense gap between SNP-level variation and tissue and body function, GWAS results concerning human traits can be challenging to interpret, in part due to the high correlations observed between studies of different traits (3). This is especially true for autoimmune diseases, which have many shared genetic risk factors (4). One way to increase GWAS variant interpretability and to discover new disease variants will be to place more focus on quantitative traits that lie at intermediate cellular and sub-tissue scales.

Histopathology has for decades remained the standard for the diagnosis and grading of many complex diseases. Advances in digital slide imaging have allowed for unprecedented resolution at tissue, cellular, and sub-cellular scales. As such, histopathology could provide the genomics community with a vast resource of quantitative traits for evaluating disease phenotype across a range of mesoscopic scales. Previous histopathological GWA studies have used discrete pathology grading schemes as quantitative traits, where disease severity is assigned an integer score by pathologists. While grading schemes are highly effective for guiding clinical decisions on the level of individual patients, and can in principle identify disease variants in GWAS (5), they do not scale well to the hundreds or thousands of samples needed due to inter- and intra-observation bias (6-8). When performing millions of statistical tests in GWAS, any source of bias will increase the likelihood of finding spurious associations and negatively affect the replicability of study findings.

Unbiased, automatically extracted, continuous features have been shown to improve the predictive power of survival analyses with pathology data in comparison to discrete grades (9,10). GWA studies using quantitative image features extracted from radiological imaging data have successfully identified COPD-relevant variants using Computed Tomography (11), and variants related to Alzheimer's Disease and Mild Cognitive Impairment using Magnetic Resonance Imaging (12). Thus, genome association studies that leverage automated imaging analysis methodologies can be highly effective. With quantitative image analysis now routinely applied to digital pathology datasets (9,10,13-15), it is clear that unbiased, continuous image features can readily be extracted for histopathological GWAS variant discovery.

Nevertheless, little attention has been paid to the use of cellular imaging to obtain quantitative traits for understanding the connection between genotype and cellular phenotypes. Part of the reason for this may be that there are few available datasets for which both genomic profiling and standardized histopathology data are available for any appreciable number of samples. The Genotype-Tissue Expression (GTEx) resource (16,17) has collected germline genotype data from over 450 autopsy research subjects. The project has also obtained biopsy samples for standardized histological imaging analysis and RNA-Seq gene expression analysis from approximately 50 different body sites with as much data as possible obtained from each individual.

One of the tissues for which there is substantial histological imaging data is the thyroid. The thyroid plays a central role in the endocrine system, producing thyroid hormones that influence metabolic rate and protein synthesis. Hashimoto's Thyroiditis (HT) is an autoimmune disease in which the thyroid gland is slowly destroyed, often leading to hypothyroidism. It affects about 5% of people at some point in their lives, typically begins between the ages of 30 and 50, and is much more prevalent in women than in men (18). Among the GTEx population there are 341 individuals for whom paired thyroid imaging and gene expression data was available; of these 31 had morphological evidence of HT according to GTEx pathology notes. We saw this dataset as an ideal place to test the hypothesis that cellular imaging features could be associated with genetic variants in an image quantitative trait locus (iQTL) analysis.

Quantitative image analysis

Image data was downloaded from the GTEx histological image archive (Materials and Methods). Images were first convolved with a Gaussian filter to smooth out pixel-level variation on a length scale smaller than observed regions of lymphocytic invasion in HT samples. Individual tissue pieces were segmented using adaptive thresholding (Figure 1A), and 117 Haralick image features were extracted for each tissue piece using the Bioconductor image processing package *EImage* (19) (Materials and Methods). Haralick features are widely used for the quantification of image texture, and take into account correlations between neighboring pixels in an image (20). Given that lymphocyte invasion has a profound effect on histological image texture, Haralick features were ideal candidates for capturing HT cellular phenotype. After removing overly small tissue pieces, averaged image feature values were calculated across pieces to give a single value for each sample and feature. In preparation for downstream model fitting, features were \log_2 transformed, and centered and rescaled using a z-score to ensure feature comparability (Fig. 1B). A fully worked through image analysis example with the applied segmentation parameters is supplied as part of a supplemental R data package.

To reduce the number of candidate features and to identify the primary directions of variability in our image feature matrix, we performed Principal Component Analysis (PCA), and found that the first two imaging Principal Components (PCs) were sufficient to explain 73% of the variance. By inspection of the images we found that low values of PC1 were associated with interior holes in the thyroid pieces (Figure S1A), an effect we adjudged to be primarily technical in origin. High values of image PC2 were visually strongly associated with the presence of invading lymphocytes (Figure S1B).

To validate this observation we identified all samples for which HT was indicated in the comments of the original GTEx pathologists (Figure 1C blue points). A ROC curve showed that image PC2 was in fact a highly performant biomarker for HT, with an AUC of 0.84 (Figure 1D). Using logistic regression, we found that HT status was significantly associated with image PC2 (OR=1.3, $P=1.4 \times 10^{-9}$) after correcting for sex and age (Materials and Methods). We also observed a higher association with HT status for women after correcting for age (OR=4.2, $P=4 \times 10^{-4}$), a known clinical characteristic of the disease (18). That image PC2 is a performant HT biomarker likely does not have direct clinical relevance since a pathologist can easily diagnose the disease by inspection of pathology images; we considered its utility to be for downstream integration with genomic data.

Integrated RNA-Seq and image analysis: imaging biomarker validation

Next we analyzed GTEx gene expression data associated with the 341 thyroid images. After the removal of lowly expressed genes, counts were \log_2 transformed and underwent quantile normalization to ensure comparability of samples (Materials and Methods). To test for associations between gene expression and image PCs we used linear models to fit each principal component against normalized gene readouts while correcting for age, sex, sample collection site, tissue autolysis score, and RNA extraction type. To remove bias caused by unknown confounders (Figure S2A) we fitted 20 PEER (Probabilistic Estimation of Expression Residuals) factors (21) to the gene expression data while including the known confounders in the PEER fit; a similar approach was used by the GTEx consortium (16). The first PEER factor was not included as a covariate in the linear model fits due to the strength of its correlation with our image PC2 biomarker (Pearson $r=0.53$, $P<2.2e-16$), which indicated that it likely contained valuable biological information. The results from fits correcting for both the known confounders and the 19 PEER factors revealed that image PC2 was systematically and highly associated with gene expression (Figure 2A), suggesting that lymphocyte infiltration is at least a major, if not the primary source of expression variability in GTEx thyroid tissue. In contrast, image PC1 had comparatively few significant associations with thyroid gene expression (Figure 2A).

To identify which biological pathways may drive the significant image PC2 associations with gene expression we performed Gene-Set Enrichment Analysis, choosing as a test set 2,913 genes with $-\log_{10}(P_G) > 10$ in the gene expression analysis and all other 21,080 genes as a background set (Materials and Methods). Significant enrichment was observed for a number of Gene Ontology terms related to immune function, and the activation and signaling activity of invading T and B lymphocytes (Table 1). To confirm the association between image PC2 and the presence of invading immune cells on the level of individual samples, we ran our gene expression data through CIBERSORT, an algorithm designed to deconvolve complex cell mixtures (22). We used the LM22 leukocyte gene signature matrix supplied by the CIBERSORT team to perform the deconvolution as it could in principle detect and distinguish between a number of T and B cell types. As expected, CIBERSORT failed to detect immune cells in the majority of healthy thyroid samples. However, in samples where immune cells were detected (CIBERSORT $P_C < 0.5$), the method's statistical power correlated significantly with image PC2 (Figure 2B; Spearman $\rho=0.67$, $P=3 \times 10^{-7}$).

We also detected infiltrating immune cells in several samples for which image PC2 was high but HT was not indicated by GTEx pathologists (Figure 2B grey points), suggesting that our imaging biomarker may be able to detect a broader spectrum of thyroid autoimmune diseases. By inspection, we confirmed that many of these images did show infiltrative phenotypes but did not exhibit the classical HT visual phenotype (for example GTEX-OXRO, GTEX-144FL), suggesting that they were not false negatives for HT. Figure 2C shows the cell type signature for samples with $P_C < 0.1$ and cell types with an average CIBERSORT frequency of more than 5% across samples. This analysis confirmed the presence of T-cell CD4 memory resting, T-cell CD8, naïve and memory B cells, as well as other immune cell types, on the level of individual samples.

To demonstrate the advantage of using PCA for the identification of a candidate imaging biomarker, as opposed to directly using the individual Haralick image features, we correlated each of the 117 image features with CIBERSORT's statistical power to detect immune cells, in the same manner as above. Image PC2 had a higher Spearman correlation coefficient than all

image features, with the closest having a coefficient of $\rho=0.60$ (Figure S2B). Thus, image PC2 combined information from multiple image features to attain a higher overall performance.

iQTL identification of Hashimoto's Thyroiditis disease variants

Based on our analysis of image features and gene expression, we decided to use image PC2 as a quantitative trait capturing the cellular features of HT in an image QTL analysis. We limited our analysis to the 292 samples that had both thyroid imaging and genotype information. The genotype data had already undergone rigorous quality control and filtering by the GTEx consortium, including but not limited to Hardy-Weinberg equilibrium and imputation quality criteria. For our analysis we also removed SNPs with a missingness frequency of more than 10%, and due to sample size considerations, minor allele frequencies (MAFs) of less than 10%.

Since we did not have sufficient sample size to achieve high power in an iQTL analysis that would include all coding and non-coding regions of the genome, we decided to reduce the search space by asking a more targeted question of the data. We hypothesized that the genotype of expressed coding genes in the invading lymphocytes could in part determine thyroid autoimmune disease risk and severity. Therefore, we further restricted our analysis to the 80,791 candidate SNPs residing in 769 coding genes that were significantly ($-\log_{10}(P_G) > 10$) and positively (regression coefficient $\beta > 2.5$) associated with image PC2 in the gene expression analysis (Figure 3A blue points).

Image QTL linear model fits were implemented using the R package *MatrixEQTL* (23) by treating image PC2 as a pseudo *trans* gene. Fits were corrected for age, sex, race, ancestry, processing center, and tissue autolysis score (Materials and Methods). A QQ plot comparing the distribution of observed p-values (denoted P_O) to the expected uniform distribution showed that the data was well behaved under the null hypothesis (Figure 3B), with a departure from the dashed line for a group of SNPs with low p-values. Independent Hypothesis Weighting (IHW) (24) using MAF as the independent covariate was used to correct for multiple testing. In a data-driven way, this method assigned higher weights to variants with a high MAF (Figure S3) to maximize the number of null hypothesis rejections while controlling for Type I error. A total of 32 SNPs across 5 haplotype blocks were identified as being significant ($P_{IHW} < 0.05$; Table S1).

Strikingly, all 32 SNPs reside in Histone Deacetylase 9 (HDAC9), an enzyme linked to epigenetic control of gene transcription, and previously proposed to be an epigenetic switch for T-cell mediated autoimmunity (25). Notably, histone deacetylase inhibitors have been effective in the treatment of hypothyroidism in mice (26), and autoimmune diseases in general (27). A plot of SNP significance along HDAC9 revealed a sharp peak at 18,313 kb on chromosome 7 (Figure 3C). To our knowledge this is the first reported association between HDAC9 variants and thyroid autoimmune disease. A search of GWAS Central (www.gwascentral.org (28)) revealed that modest associations between Ulcerative Colitis (UC) and two of our significant image QTL variants ($-\log_{10}(P)=2.3$ for *rs215122*, $-\log_{10}(P)=2.7$ for *rs2529749*) have previously been reported (29). While further work will be required to investigate the biological effects of HDAC9 genotype in both UC and HT, a common genetic basis would be consistent with evidence that these conditions present concurrently at higher rates than would be expected in the general population (30).

Validation of iQTL results using an independent dataset

To validate the observed association between HDAC9 genotype and thyroid autoimmune disease in an independent dataset, we obtained genotype and phenotype data, and clinical metadata from the Electronic Medical Records and Genomics (eMERGE) network (31) (Materials and Methods). We selected 1,261 chronic autoimmune hypothyroidism (presumptive Hashimoto's hypothyroidism) cases and 4,457 non-hypothyroidism controls (Materials and Methods). As over 96% of the hypothyroidism cases were from the Caucasian cohort, and 250/292 of the GTEx samples used for the image QTL analysis were of European American descent, to simplify the analysis we removed all non-Caucasian cohorts, leaving 1,213 cases and 3,789 controls. For consistency with the image QTL analysis, SNPs with a minor allele frequency of less than 10% and a missingness frequency of more than 10% were removed. One SNP in HDAC9 was removed due to deviation from Hardy-Weinberg equilibrium (Materials and Methods).

We used logistic regression to test for associations between hypothyroidism status and HDAC9 genotype while correcting for patient age (decade of birth), sex, and collection site (Materials and Methods). While the maximum strength of association was modest ($-\log_{10}(P)=2.3$), the location of the highest eMERGE peak coincided closely with the highest GTEx image QTL peak (Figure 3C,D vertical line), and we noted a high degree of similarity between the association profiles. To quantitatively assess the concordance of the HDAC9 association signature between the GTEx and eMERGE data, we first selected the 1,071 HDAC9 SNPs common to both datasets, and grouped neighboring SNPs into bins of size 15 SNPs based on their genomic coordinates. The lowest p-value in each bin was then selected for both datasets. A correlation test revealed that the association signature between the eMERGE and GTEx data was highly concordant across HDAC9 (Figure S4A; Spearman $\rho=0.68$, $P<2.2\times 10^{-16}$).

To verify that it was hypothyroidism status driving the concordance, we performed 1,000 permutations of the hypothyroidism case/control labels in the eMERGE dataset and repeated the binning and correlation analysis. A total of 992 of the 1,000 eMERGE permutations resulted in a lower Spearman correlation coefficient than the non-permuted data (Figure S4B), demonstrating that it was unlikely to obtain the observed concordance between the eMERGE and GTEx HDAC9 profiles by chance ($P=0.009$, Materials and Methods).

This validation also showed that image QTL approaches that use continuous phenotypic measurements can attain higher power for the detection of disease-relevant genomic regions as compared to traditional cases versus controls analyses. One explanation for this might be that quantitative histological image analysis captures subclinical, as well as clinical disease phenotype as a smooth function, while binary categorizations of disease status do not take into account informative subclinical disease cases, or the morphological severity of the disease phenotype.

Outlook

Our intuition, along with over a decade of progress in GWAS, tells us that genotype and phenotype are associated in ways that can be quantitatively assessed. The growth of integrated data resources such as GTEx, and the advancement of digital slide imaging technologies, provides the opportunity to explore that association at the cellular level. This study demonstrates the potential of histopathological image-based QTL profiling for *de novo* discovery of disease

variants, especially when unbiased, continuous, automatically extracted image features are used as quantitative traits. While we applied histopathological iQTLs to a thyroid dataset, the method is completely generalizable to other tissues and disease contexts. Therefore, we foresee that with larger sample sizes and study designs tailored for specific diseases, genome-wide scans that use similar approaches can be extended to whole-exome and non-coding regions. Ultimately, digital pathology image analysis approaches could revolutionize genome-wide association studies by providing a wealth of unbiased quantitative traits that have high levels of biological interpretability. Our results also suggest that HDAC9 may be a viable drug target for patients with thyroid autoimmune disease.

References

1. Pickrell JK, Marioni JC, Pai AA, Degner JF, Engelhardt BE, Nkadori E, et al. Understanding mechanisms underlying human gene expression variation with RNA sequencing. *Nature* 2010;464:768-72.
2. Westra H-J, Peters MJ, Esko T, Yaghootkar H, Schurmann C, Kettunen J, et al. Systematic identification of trans eQTLs as putative drivers of known disease associations. *Nature Genetics* 2013;45:1238-43.
3. Bulik-Sullivan B, Finucane HK, Anttila V, Gusev A, Day FR, Loh P-R, et al. An atlas of genetic correlations across human diseases and traits. *Nature Genetics* 2015;47:1236-41.
4. Gutierrez-Arcelus M, Rich SS, Raychaudhuri S. Autoimmune diseases — connecting risk alleles with molecular traits of the immune system. *Nature Reviews Genetics* 2016;17:160-74.
5. Chalasani N, Guo X, Loomba R, Goodarzi MO, Haritunians T, Kwon S, et al. Genome-Wide Association Study Identifies Variants Associated with Histologic Features of Nonalcoholic Fatty Liver Disease. *Gastroenterology* 2010;139:1567-76.e6.
6. Layfield LJ, Esebua M, Frazier SR, Hammer RD, Bivin WW, Nguyen V, et al. Accuracy and Reproducibility of Nuclear/Cytoplasmic Ratio Assessments in Urinary Cytology Specimens. *Diagnostic Cytopathology* 2017;45:107-12.
7. Qureshi A, Lakhtakia R, AL Bahri M, Al Haddabi I, Saparamadu A, Shalaby A, et al. Gleason's Grading of Prostatic Adenocarcinoma: Inter-Observer Variation Among Seven Pathologists at a Tertiary Care Center in Oman. *Asian Pacific journal of cancer prevention: APJCP* 2016;17:4867-68.
8. Colquhoun P, Nogueras JJ, Dipasquale B, Petras R, Wexner SD, Woodhouse S. Interobserver and intraobserver bias exists in the interpretation of anal dysplasia. *Diseases of the Colon and Rectum* 2003;46:1332-36; discussion 36-38.
9. Yu K-H, Zhang C, Berry GJ, Altman RB, Ré C, Rubin DL, et al. Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features. *Nature Communications* 2016;7.
10. Khan FM, Scott R, Donovan M, Fernandez G. Predicting and replacing the pathological Gleason grade with automated gland ring morphometric features from immunofluorescent prostate cancer images. *Journal of Medical Imaging (Bellingham, Wash)* 2017;4:021103.
11. Cho MH, Castaldi PJ, Hersh CP, Hobbs BD, Barr RG, Tal-Singer R, et al. A Genome-Wide Association Study of Emphysema and Airway Quantitative Imaging Phenotypes. *American Journal of Respiratory and Critical Care Medicine* 2015;192:559-69.
12. Shen L, Kim S, Risacher SL, Nho K, Swaminathan S, West JD, et al. Whole genome association study of brain-wide imaging phenotypes for identifying quantitative trait loci in MCI and AD: A study of the ADNI cohort. *NeuroImage* 2010;53:1051-63.
13. Popovici V, Budinská E, Čápková L, Schwarz D, Dušek L, Feit J, et al. Joint analysis of histopathology image features and gene expression in breast cancer. *BMC Bioinformatics* 2016;17:209.

14. Yuan Y, Failmezger H, Rueda OM, Ali HR, Gräf S, Chin S-F, et al. Quantitative Image Analysis of Cellular Heterogeneity in Breast Tumors Complements Genomic Profiling. *Science Translational Medicine* 2012;4:157ra43.
15. Beck AH, Sangoi AR, Leung S, Marinelli RJ, Nielsen TO, Vijver MJvd, et al. Systematic Analysis of Breast Cancer Morphology Uncovers Stromal Features Associated with Survival. *Science Translational Medicine* 2011;3:108ra13-08ra13.
16. Consortium G. Human genomics. The Genotype-Tissue Expression (GTEx) pilot analysis: multitissue gene regulation in humans. *Science (New York, NY)* 2015;348:648-60.
17. Lonsdale J, Thomas J, Salvatore M, Phillips R, Lo E, Shad S, et al. The Genotype-Tissue Expression (GTEx) project. *Nature Genetics* 2013;45:580-85.
18. Pyzik A, Grywalska E, Matyjaszek-Matuszek B, Roliński J. Immune disorders in Hashimoto's thyroiditis: what do we know so far? *Journal of Immunology Research* 2015;2015:979167.
19. Pau G, Fuchs F, Sklyar O, Boutros M, Huber W. EBImage—an R package for image processing with applications to cellular phenotypes. *Bioinformatics* 2010;26:979-81.
20. Haralick RM. Statistical and structural approaches to texture. *Proceedings of the IEEE* 1979;67:786-804.
21. Stegle O, Parts L, Piipari M, Winn J, Durbin R. Using probabilistic estimation of expression residuals (PEER) to obtain increased power and interpretability of gene expression analyses. *Nature Protocols* 2012;7:500-07.
22. Newman AM, Liu CL, Green MR, Gentles AJ, Feng W, Xu Y, et al. Robust enumeration of cell subsets from tissue expression profiles. *Nature Methods* 2015;12:453-57.
23. Shabalin AA. Matrix eQTL: ultra fast eQTL analysis via large matrix operations. *Bioinformatics (Oxford, England)* 2012;28:1353-58.
24. Ignatiadis N, Klaus B, Zaugg JB, Huber W. Data-driven hypothesis weighting increases detection power in genome-scale multiple testing. *Nature Methods* 2016;13:577-80.
25. Yan K, Cao Q, Reilly CM, Young NL, Garcia BA, Mishra N. Histone Deacetylase 9 Deficiency Protects against Effector T Cell-mediated Systemic Autoimmunity. *The Journal of Biological Chemistry* 2011;286:28833-43.
26. Kim DW, Park JW, Willingham MC, Cheng S-y. A histone deacetylase inhibitor improves hypothyroidism caused by a TRα1 mutant. *Human Molecular Genetics* 2014;23:2651-64.
27. Tang J, Yan H, Zhuang S. Histone deacetylases as targets for treatment of multiple diseases. *Clinical Science* 2013;124:651-62.
28. Beck T, Hastings RK, Gollapudi S, Free RC, Brookes AJ. GWAS Central: a comprehensive resource for the comparison and interrogation of genome-wide association studies. *European journal of human genetics: EJHG* 2014;22:949-52.
29. Anderson CA, Boucher G, Lees CW, Franke A, D'Amato M, Taylor KD, et al. Meta-analysis identifies 29 additional ulcerative colitis risk loci, increasing the number of confirmed associations to 47. *Nature Genetics* 2011;43:246-52.
30. Shizuma T. Concomitant Thyroid Disorders and Inflammatory Bowel Disease: A Literature Review. *BioMed Research International* 2016;2016:e5187061.
31. Gottesman O, Kuivaniemi H, Tromp G, Faucett WA, Li R, Manolio TA, et al. The Electronic Medical Records and Genomics (eMERGE) Network: past, present, and future. *Genetics in Medicine* 2013;15:761-71.

Notes: Code and analysis reproducing paper figures, as well as reported image and statistical analyses performed on non-protected data, are supplied as part of an R package that will be made available upon publication.

Author Contributions: JDB, HJWLA and JQ conceived of the project. JDB performed all image and data analysis, with input from MF (CIBERSORT analysis), JNP (statistical analysis), and JP (QTL pipeline). JDB wrote the paper with input from all authors.

Acknowledgements: This work was supported by grants from the US National Institutes of Health, including 5U01CA190234 and 5U24CA194354 (JDB, HJWLA, JQ), 5R01AI099204 (JP, JQ); 5R01HL111759 (JQ, JNP), and 5P50CA127003 (JQ). This work was conducted under dbGaP approved protocols #9112 and #13896.

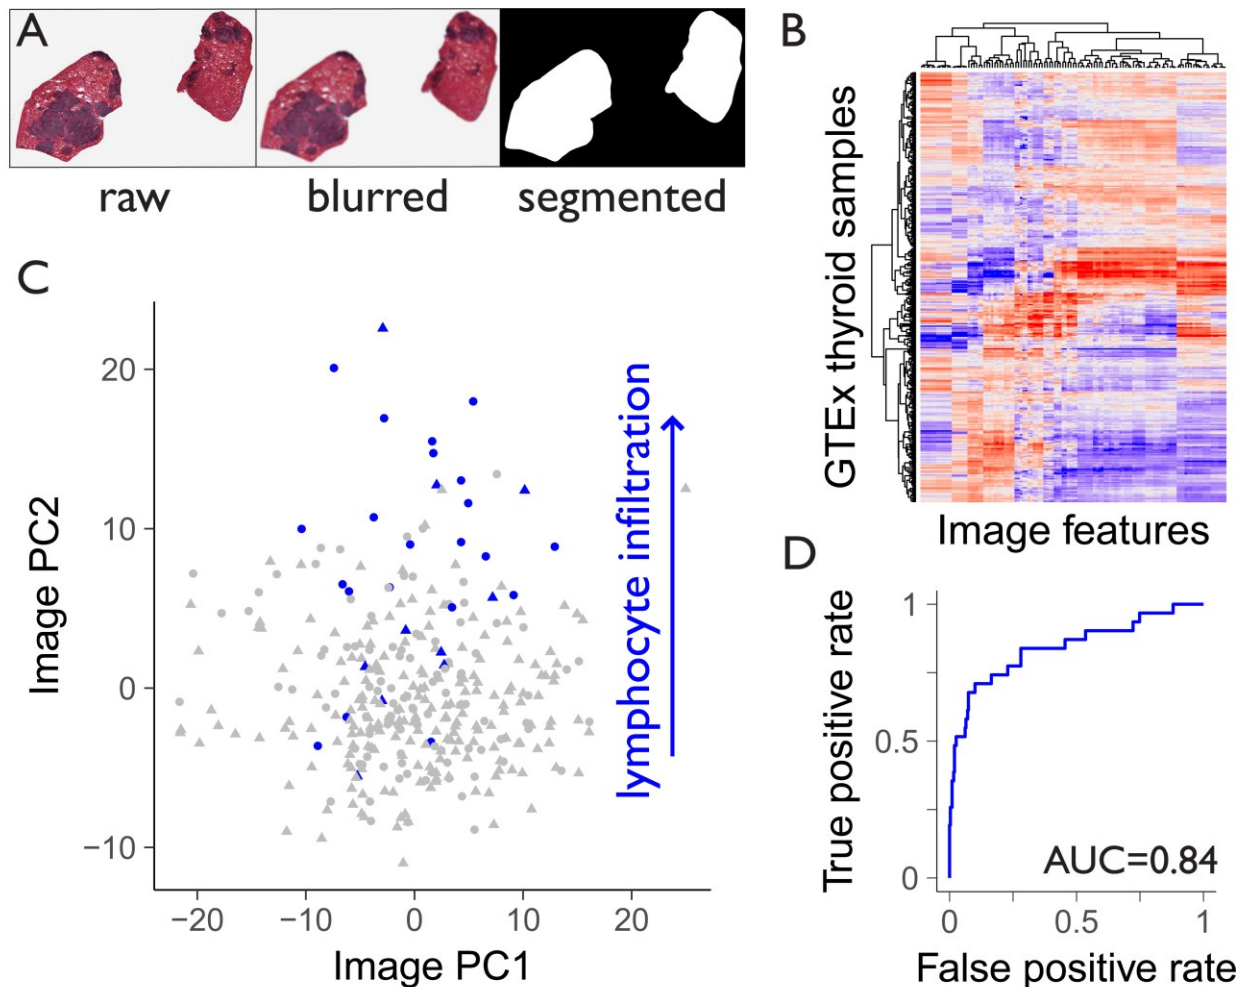


Figure 1. Thyroid image processing and establishment of a quantitative imaging biomarker for Hashimoto's Thyroiditis (HT). (A) Digital pathology slide for thyroid sample GTEX-11NV4. Raw: pre image processing. Blurred: post Gaussian convolution. Segmented: tissue piece masks after adaptive thresholding. (B) Heatmap of 117 log₂ transformed and standardized Haralick image features against 341 thyroid samples. (C) PC1 versus PC2 from a PCA of the image feature matrix. Blue points indicate patients with HT, as identified from pathology notes. Circles indicate females, and triangles males. (D) ROC curve showing biomarker performance of PC2 for predicting HT.

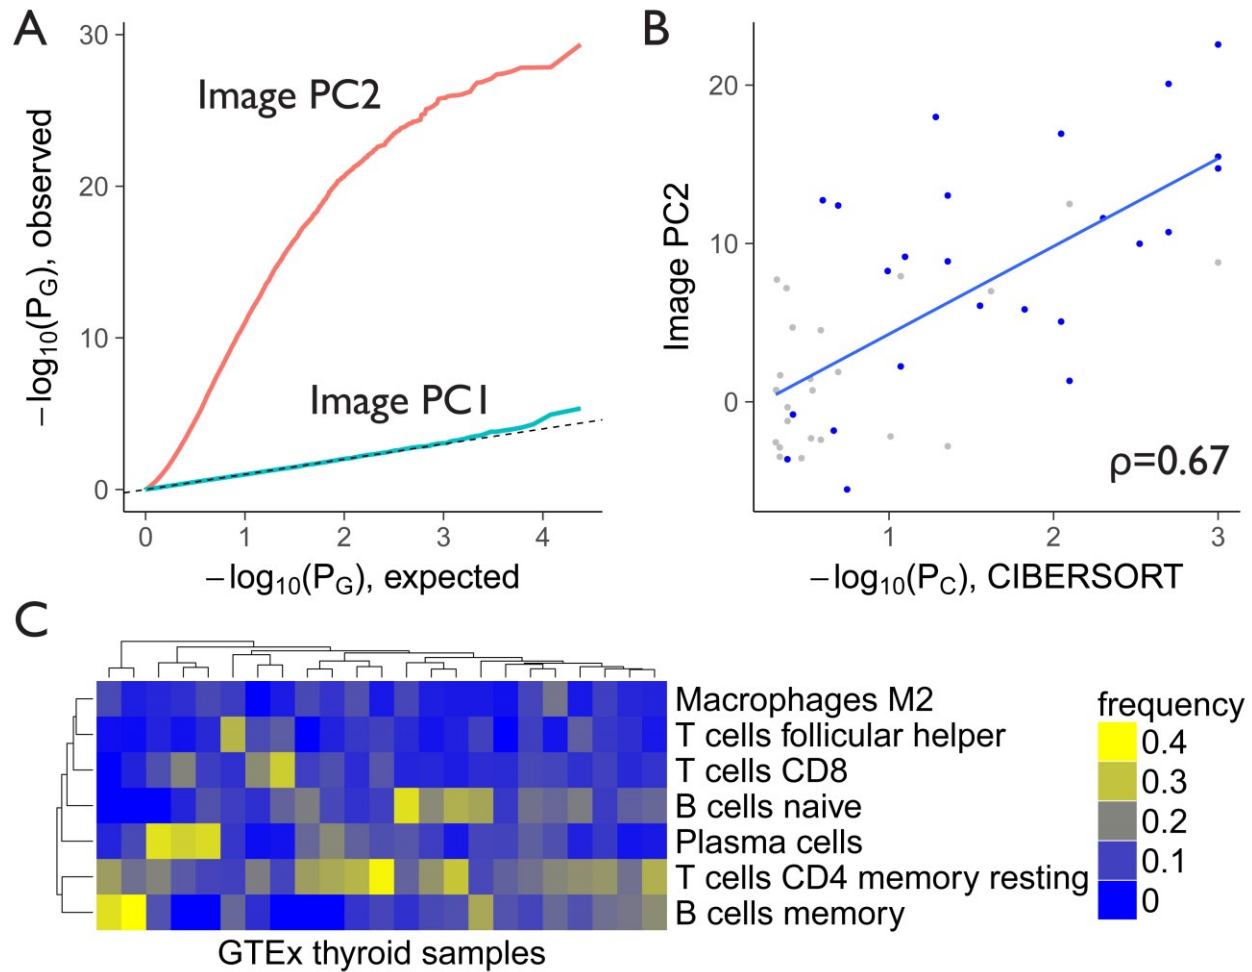


Figure 2. Integration of imaging data with gene expression analyses verifies that image PC2 is highly associated with autoimmune response. **(A)** A QQ plot of p-values (P_G) from the regression analysis of image PC1 and PC2 against thyroid gene expression for 23,993 genes. **(B)** Correlation of image PC2 with $-\log_{10}(P_C)$ from the CIBERSORT analysis for samples with $P_C < 0.5$. Blue points are samples for which Hashimoto's Thyroiditis was indicated in GTEx pathology notes. **(C)** Frequencies of immune cell types reported from CIBERSORT for samples with $P_C < 0.1$. Cell types with an average frequency of 5% or more are shown.

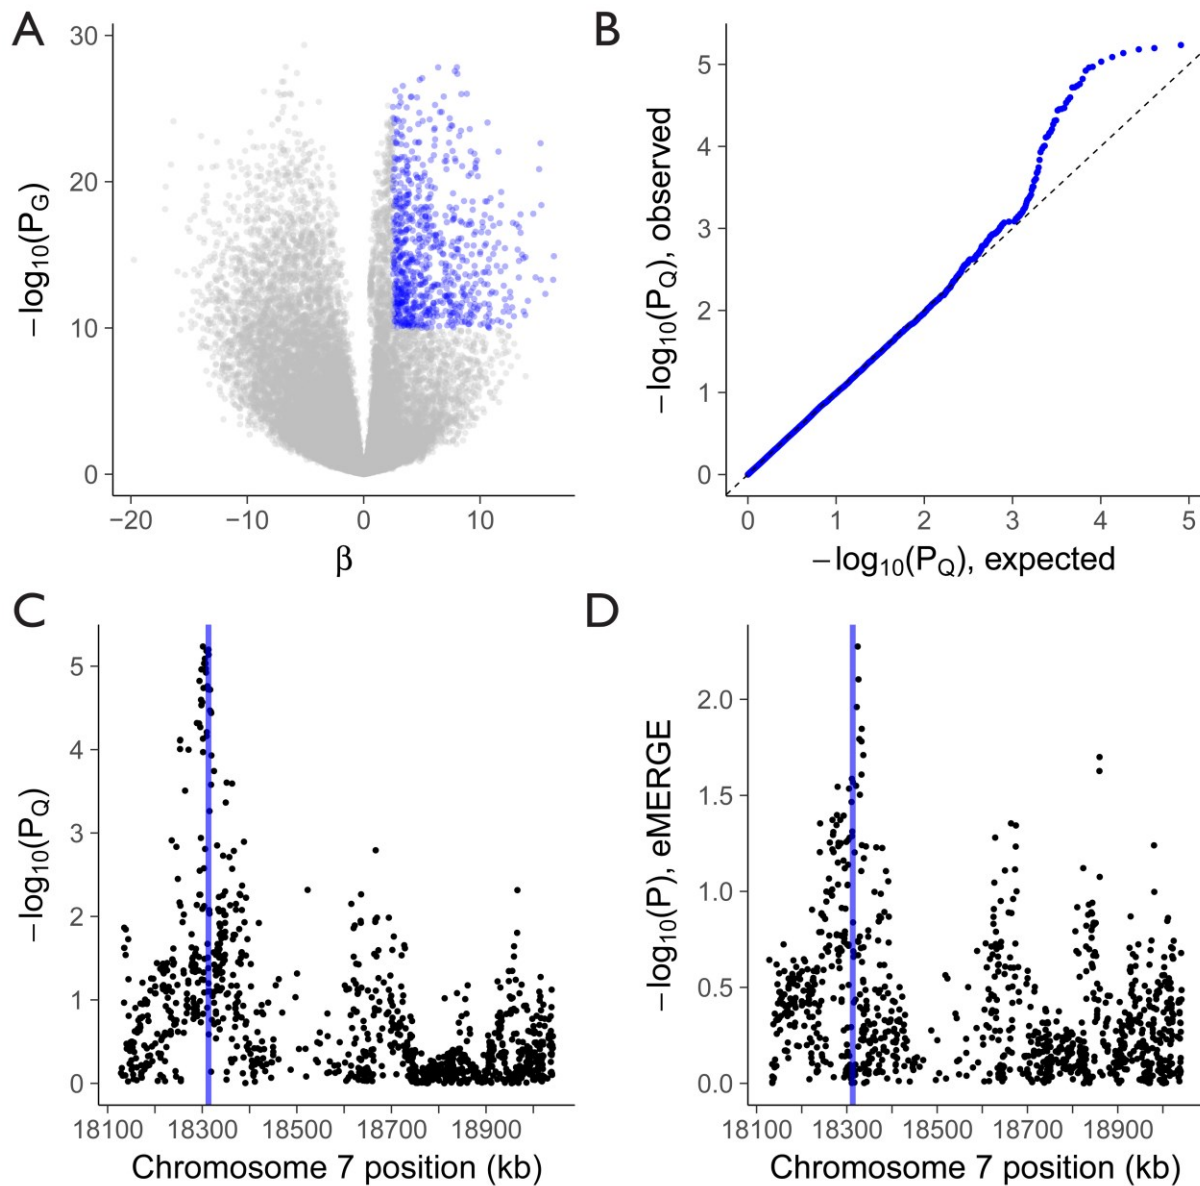


Figure 3. An image QTL analysis finds associations between HDAC9 variants and our image PC2 biomarker for thyroid autoimmune disease. **(A)** Selection of 769 candidate genes (blue points) based on their positive (regression coefficient $\beta > 2.5$) and significant ($-\log_{10}(P_G) > 10$) association with image PC2 in the gene expression analysis. **(B)** A QQ plot showing expected vs observed p-values (P_Q) from image QTL fits of 80,791 candidate SNPs residing in the selected genes highlighted blue in panel A. **(C)** A plot of $-\log_{10}(P_Q)$ for all tested SNPs in HDAC9. A sharp peak was identified at 18,313 kb on Chromosome 7. Vertical line indicates the location of the most significant SNP after multiple testing correction using the IHW method described in the main text. Plot range is mapped to the start and end positions of HDAC9. **(D)** Significance of associations between HDAC9 variants and hypothyroidism status in the eMERGE dataset used for independent validation. Vertical line indicates the location of the top image QTL SNP.

Table 1. Significant gene ontology terms for genes highly associated with image PC2 ($-\log_{10}(P_G) > 10$) in the gene expression regression analysis. All terms with an odds ratio of greater than 5 and an adjusted p-value of less than 0.01 are shown. BP: Biological Process, CC: Cellular Component.

GO id	Type	Description	Odds	P	P _{adj}	N
GO:0050776	BP	regulation of immune response	6.10	1.3E-14	5.8E-11	43
GO:0031295	BP	T cell costimulation	5.83	6.0E-11	2.1E-07	32
GO:0042110	BP	T cell activation	5.57	2.6E-08	3.6E-05	24
GO:0042113	BP	B cell activation	8.67	7.0E-08	8.6E-05	17
GO:0045060	BP	negative thymic T cell selection	25.51	6.7E-07	5.4E-04	10
GO:0050690	BP	regulation of defense response to virus by virus	6.96	2.1E-06	1.5E-03	15
GO:0042101	CC	T cell receptor complex	17.00	2.5E-06	1.6E-03	10
GO:0050853	BP	B cell receptor signaling pathway	5.89	7.2E-06	3.8E-03	15

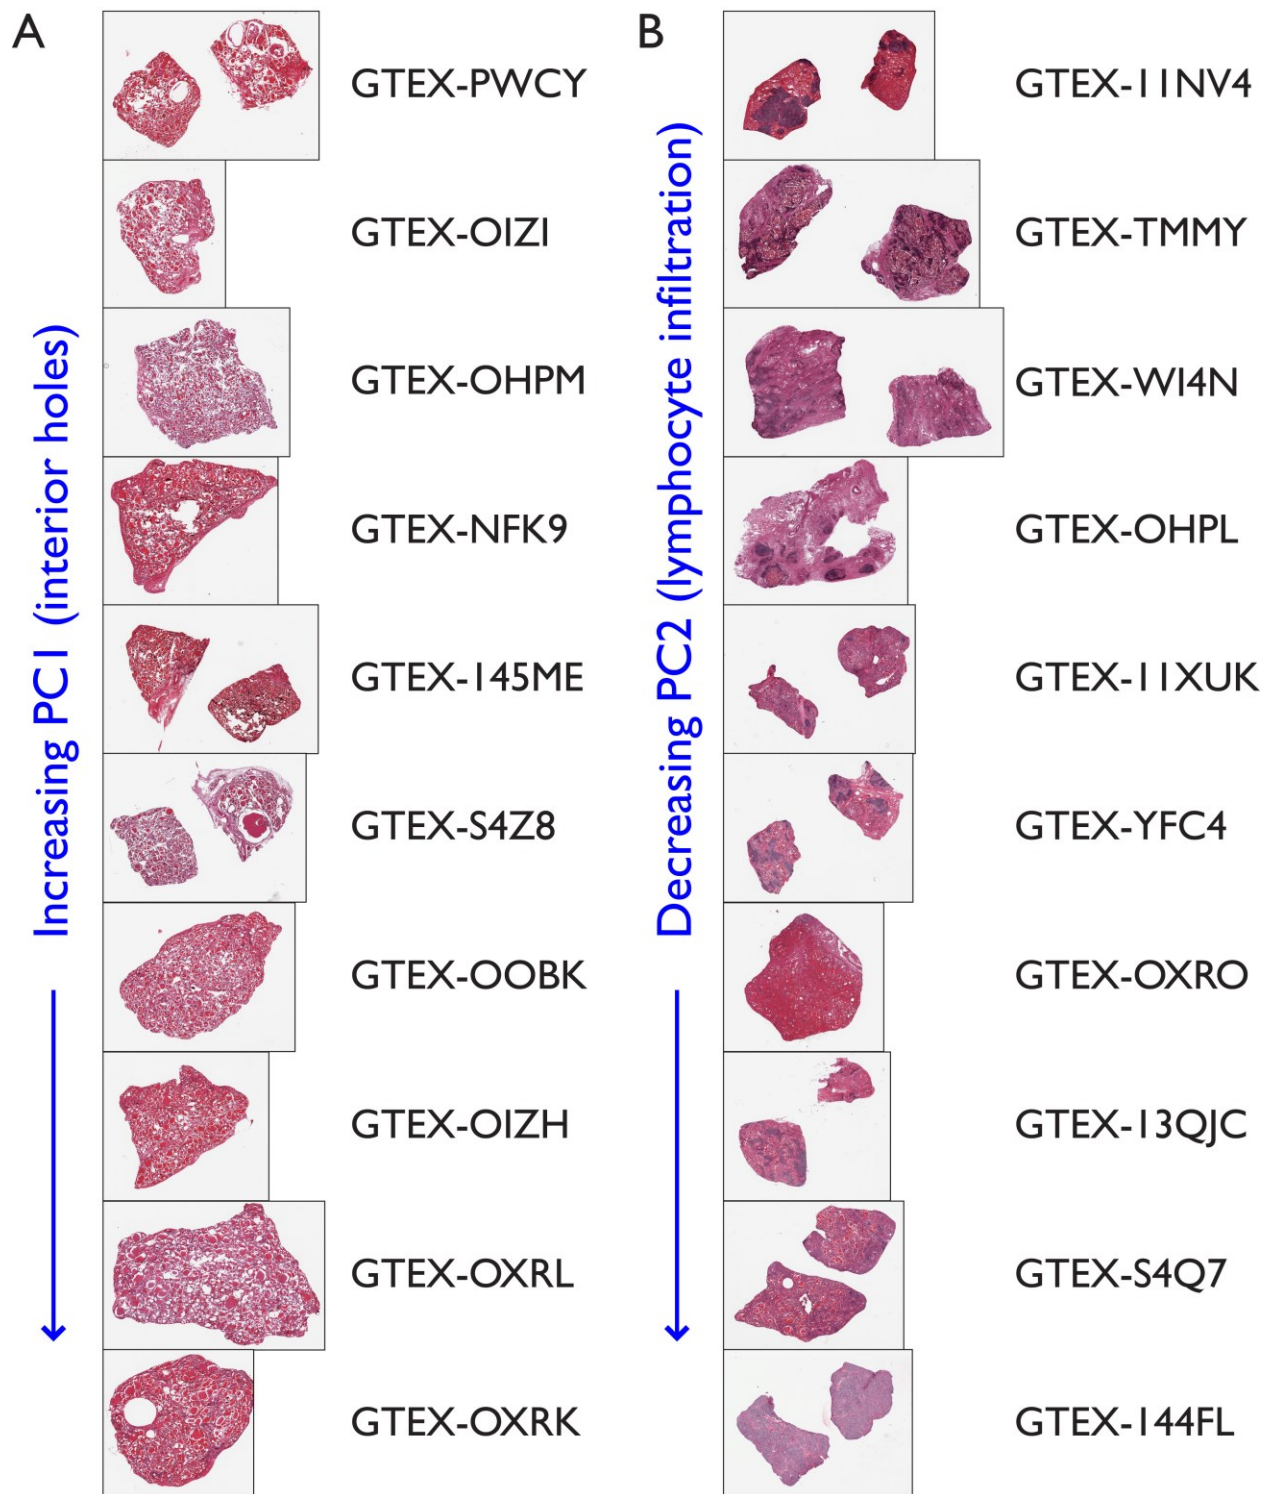


Figure S1. Examples of phenotypes associated with image PC1 and PC2. (A) Images with the 10 lowest values of PC1. Interior holes were observed, many of which were likely associated with tissue damage incurred during sample preparation. (B) Images with the 10 highest values of PC2. Lymphocyte invasion phenotypes were apparent.

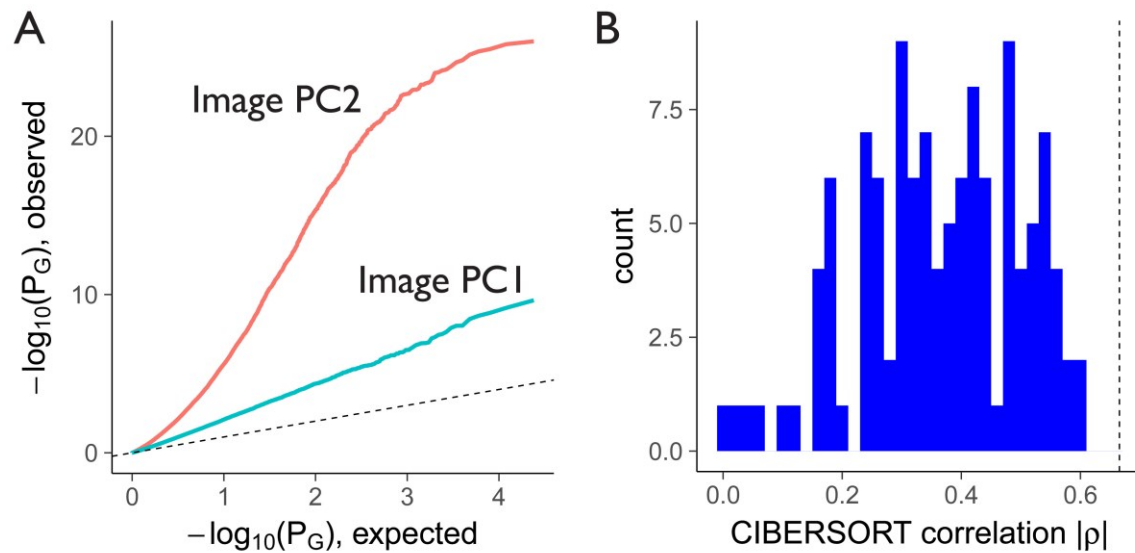


Figure S2. (A) QQ plot showing deviation from the dashed line (of slope 1) when fitting image PCs against gene expression while only correcting for the known confounders listed in the main text. Additionally including PEER factors as covariates removed the observed bias (Figure 2A). (B) Histogram of the absolute value of Spearman correlation coefficients between all 117 Haralick image features and CIBERSORT's statistical power to detect immune cells. Dashed line indicates the Spearman correlation coefficient for image PC2 ($\rho=0.67$).

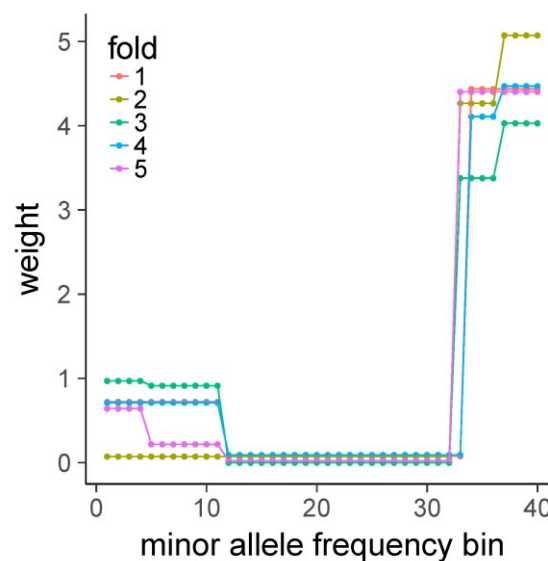


Figure S3. Independent Hypothesis Weighting weights as a function of minor allele frequency bin. Each bin corresponds to a change of 0.01 in minor allele frequency between 0.1 and 0.5. Similar profiles for each fold demonstrate the robustness of the IHW fitting procedure.

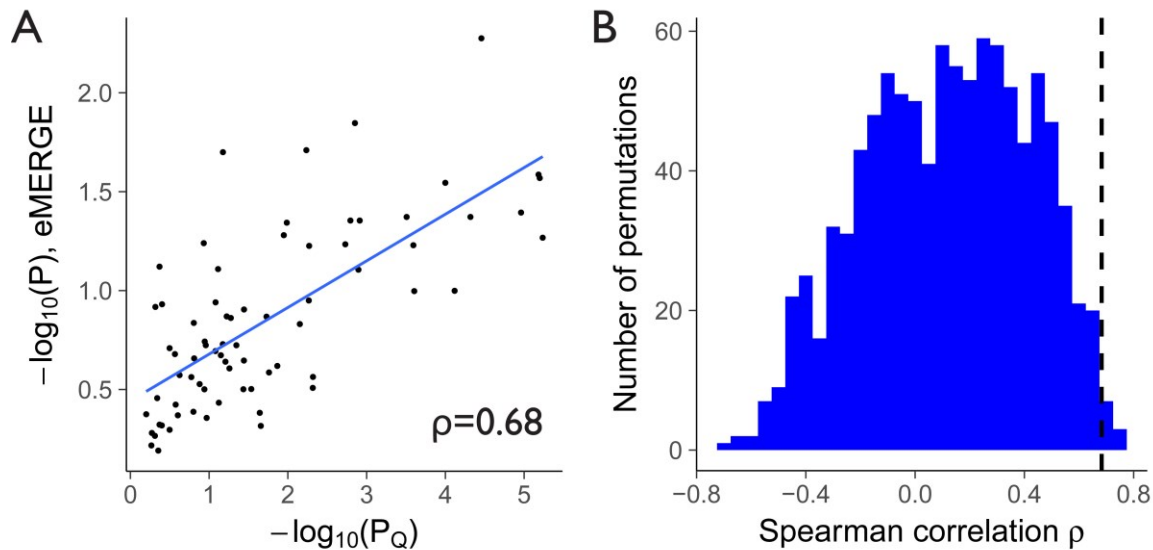


Figure S4. (A) Concordance of the HDAC9 association signature between the GTEx (x-axis) and eMERGE (y-axis) data. Each point represents the minimum p-value in each HDAC9 bin of size 15 SNPs for each dataset. (B) Histogram of the GTEx and eMERGE correlation coefficients for 1,000 permutations of the eMERGE hypothyroidism case/control labels. Dashed line indicates the Spearman correlation coefficient for the non-permuted data displayed in panel A.

Table S1. Image QTLs with significant associations ($P_{IHW} < 0.05$) with our quantitative imaging biomarker for thyroid autoimmune disease.

rsid	GRCh37/hg19 Coordinates	Gene	P_{IHW}	P	β	MAF	Haplotype
rs12154277(G>T)	Chr7:18313167	HDAC9	0.020	6.3E-06	1.912	0.434	hap1
rs12672535(G>A)	Chr7:18308913	HDAC9	0.020	6.6E-06	1.904	0.438	hap2
rs6972958(G>T)	Chr7:18313358	HDAC9	0.020	7.3E-06	1.896	0.436	hap1
rs6973137(G>A)	Chr7:18313517	HDAC9	0.020	7.3E-06	1.896	0.436	hap1
rs4416737(C>T)	Chr7:18305489	HDAC9	0.020	8.1E-06	1.913	0.431	hap2
rs73066383(C>A)	Chr7:18303874	HDAC9	0.020	9.3E-06	1.895	0.429	hap2
rs12154412(G>A)	Chr7:18304326	HDAC9	0.020	9.3E-06	1.895	0.429	hap2
rs4361680(T>A)	Chr7:18306817	HDAC9	0.020	9.3E-06	1.895	0.429	hap2
rs7459237(C>A)	Chr7:18307448	HDAC9	0.020	1.1E-05	1.889	0.427	hap2
rs302193(G>A)	Chr7:18298184	HDAC9	0.020	1.1E-05	1.822	0.445	hap2
rs1107693(G>A)	Chr7:18307906	HDAC9	0.020	1.2E-05	1.886	0.427	hap2
rs17138868(A>T)	Chr7:18294050	HDAC9	0.025	1.5E-05	1.834	0.433	hap2
rs302178(G>A)	Chr7:18310791	HDAC9	0.025	1.7E-05	1.798	0.440	hap2
rs73066378(G>A)	Chr7:18302581	HDAC9	0.025	1.8E-05	1.895	0.423	hap2
rs7794278(G>A)	Chr7:18316661	HDAC9	0.025	1.9E-05	1.822	0.429	hap1
rs28410383(A>G)	Chr7:18312712	HDAC9	0.027	1.9E-05	1.849	0.415	hap1
rs17138874(T>C)	Chr7:18297439	HDAC9	0.027	2.5E-05	1.792	0.461	hap2
rs6953847(C>T)	Chr7:18300086	HDAC9	0.029	2.7E-05	1.799	0.435	hap2
rs302194(T>C)	Chr7:18298160	HDAC9	0.029	2.9E-05	1.740	0.443	hap2
rs6956728(C>T)	Chr7:18316305	HDAC9	0.029	3.4E-05	1.755	0.433	hap1
rs17138904(T>C)	Chr7:18317529	HDAC9	0.029	3.5E-05	1.790	0.435	hap3
rs17138905(G>T)	Chr7:18317766	HDAC9	0.029	3.5E-05	1.790	0.435	hap3
rs1469575(G>A)	Chr7:18318382	HDAC9	0.029	3.5E-05	1.789	0.435	hap3
rs6461375(G>C)	Chr7:18318874	HDAC9	0.029	3.6E-05	1.760	0.448	hap3
rs917325(A>G)	Chr7:18288593	HDAC9	0.031	4.8E-05	1.742	0.447	hap4
rs302137(C>T)	Chr7:18293221	HDAC9	0.033	4.9E-05	1.686	0.491	hap2
rs215122(C>T)	Chr7:18295790	HDAC9	0.033	5.4E-05	1.688	0.466	hap2
rs12666671(T>C)	Chr7:18309456	HDAC9	0.041	6.2E-05	1.713	0.455	hap2
rs12666653(T>A)	Chr7:18309374	HDAC9	0.047	6.9E-05	1.704	0.458	hap2
rs2529749(A>C)	Chr7:18253060	HDAC9	0.047	7.8E-05	1.708	0.448	hap5
rs7779273(T>C)	Chr7:18301410	HDAC9	0.048	7.4E-05	1.692	0.460	hap2
rs917331(T>A)	Chr7:18253378	HDAC9	0.048	7.6E-05	1.709	0.448	hap5

Materials and Methods:

Image Processing: 341 publicly available GTEx histopathological thyroid images were downloaded from the Biospecimen Research Database (<http://brd.nci.nih.gov/image-search/searchhome>). For the purpose of segmenting individual tissue pieces only, the average intensity across color channels was calculated, and adaptive thresholding was performed to distinguish tissue from background. Interior holes in the tissue piece masks were filled using the *fillHull* function from the Bioconductor package *EBImage* (19). In preparation for feature extraction, Gaussian blurring was used on each color channel to smooth out pixel-level variation on a length scale smaller than the observed lymphocyte invasion phenotypes. 117 textural Haralick features were then extracted using *EBImage* by calculating 13 base Haralick features for each of the three RGB color channels, and across 3 Haralick scales by sampling every 1, 10, or 100 pixels. Image metadata was extracted to verify identical pixel dimensions and scanning instrumentation across the dataset.

Image Data Analysis: 702 thyroid pieces were identified from the image segmentation. 49 pieces with overly small areas were removed. Averaged Haralick feature values were then calculated across tissue pieces for each sample. In preparation for downstream model fitting, each distribution was \log_2 transformed and rescaled using a z-score. Five NA feature values were set to zero to avoid technical issues with the PCA fit.

Gene Expression Analysis: GTEx v6 RNA-Seq data was downloaded from dbGaP (phs000424.v6.p1, 2015-10-05 release) under approved protocol #9112. Lowly expressed genes were defined as those which had a read count of less than 1 CPM in more than half of the samples, and were removed. Counts were \log_2 transformed after adding 1 to the counts to avoid the problem of taking the logarithm of zero. Quantile normalization was performed using the R package *limma*. Prior to running the gene expression data through CIBERSORT, RNA extraction batch was corrected for using *limma*. PEER factors were calculated using the R package *peer* (github: PMBio/peer). The regression analysis with the image PCs was performed using linear model fits in R with the formula $PC = \alpha + \beta gene + \sum_k \gamma_k covariate_k + \epsilon$ where covariates included the known confounders and PEER factors, as discussed in the main text. For each gene, p-values were extracted from *t*-statistics associated with the gene expression coefficient β .

Genotype Analysis: GTEx genotype data was downloaded from dbGaP (phs000424.v6.p1, 2015-10-05 release) under approved protocol #9112. A total of 292 genotyped samples overlapped with the thyroid imaging samples. Minor allele frequency and missingness frequency filtering was performed using PLINK. QTL fits were performed using *MatrixEQTL*, which used the following model: $PC2 = \alpha + \beta genotype + \sum_k \gamma_k covariate_k + \epsilon$. Genotype could take on the values 0/1/2/NA, and covariates were as described in the main text. QTL significance was extracted from the *t*-statistics associated with the genotype coefficient β . The correction for ancestry was performed by including as covariates the first 3 PCs of a PCA on the genotype matrix computed from the full set of 450 GTEx patients. Haplotypes were identified using PLINK, with a maximum block size of 5,000 kb. The haplotype identification utilized the full set of 450 GTEx samples. Reported genomic coordinates are for human genome build GRCh37/hg19.

eMERGE data was obtained from the study “eMERGE Network Imputed GWAS for 41 Phenotypes” (dbGaP study accession phs000888.v1.p1) under approved protocol #13896 for the consent cohorts Health/Medical/Biomedical (N=18,621) and Health/Medical/Biomedical-

Genetic Studies Only-No Insurance Companies (N=15,911). Chromosome 7 genotype data from both cohorts was merged and 5,718 samples were selected based on their status as Chronic Autoimmune Hypothyroidism (accession phd004989.1) cases or controls. Inclusion criteria for hypothyroidism cases (C99269) included, but were not limited to, abnormal TSH/FT4 levels and the use of thyroid replacement medication. Exclusion criteria for hypothyroidism cases included secondary causes of hypothyroidism, hypothyroidism induced by surgery or radiation treatment, evidence of other thyroid diseases, or the use of thyroid-altering medication. Inclusion criteria for hypothyroidism controls (C99270) included no billing codes for hypothyroidism, no use of thyroid replacement medication, and normal TSH/FT4 levels. Exclusion criteria for controls included any evidence of hypo- or hyper-thyroidism, as well as other thyroid diseases, or the use of thyroid-altering medication. The Caucasian cohort was selected for further analysis, as described in the main text. Tests for Hardy-Weinberg Equilibrium (HWE) were performed using PLINK on the Caucasian cohort. One SNP in HDAC9 (*rs2058074*) had a HWE p-value less than 10^{-6} and was removed. Logistic regression models were fitted with the *glm* function in R using the following model: $\text{logit}(Y) = \alpha + \beta \text{genotype} + \sum_k \gamma_k \text{covariate}_k$. Y indicated the binary case/control disease status, genotype could take on the values 0/1/2/NA, and covariates were as described in the main text. Significance was extracted from the z-statistics associated with the genotype coefficient β . Reported genomic coordinates are for human genome build GRCh37/hg19.

Statistical Analysis: To test for associations between HT status and image PC2, we used the following logistic regression model: $\text{logit}(Y) = \alpha + \beta_1 \text{PC2} + \beta_2 \text{sex} + \beta_3 \text{age}$ where Y indicated the HT status according to GTEx pathology notes. The significance of association was extracted from the z-statistics associated with the beta coefficients. Due to co-linearity between sex and image PC2, associations between HT status and sex were assessed using a reduced model without the image PC2 term. Gene-set enrichment tests were performed using a Fisher's Exact Test for each Gene Ontology term. Multiple-testing correction across terms was performed using the method of Benjamini and Hochberg. Independent Weighting Hypothesis p-value correction on the image QTL results was performed using the Bioconductor software *IHW* (24). The permutation p-value reported for the eMERGE validation was estimated using the formula $(M + 1)/(N + 1)$ where M is the number of permutations with a correlation coefficient exceeding that of the non-permuted data, and N is the total number of permutations.