

Discovering and interpreting transcriptomic drivers of imaging traits using neural networks

Nova F. Smedley^{1,2,3}, Suzie El-Saden^{1,2}, and William Hsu^{1,2,3,4}

1. Medical & Imaging Informatics, University of California Los Angeles, USA
2. Department of Radiological Sciences, University of California Los Angeles, USA
3. Department of Bioengineering, University of California Los Angeles, USA
4. Bioinformatics IDP, University of California Los Angeles, USA

Abstract

Motivation: Cancer heterogeneity is observed at multiple biological levels. To improve our understanding of these differences and their relevance in medicine, approaches to link organ- and tissue-level information from diagnostic images and cellular-level information from genomics are needed. However, these “radiogenomic” studies often use linear, shallow models, depend on feature selection, or consider one gene at a time to map images to genes. Moreover, no study has systematically attempted to understand the molecular basis of imaging traits based on the interpretation of what the neural network has learned. These current studies are thus limited in their ability to understand the transcriptomic drivers of imaging traits, which could provide additional context for determining clinical traits, such as prognosis.

Results: We present an approach based on neural networks that takes high-dimensional gene expressions as input and performs nonlinear mapping to an imaging trait. To interpret the models, we propose gene masking and gene saliency to extract learned relationships from radiogenomic neural networks. In glioblastoma patients, our models outperform comparable classifiers (>0.10 AUC) and our interpretation methods were validated using a similar model to identify known relationships between genes and molecular subtypes. We found that imaging traits had specific transcription patterns, e.g., edema and genes related to cellular invasion, and 15 radiogenomic associations were predictive of survival. We demonstrate that neural networks can model transcriptomic heterogeneity to reflect differences in imaging and can be used to derive radiogenomic associations with clinical value.

Availability and implementation: <https://github.com/novasmedley/deepRadiogenomics>.

Contact: whsu@mednet.ucla.edu

1 Introduction

Radiogenomic mapping is the integration of traits observed on medical images and traits found at the molecular level, such as gene expression profiling [1, 2]. As such, the study of radiogenomics plays a role in precision medicine, where associations can describe prognosis or therapy response [3, 4, 5]. A common approach to radiogenomic mapping involves dimensionality reduction and pairwise comparisons [2, 6, 7, 8, 9, 10, 11] or predictive models, such as decision


trees [1, 12, 13, 14, 15] or linear regression [16, 17, 18, 19, 11]. Markedly, these approaches often require feature selection; assume linearity; and/or depend on pairwise associations, limiting their capacity to represent complex biological relationships.

Neural networks, with their ability to automatically learn nonlinear, hierarchical representations of large input spaces [20, 21], are alternate approaches for radiogenomics [5, 22, 23, 24, 25]. However, current applications of neural networks have focused on diagnostic images as inputs to predict a single gene status and have excluded gene interactions [5, 22, 24]. To the best of our knowledge, no prior studies have interpreted the neural networks to ascertain what radiogenomic associations are learned. Towards understanding the biological basis of imaging traits, we thus present an approach using the representational power of neural networks to model tumor transcriptomes and nonlinearly map genes to tumor imaging traits. No *a priori* selection is used on the transcriptome.

Importantly, we provide approaches for understanding radiogenomic neural networks based on visualization techniques, such as input masking [26] and saliency maps [27]. Although neural networks are considered “black boxes”, these methods probe the trained neural networks to understand relationships between gene expressions and imaging traits. The methods can identify cohort-level imaging-transcriptomic associations, which we refer to as *radiogenomic associations*, and patient-level radiogenomic associations, which we refer to as *radiogenomic traits*. We validated the associations and traits generated by a network in a classification task with known relationships, i.e., gene expressions (model input) and molecular subtypes (model output) [28].

As a use case, we study radiogenomic associations related to human-understandable imaging traits in magnetic resonance imaging (MRI) scans from patients with glioblastoma (GBM). GBM is a grade IV malignant brain tumor with poor prognosis, and for which imaging is heavily used in diagnosis, prognosis, and treatment assessment. MRI traits like tumor enhancement, non-contrast enhancement, edema, and necrosis in Fig. 1, describe some of the visual, phenotypic variations between patients as they are diagnosed and treated. We present here extracted associations found by the radiogenomic neural networks using our approach, compare against previous work to show both new and consistent findings, and establish the radiogenomic associations’ clinical value in estimating patient survival over clinical or imaging traits alone.

Figure 1: Examples of phenotypic differences observed in glioblastoma patients. Shown are single, axial images of pre-operative MRI scans from the TCGA-GBM cohort. Four MRI se-



2 Materials and methods

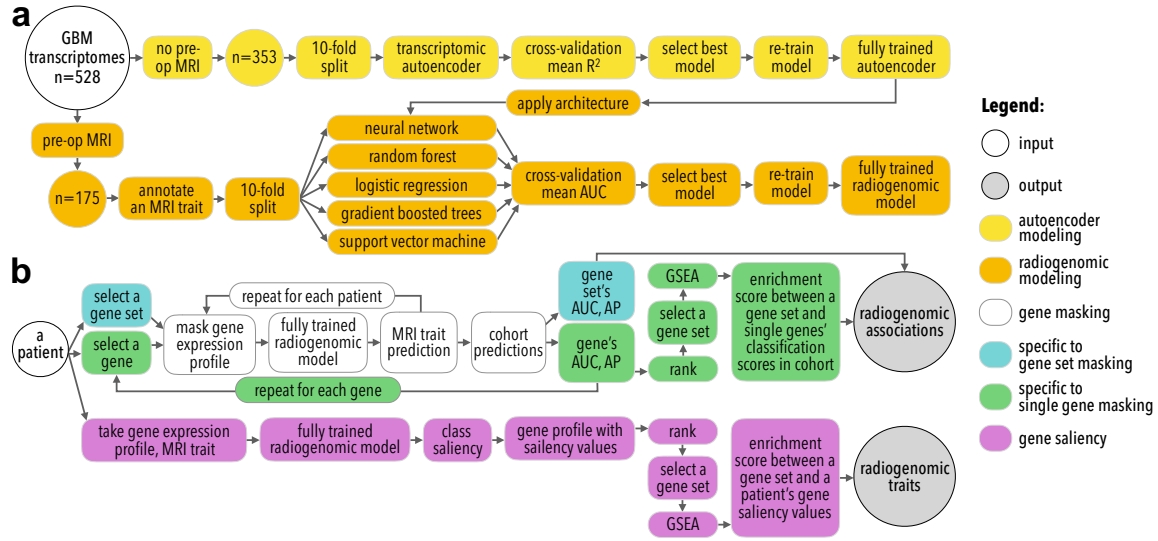
2.1 Gene expression

Transcriptomes were available for 528 GBM patients as part of The Cancer Genome Atlas (TCGA), see Supp. Table S1. Samples were primary tumors, untreated, and had $\geq 80\%$ cancer and $\leq 50\%$ necrotic cells [29]. Samples were analyzed by the Broad Institute on Affymetrix arrays, quantile normalized, and background corrected. Level 3 data were downloaded from the Genomic Data Commons at <https://gdc.cancer.gov/>. Each profile had 12,042 genes.

2.2 Magnetic resonance imaging

Medical images for 262 GBM patients were downloaded from The Cancer Imaging Archive (TCIA) [30]. Patients were matched based on barcodes shared by TCGA and TCIA. A board-

Figure 2: An overview of the study’s approaches to radiogenomic neural network (a) training and (b) interpretation, i.e., gene masking and gene saliency to extract radiogenomic associations and radiogenomic traits.




certified neuroradiologist, Dr. El-Saden (26 years of experience), evaluated images using the Osirix medical image viewer. An electronic form was used to record MRI traits according to the Visually Accessible Rembrandt Images (VASARI) feature guide [31]. 175 patients had pre-operative (pre-op) MRIs and transcriptomes, see Supp. Table S1. Six MRI traits were annotated from the pre-op studies. Traits were modeled as binary labels given the small sample sizes. Supp. Table S2 describes all labels and their percentages in the cohort.

2.3 Radiogenomic modeling

To map relationships between gene expression profiles and MRI traits, feed-forward neural networks [21] were used, see Fig. 3a–b. Each MRI trait was a binary classification task, where the positive class was the least frequent label (Supp. Table S2). Models were provided all 12,042 gene expressions as input vectors to classify each imaging trait, resulting in one model per trait. During training, early stopping with a patience of 200 epochs was used while monitoring the receiver operating characteristic curve, and the area under the curve (AUC) calculated. To help learning in the radiogenomic model, an autoencoder was used since many more patients had transcriptomic data. The radiogenomic neural networks were pretrained using weights transferred from a deep transcriptomic autoencoder trained on a separate subset of 353 patients. The transcriptome dataset consisted of TCGA-GBM patients with transcriptomes, but no pre-op MRIs and thus excluded the radiogenomic dataset.

Figure 3: The radiogenomic neural network’s (a) architecture, (b) transfer learning with a deep transcriptomic autoencoder, and interpretation methods through (c) gene masking and (d) gene saliency. Pretrained weights learned in the autoencoder were transferred to a radiogenomic model, where weights were frozen (non-trainable, long red arrows) and/or fine-tuned (trainable, dashed red arrow) during radiogenomic training.



The transcriptomic autoencoder takes a gene expression profile as input, compresses the information through three encoding layers, and then decodes the information to reconstruct the transcriptome. Early stopping was used in training and monitored the mean coefficient of determination (R^2) of each gene. The trained autoencoder weights, along with the gene pre-processing parameters, were then used as non-random weight initialization (i.e., weights *can* be fine-tuned during training) and/or frozen weights (i.e., weights *cannot* be fine-tuned during training) in the radiogenomic models.

Performance was estimated with 10-fold cross-validation with sample weighting and stratified fold splitting. Each gene expression was mean subtracted and divided by its standard deviation, a process performed at each fold split. Hyperparameters were optimized via grid search. An illustration of the overall methods for radiogenomic model training is shown in Fig. 2a. The autoencoder achieved a mean validation of 0.45 R^2 in cross-validation and 0.61 R^2 in retraining, see Supp. Fig. S2. Radiogenomic models were then constructed to have the same encoding architecture, activation function, and optimizer as the best performing autoencoder.

Neural networks were trained on NVIDIA Tesla K80 and V100 GPUs through Amazon Web Services using Python 3.6, Keras 2.2.4 [32], and TensorFlow 1.12.0 on a Ubuntu 16.04 machine. Other classifiers were implemented via XgBoost 0.80 [33] and sklearn

0.20.0 [34]. For more modeling details, see Supp. Table S3.

2.4 Bootstrapping

The best performing models of each model type (see Supp. Tables S3) were evaluated on bootstrapped datasets to measure classification performance variability within and between model types. For each bootstrap, a radiogenomic dataset was split for 10-fold cross-validation using a different seed. In each split, training and validation samples were separately resampled with replacement to obtain bootstrapped sets. Sets were resampled if not all classes were observed. Each model was trained and compared on the same bootstrapped dataset using the same procedure in the aforementioned methods. This process was repeated 100 times for each MRI trait.

2.5 Molecular subtype modeling

Molecular subtypes and their gene sets were downloaded from [28], which contained 840 genes to describe four GBM subtypes, i.e., classical, mesenchymal, proneural, and neural. Of the 528 patients, 171 had subtype labels and Affymetrix profiles, see Supp. Tables S1 and S2. A neural network was trained to predict subtypes using gene expression profiles. The four subtypes were modeled as multi-class classification task via one-hot-encoded labels. Gene expressions pre-processing and model training were performed in the same manner as the radiogenomic models, mainly, 10-fold cross-validation and hyperparameter grid search, see Supp. Tables S4.

2.6 Gene masking

Masking is a sensitivity analysis where the actual value of one or more components of the input are kept while all others are replaced with zeros. The goal is to determine the impact that the kept input components have on the end classification; this procedure was previously described in [26]. Here, we define “gene masking” to extract radiogenomic associations from a neural network, see Fig. 3c. For each individual, the gene expression values for a particular gene set were kept while all other expressions were replaced with zeros. The masked profiles were pushed through a fully trained neural network and the output, i.e., a class probability based on using genes from the gene set, was recorded. After repeating this process for each patient, classification performance was calculated. Each gene set was evaluated by AUC and average

precision (AP) to measure the strength of a radiogenomic association. As such, gene masking reported radiogenomic associations that were generally predictive of an MRI trait in the entire cohort.

In *single gene* masking, each gene was additionally used in gene set enrichment analysis (GSEA) [35] (ranked by AP or AUC). We also used *gene set* masking, where predefined gene sets smaller than 500 genes (see GSEA methods) from the Broad Institute’s molecular signatures database (MSigDB, v6.2) [36], molecular subtypes [28], and brain cell types and phenotypes [37, 38, 39] taken from [40] were used. MSigDB was also queried for gene sets that include the 22 genes characterized as potential contributors of GBM tumorigenesis [29, 41], see Supp. information. The top performing genes or gene sets for each MRI trait were visualized together by clustering classification scores using *pheatmap* in R.

2.7 Gene saliency

Class saliency is a visualization technique used to compute the gradient an output class prediction with respect to an input via back-propagation [27, 42]. Thus, class saliency identifies the relevant input components whose values would affect the positive class probability in a neural network. Here, we define “gene saliency” as the genes whose change in expression would increase the model’s belief of the positive class label, see Fig. 3d. In each model, salient genes are derived for each patient, ranked, and used in GSEA to determine if a gene set is relevant to predicting his/her MRI trait. Subsequently, positive enrichment between a single patient’s salient genes and a gene set is defined as a “radiogenomic trait.” For example, the edema model was probed to identify a single patient’s salient genes. The most salient genes were the genes that increased the probability of edema being $\geq 1/3$. If GSEA found the salient genes were enriched by a gene set, then the prediction of the patient’s edema was related to the gene set, i.e., the patient has the radiogenomic trait between the gene set and $\geq 1/3$ edema. Saliency was implemented using *keras-vis* 0.4.1 [42]. The input range was determined by the gene expression range in the dataset, and *guided* was used as the *backprop_modifier*; other parameters were set to default. In the subtype neural network, gene saliency was repeated for each class as one-versus-others. Fig. 2b depicts the overall process for gene masking compared to gene saliency.

2.8 Gene set enrichment analysis

Pre-ranked GSEA [35] was implemented using `fgsea` [43]. Gene sets were parameterized by the recommended maximum size of 500 genes in a gene set, a minimum size of 15, and 10,000 permutations. Genes were ranked via single gene masking classification scores or gene saliency values for a patient. Enrichments were significant at an adjusted p-value < 0.05 [43]. Correlation between a gene expression and an imaging trait was used for comparison. Clustering of enrichment scores was performed as previously defined.

2.9 Survival

Clinical data from TCGA was used to define patient outcomes. Patient covariates included gender, race (binned as white and non-white), and age at initial diagnosis (binned above or below the median value in all 528 patients). Overall survival (OS) and time-to-death events were in the `patient` file. Progression-free survival (PFS) outcomes were defined by the `followup` file, where all event types with days-to-event data were considered, but only the earliest event was used (see also Supp. information).

Cox proportional hazards models were used to estimate univariate and adjusted hazard ratios (HR). Radiogenomic traits enriched in at least five patients were kept. Adjusted HRs were estimated by a Cox model with backward feature selection and forced to keep all three patient covariates, while free to choose any MRI or radiogenomic trait.


Survival analyses were done in R using the `survival` and `survminer`. Kaplan-Meier estimates were obtained with `survfit`, `ggsurvplot`, and `survdiff`. Patients with missing information were removed and the log-rank test was used to test for significant differences between groups. Cox models were obtained with `coxph` and feature selection performed with `step`.

3 Results

3.1 Neural networks achieve best performance in classifying MRI traits

Neural networks were better at estimating MRI traits than all other classifiers, see Fig. 4a. In predicting proportions of nCET, necrosis, and edema, the first hidden layer used frozen pre-trained weights and only the last two hidden layers were used for fine-tuning, see Supp. Table

Figure 4: Radiogenomic models performances. (a) Observed 10-fold cross-validation performances. (b) Performance differences between a neural network and another model in 100 bootstrapped datasets. Notation: neural network (nn), gradient boosted trees (gbt), random forest (rf), support vector machines (svm), logistic regression (logit).



1. Bootstrapping showed neural networks had higher performances, where 95% confidence intervals (CI) in Fig. 4b indicate neural networks outperformed other models by more than 0.10 AUC. Supp. Table S5 and Figs S3–S5 have further modeling results.

3.2 Neural networks correctly learned known associations between gene expressions and molecular subtypes

To verify the relationships learned in the neural network's layers, a molecular subtype model was trained. The model achieved a micro-averaged AUC of 0.994 in 171 patients, see Supp. Table S6. Gene masking with subtypes gene sets showed the neural network was able to predict subtype classes with high certainty, see Fig. 5. For example, when the model was given the expressions of the 216 mesenchymal genes, subtype probabilities approached 1 or 0 and often corresponded to the correct subtype (0.90 AUC, 0.89 AP). Performance improved when the model was given all 840 subtype genes (0.99 AUC, 0.98 AP). Conversely, given the expressions of 200 random, non-subtype genes, the model was less certain (probabilities away from 1 or 0) and a fully trained performance of 1.0 AP dropped to 0.68 AP, see Supp. Table S7.

The majority of the top 20 predictive genes in each subtype class belonged to the original subtype class definition, see Fig. 6a. For example, 18 of the top 20 genes for predicting the

Table 1: Neural network hyperparameters. Layers refers to the depth of hidden layers in the radiogenomic model that used pretrained weights from the autoencoder (AE), e.g., two AE layers indicate the first two layers used pretrained weights. Retrain refers to models trained on the full dataset.

label	CV means		retrain		architecture	opt	act	drop		
	R^2	eph	R^2	eph						
transcriptome	0.45	467	0.61	486	4000, 2000, 1000	Adadelata	tanh	0.0	layers	
label	AUC	eph	AUC	eph	architecture	opt	act	drop	AE	frozen
enhancing	0.72	38	1.00	14	4000, 2000, 1000	Adadelata	tanh	0.6	3	0
nCET	0.83	38	1.00	11	4000, 2000, 1000	Adadelata	tanh	0.0	1	1
necrosis	0.75	44	1.00	11	4000, 2000, 1000	Adadelata	tanh	0.0	1	1
edema	0.78	109	1.00	16	4000, 2000, 1000	Adadelata	tanh	0.0	1	1
infiltrative	0.78	70	1.00	12	4000, 2000, 1000	Adadelata	tanh	0.0	2	1
focal	0.85	44	1.00	12	4000, 2000, 1000	Adadelata	tanh	0.6	3	0
subtype	0.994	14	0.998	66	3000, 1500, 750	Nadam	sigmoid	0.4	-	-

eph (epoch), opt (optimizer), act (activation), drop (dropout).

proneural subtype were a part of the proneural gene set and each had least 0.80 AP and a 0.80 AUC. Of the top 500 genes ranked by AP, 270 genes (54%) were subtype genes; this represented 32% of all subtype genes, see Fig. 6b. As expected, subtype genes were predictive of more than one subtype. Similarly, GSEA showed the most predictive single genes for each subtype prediction were significantly (adjusted p-value < 0.05) and positively enriched by the corresponding subtype gene set, see Fig. 6c. This observation was corroborated in GSEA using ranked genes based on the correlation, see Fig. 6d.


3.3 Genes driving the prediction of MRI traits

3.3.1 Enhancing tumor

Tumor enhancement was measured on T1W+Gd images. Low grade or well-differentiated brain tumors tend to generate blood vessels with intact blood-brain barriers (BBB) and do not enhance. Poorly differentiated or more aggressive tumors, like GBM, generate leaky blood vessels without an intact BBB and enhance on T1W+Gd images.

Enhancement was found to be associated with growth, immune responses, hormones, the extracellular matrix (ECM), vasculature, and kinase activity in gene masking, see Table 2. The ECM association included gene expressions of ECM-related proteins [44], see Supp. Fig. S18. Of the MSigDB hallmark gene sets, apical junction (cellular components, adherens and tight junctions), IL2/STAT5 signaling immune response activation), complement system, early and late responses to estrogen (associated with ESR1 expression), and heme metabolism (erythroid


Figure 5: Gene masking in the fully trained subtype neural network. The model’s (a) estimated subtype probabilities, where each row was a patient and grouped by their true subtype; and (b) classification performance measured by AP in gene set masking, where each row was a gene set and each column was the subtype prediction. The random gene set excluded any gene included in a subtype set. For visualization purposes, rows were sorted by the model’s estimated mesenchymal probability. Notation: classical (CL), mesenchymal (MES), neural (NL), proneural (PN), coverage (percent of gene set that exist in gene expression profiles). For more gene masking, see Supp. Figs. S6–S8.



differentiation, STAT5 activation) [45] were most predictive of enhancement, see Fig. 7a. Gene Ontology (GO) gene sets related to GBM-abnormalities support the association of growth, immune system, and hormones with enhancement.


In single gene masking, enhancement was best predicted by *SNTB1* (0.67 AUC, 0.58 AP) and *B4GALT6* (0.64 AUC, 0.60 AP), see Supp. Fig. S10. *SNTB1*, a cytoskeletal protein, was down-regulated in a GBM cells study [46] and a potential binder to *PTPRZ*, a protein contributing to glioma cell growth [47].

Figure 6: Single gene masking in the subtype model: (a) the top 20 genes used to predict each subtype; (b) the percent of subtype genes covered in the top N genes; (c) GSEA with genes ranked by average precision, where positive enrichment indicated the subtype gene set was correlated with high average precision and vice versa; and (d) GSEA with genes ranked by correlation with label, where positive enrichment indicated the subtype gene set was correlated with subtype and vice versa.



In previous GBM radiogenomic studies, enhancement was associated with hypoxia, ECM, angiogenesis in 22 patients [2]; Biocarta pathways and genes, *C1orf172*, *CAMSAP2*, *KCNK3*, and *LTBP1* in 23 patients [10]; and *EGFR* copy number amplification in 75 patients [7]. Gene sets involving the ECM, *EGFR*, *C1orf172*, *KCNK3*, and *LTBP1* were confirmed to have performances greater than 0.70 in both AUC and AP, see Supp. Fig. S12a.

Figure 7: Gene masking of the radiogenomic models with the MSigDB hallmark gene sets [45]. **(a)** Model performance in gene set masking. Shown are the top five gene sets ranked by average precision in each MRI trait, see also Supp. Fig. S9. **(b)** Enrichment among genes ranked by average precision in single gene masking. Positive enrichment indicated gene sets were predictive of an MRI trait and negative enrichment indicated the opposite. Shown are hallmarks with at least one significant enrichment.



3.3.2 Edema

Tumor edema was identified as abnormal hyperintensity on FLAIR or T2W images. Edema often co-occurs with enhancement, implying a more aggressive tumor and does not typically occur in low grade brain tumors. Edema also results from leaky capillaries and usually surrounds the tumor, spreading within the white matter. Edema suggests an inflammatory and/or immune response to a malignant tumor, which is essentially a foreign body when highly dedifferentiated.

Edema was associated with epithelial mesenchymal transition (EMT, metastasis and invasion), cell differentiation, and growth, see Table 2. p53 pathway (cell cycle, death), myogenesis, apical junction, heme metabolism, and glycolysis (cell metabolism) were the top hallmark gene sets, see Fig. 7a. GO terms relating to cell differentiation, death and adhesion with ≥ 0.80

Table 2: Summary of transcriptomic drivers of MRI traits in GBM patients. Shown are the top five hallmark gene sets ranked by AUC or AP and compared against gene sets related to prior GBM work, see Supp. Figs S11 (gene abnormalities [29, 41]) and S12–S15 (radiogenomics). Note: common themes can comprise of different gene sets.

MRI trait	transcriptomic drivers		AUC	AP	see also
	theme	gene set (collection*, query†)			
enhancing	growth/death	growth (GO, <i>PTEN</i>)	0.86	0.84	
		sensory organ development (GO, <i>EGFR</i> , <i>KCNK3</i>)	0.85	0.84	[10, 7]
	immune system	IL2/STAT5 signaling (H)	0.77	0.76	
		complement system (H)	0.79	0.75	
		activation of immune response (GO, <i>PTEN</i>)	0.90	0.89	
		leukocyte & lymphocyte activation (GO, <i>PIK3R1</i>)	0.86, 0.85	0.85, 0.83	
	hormones	immune effector process (GO, <i>PIK3CA</i>)	0.87	0.84	
		early & late responses to estrogen (H)	0.79, 0.78	0.73, 0.73	
		response to steroid hormone (GO, <i>RBI</i>)	0.88	0.88	
		regulation of hormone levels (GO, <i>PARK2</i>)	0.87	0.84	
	ECM	related to ECM proteins (C, ECM)	0.77–0.84	0.73–0.76	[2]
		apical junction (H)	0.80	0.75	
	vasculature	heme metabolism (H)	0.77	0.65	
		vasculature & heart development (GO, <i>LTBP1</i>)	0.81, 0.78	0.80, 0.77	[10]
edema	kinases activity	multiple (GO, <i>EGFR</i> , <i>LTBP1</i> , <i>KCNK3</i>)	all 0.87	all 0.85	[10, 7]
		EMT (H)	0.80	0.80	
		positive regulation of locomotion (GO, <i>EGFR</i> , <i>POSTN</i>)	0.80	0.81	[6]
		taxis (GO, <i>CXCL12</i> , <i>KIF5C</i>)	0.80	0.80	[6]
		apical junction (H)	0.77	0.74	
		related to cell adhesion (GO, <i>CDKN2A</i> , <i>EGFR</i> , <i>CTNNA2</i>)	0.75–0.83	0.78–0.82	[6]
	growth/death	p53 pathway (H)	0.77	0.77	
		autophagy (GO, <i>CDKN2A</i>)	0.75	0.76	
		myogenesis (H)	0.75	0.76	
		urogenital system development (GO, <i>PTEN</i>)	0.81	0.81	
		muscle structure development (GO, <i>COL6A3</i>)	0.80	0.80	[6]
		response to growth factor (GO, <i>EGFR</i> , <i>POSTN</i>)	0.76	0.81	[6]
	vasculature	heme metabolism (H)	0.77	0.73	
	differentiation	central nervous system neuron differentiation (GO, <i>PTEN</i>)	0.79	0.81	
		cell differentiation (GO, <i>MET</i>)	0.79	0.81	
		stem cell differentiation (GO, <i>CDK6</i>)	0.80	0.80	
		regulation of cell activation (GO, <i>CDKN2A</i>)	0.83	0.82	
	immune system	neg. regulation of immune system (GO, <i>CDKN2A</i> , <i>CXCL12</i>)	0.82	0.81	[6]
		immune effector process (GO, <i>PIK3CA</i>)	0.80	0.81	
	other	glycolysis (H)	0.76	0.70	
		mitotic spindle (H)	0.78	0.70	
		DNA repair (H)	0.77	0.66	
		G2M checkpoint (H)	0.72	0.63	
nCET	cell cycle	regulation of mitotic cell cycle (GO, <i>TP53</i>)	0.78	0.71	
		p53 pathway (H)	0.76	0.65	
		urogenital & vasculature development (GO, <i>PTEN</i>)	0.81, 0.80	0.76, 0.71	
		neg. regulation of cell cycle (GO, <i>TP53</i>)	0.80	0.74	
	growth/death	reproductive system development (GO, <i>EGFR</i>)	0.80	0.72	
		UV response down (H)	0.78	0.59	
		response to radiation (GO, <i>TP53</i>)	0.83	0.73	
		glycerphospholip metabolism process (GO, <i>EGFR</i>)	0.75	0.74	
	UV response	small molecule catabolic process (GO, <i>PTEN</i>)	0.78	0.73	
	other				

* an MSigDB collection, where H=hallmarks, GO=Gene Ontology, and C=Canonical. † Gene sets queried from MSigDB using gene names or functions reported by previous work as keyword(s). neg.=negative.

AP, see Supp. Fig. S11b. Similar to the enhancement model, vasculature, immune system and *EGFR*-related processes (albeit through different GO terms) were apart of the most predictive gene sets.

Growth and metastasis was also found to be predictive of edema in single gene masking. *RAI2*, *ANXA2*, *POSTN* genes, all related to cell growth, were the top three single most predictive

Table 3: Summary of transcriptomic drivers of MRI traits in GBM patients, continued.

MRI trait	transcriptomic drivers		AUC	AP	see also
	theme	gene set (collection*, query†)			
necrosis	vasculature	heme metabolism (H)	0.72	0.61	
		growth/death			
		apoptosis (H)	0.71	0.58	
		apoptotic signaling pathway (GO, <i>IL4</i> , <i>TP53</i>)	0.72	0.63	[9]
		related to <i>TP53</i>	0.75–0.76	0.63–0.68	
		gland development (GO, <i>EGFR</i>)	0.76	0.65	
	immune system	IL6/JAK/STAT3 signaling (H)	0.67	0.56	
		leukocyte cell cell adhesion (GO, <i>IL4</i> , <i>ITGA5</i>)	0.76	0.59	[9, 10]
		regulation of leukocyte proliferation (GO, <i>CDKN2A</i>)	0.77	0.61	[7]
	others	xenobiotic metabolism (H)	0.76	0.65	
		related to <i>PTEN</i>	0.73–0.78	0.63–0.67	
		regulation of homeostatic process (GO, <i>NF1</i>)	0.78	0.65	
		glycolysis (H)	0.76	0.56	
		regulation of anatomical structure size (GO, <i>PTEN</i>)	0.96	0.88	
focal	growth/death	response to growth factor (GO, <i>EGFR</i>)	0.92	0.83	
		transport			
		secretion by cell (GO, <i>NF1</i>)	0.95	0.88	
		neg. regulation of transport (GO, <i>PTEN</i>)	0.96	0.87	
		regulation of cytoplasmic transport (GO, <i>TP53</i>)	0.95	0.85	
		monovalent inorganic cation transport (GO, <i>PARK2</i>)	0.95	0.81	
	response to	steroid hormone, lipid, & organic cyclic compound (GO, <i>RB1</i>)	0.94–0.96	0.81–0.87	
		vasculature development (GO, <i>PTEN</i>)	0.93	0.84	
	vasculature	muscle & circulatory system process (GO, <i>PIK3CA</i>)	0.94, 0.94	0.82, 0.83	
		oxygen			
		hypoxia (H)	0.85	0.61	
	others	genes down-regulated by <i>KRAS</i> (H)	0.88	0.64	
		protein heterodimerization activity (GO, <i>TP53</i>)	0.97	0.85	
infiltrative	oxygen	neg. regulation of intracellular signaling transduction (GO, <i>PTEN</i>)	0.96	0.84	
		synaptic signaling (GO, <i>PTEN</i>)	0.92	0.82	
		reactive oxygen species pathway (H)	0.71	0.50	
		response to oxygen levels (GO, <i>TP53</i>)	0.67	0.60	
	transport	neg. regulation of transport (GO, <i>PTEN</i> , <i>NFKBIA</i>)	0.74	0.64	[48]
		healing			
		wound healing (GO, <i>NF1</i>)	0.80	0.64	
		hemostasis (GO, <i>PIK3CA</i>)	0.78	0.59	
	growth/death	developmental growth (GO, <i>PTEN</i>)	0.74	0.62	
		spinal cord development (GO, <i>NF1</i>)	0.66	0.60	
	response to	response to drug (GO, <i>MDM2</i> , <i>MYC</i>)	0.75	0.60	
		response to inorganic substance (GO, <i>PTEN</i>)	0.73	0.61	[48]
	others	DNA repair (H)	0.70	0.58	
		ligase activity (GO, <i>MDM2</i>)	0.75	0.64	
		ubiquitin like protein transferase activity (GO, <i>MDM2</i>)	0.70	0.59	
		ubiquitin like protein ligase binding (GO, <i>NFKBIA</i>)	0.67	0.57	[48]
		related to protein & transcription factor complex (GO, <i>TP53</i>)	0.74–0.75	0.60–0.63	
		WNT signaling pathway (GO, <i>PTEN</i> , <i>MYC</i>)	0.73	0.62	[48]

* an MSigDB collection, where H=hallmarks, and GO=Gene Ontology. † Gene sets queried from MSigDB using gene names or functions reported by previous work as keyword(s). neg.=negative.

gene with 0.68–0.70 AUC and 0.60–0.64 AP. The top three ranked by AP were *MTSS1*, *LAMA5*, and *KLHDC3*, with 0.65–0.66 AP and 0.63–0.67 AUC, see Supp. Fig. S10. Both *MTSS1* and *LAMA5* were associated with metastasis [49].

GSEA showed significant enrichment in EMT, angiogenesis, androgen response (hormone), hedgehog signaling (including *MTSS1*), and xenobiotic metabolism (drug metabolism), see Fig. 7b. The appearance of drug metabolism could be due to the use of symptomatic relief drugs prior to surgery, such as corticosteroids for patients with neurologic symptoms caused by edema [50, 51].

Previously, *POSTN* was associated with edema in 78 GBM patients; the authors suggested

POSTN was regulated by miR-219 and contributed to cell migration or invasion [6]. GO gene sets related to the study's top five upregulated genes and microRNAs were found to be predictive of edema, see Supp. Figs. S13. Table 2 shows an overlap of gene set patterns between the study's findings and gene masking of the edema model. In particular, gene sets associated *POSTN*, cell taxis, and cell adhesion added to the association between edema and EMT.

3.3.3 Non-contrast enhancing tumor

Non-contrast enhancing tumor (nCET) was best identified on contrast-enhanced T1W and FLAIR or T2W images. nCET is typically lower grade tumor (better cellular differentiation, more closely resembling normal brain tissue), generating vessels with an intact BBB, and absent of contrast enhancement. While the abnormality on images is a mass-like neoplastic tissue, it is not rapidly dividing or aggressively dedifferentiating.

Cell cycle, growth, and radiation response were themes among the most predictive gene sets for nCET. Mitotic spindle, UV response down (genes down-regulated in response to ultraviolet radiation), DNA repair, and p53 pathway were predictive of nCET, see Fig. 7a. Of the gene sets related to GBM genomic alterations, the nCET model had a mix of ones found in the enhancing and edema model, see Fig. S11c, and supported transcription patterns found in hallmark gene sets in Table 2.

SGPL1 (0.73 AUC, 0.56 AP) and *DDR1* (0.66, 0.61 AP) were the top performing genes in single gene masking of the nCET model, see Fig. S10, where hypoxia, TNFA signaling via NFKB, and oxidative phosphorylation were significantly enriched, see Fig. 7b. The latter two gene sets were also identified in gene set masking, see Supp. Fig. S9.

3.3.4 Necrosis

Tumor necrosis was evaluated as the area of fluid signal intensity on T1W+Gd images. As tumors proliferate, they create new blood supply (angiogenesis) and/or expands to recruit blood from adjacent tissue. Subsequently, necrosis occurs, typically within the central portions of an aggressive tumor as the outer rim of enhancing surviving cells can be observed on MR images.

Vasculature, apoptosis, immune system, and homeostasis was associated with necrosis, see Table 3. Predictive GO terms for necrosis included several *TP53* and *PTEN* related processes.

Gene masking found some gene sets related to *IL4*, *WWTR1*, *RUNX3*, *ITGA5*, and *CDKN2A* (found in previous radiogenomic studies [9, 10, 7]) support the association between necrosis and apoptosis and the immune system. Earlier, drug metabolism was predictive of edema, but was also predictive of necrosis. Besides corticosteroids, antiepileptics can be prescribed patients who experience seizures from tumors [50].

In single gene masking, *CACNB2* (0.67 AUC, 0.51 AP) and *PACSIN3* (0.64 AUC, 0.51 AP) the most predictive genes for necrosis, see Supp. Fig. S10. Notably, the MYC targets hallmark was significantly enriched, see Fig. 7b.

3.3.5 Focal vs. non-focal

Focal vs. non-focal traits were determined via T1W+Gd and FLAIR or T2W images. Focal tumors appear in one region. Non-focal tumors included those described as multifocal, multicentric, or with gliomatosis cerebri. A *multifocal* tumor is one with separate enhancing regions that appear connected on FLAIR/T2W images with contiguous hyperintensity spreading via white matter tracts. A *multicentric* GBM has multiple enhancing or non-enhancing tumors growing synchronously without contiguity on FLAIR/T2W. Gliomatosis cerebri is a rare, diffusely infiltrating subtype and involve at least three cerebral lobes.

Focal traits were associated with growth, transport, vasculature, and hypoxia, see Table 3. Several of these involved *PTEN* and *RB1*. Secretion by cell includes genes potentially overlapped by others, e.g., *NF1*, *EGF*, and *VEGF*.

In general, focal traits were better predicted by GO gene sets related to GBM genes (Supp Fig. S11e) than hallmark gene sets (Fig. 7) or single genes (Supp. Fig. S10). These highly predictive GO gene sets (≥ 0.90 AUC and ≥ 0.80 AP) indicated broad tumor characteristics, e.g., proliferation (growth, response to growth factors, secretion of growth factors) resulting in the need for angiogenesis (vasculature development), were used by the focal model to determine focal vs. non-focal tumors.

3.3.6 Expansive vs. infiltrative

Expansive vs. infiltrative was measured as the ratio of T1/FLAIR abnormality on T1W and FLAIR or T2W images. Expansive tumors have similar distribution on T1W as on FLAIR/T2W; the closer the two, the better defined the tumor margins and the better for surgical resection.

Infiltrative tumors have FLAIR/T2W abnormality that is large compared to T1W abnormality, where the tumor is spreading through white matter tracts to cause large edema relative to the core tumor mass. Infiltrative traits indicate ill-defined tumor margins, less successful surgical debulking, and worse prognosis.

Infiltrative traits were best predicted by gene sets related to oxygen, transport, healing, and growth, see Table 3. Gene masking showed that GO gene sets were more predictive than hallmark gene sets (Fig. 7). Of the top GO gene sets, wound healing and hemostasis were the most predictive and several included *TP53*, *MDM2*, and *PTEN*. Notably, *MDM2* transcription is regulated by *TP53*.

Previous radiogenomic studies related to expansiveness or infiltrative traits found associations with *MYC*, *NFKBIA*, and immune cell gene modules [48]. The infiltrative model was masked with related gene sets and was able to predict infiltrative tumors with 0.50–0.70 AP, see Supp. Fig. S15.

In single gene masking, *ZBTB48* (0.68 AUC, 0.53 AP) and *PRTN3* (0.67 AUC, 0.57 AP) as the best single gene predictors in Supp. Fig. S10. For more gene masking, see Supp. Figs S16–S21.


3.4 Radiogenomic traits: Patient-specific radiogenomic associations

Gene masking was used to identify cohort-level radiogenomic associations as genes were ranked by their overall classification performance among all tumors. In contrast, gene saliency was measured for each patient and identified patient-level radiogenomic associations, termed 'radiogenomic traits.' For salient genes in the subtype model, see Supp. Fig. S25.

Classical subtype genes were salient in predicting larger proportions of necrosis, where 77 patients were enriched with classical genes in the necrosis model, see Fig. 8. Similarly, neural and proneural gene sets were associated with greater edema and nCET proportions, respectively.

The nCET model showed more than 40 patients had salient genes enriched by oligodendrocytes, mature astrocytes, and hypoxia gene sets. Larger edema proportions were associated with neurons and replicating fetal neurons genes. The anti-cell cycle genes (negatively correlated with the cell cycles genes, some of which were a part of the hypoxia gene set in [37]) were associated with prediction of the non-focal class (35 patients enriched). Patients were not

Figure 8: Radiogenomic traits: results of gene saliency applied to the neural networks. Patients were considered enriched for a gene set at an adjusted p-value < 0.05. Gene sets with at least 10 enriched patients in a model were shown. All enriched patients were positively enriched. Positive enrichment indicated that the gene set was among the most salient genes for predicting a single patient's imaging trait.



significantly enriched with cell type or phenotype in the enhancing model, possibly reflecting more tumor heterogeneity in patients with more enhancing and aggressive tumor.

The hypoxia and nCET association (66 enriched patients) was consistent with the aforementioned anti-cell and hypoxia findings and with gene masking analysis, where vasculature development was predictive of nCET at the cohort-level. The association between hypoxia and greater proportions of nCET may be linked to lower-grade cells in the beginning stages of aggressive tumor growth and therefore responding to the beginning stages of hypoxic conditions and driving angiogenesis. Xenobiotic metabolism was enriched in 15 patients for predicting

larger nCET, suggesting an increase in tumor size will result in an increased dosage of drugs administered prior to surgical resection or biopsy [51].

Interestingly, the edema model showed only two hallmarks enriched by more than 10 patients. Although gene masking showed the edema model had high overall performance with the EMT gene set, other genes that are not associated with a predefined gene set may have been more influential in predicting each individual patient's edema proportions. In fact, the EMT hallmark was more associated with the radiogenomic model's belief of an infiltrative tumor in 27 patients. This subset of patients support the hypothesis that tumor cells with alterations in EMT-related genes are driving the observation of higher edema proportions than tumor cell proportions. Glycolysis and hypoxia hallmarks were also moderately (< 25 patients) associated with infiltrative tumors. TNFA signaling via NFkB was associated with larger proportions of necrosis in 34 patients.

Chromosomal aberrations have been reported in GBM [52, 28]. There were 86 and 82 patients who were enriched by the chr19p13 gene set in predicting their necrosis and infiltrative traits, respectively. Genes in chr19q13, chr1p35, and chr6q27 were also salient to infiltrative tumors. chr3p21 and chr3p22 gene sets were also salient for greater nCET proportions.

3.4.1 Radiogenomic traits with survival implications

Of the 175 patients with radiogenomic data, 127 had all six MRI traits labeled and outcomes data. Given that only a subset of these individuals had clinical, imaging, and transcriptomic data to perform a survival analysis, we tested whether this subset of patients had any differences in outcomes compared to the transcriptome cohort ($n=528$). There were no overall survival (OS) or progression-free-survival (PFS) differences between the transcriptome cohort, the subset of patients with MRI traits, and the subset of patients with all six MRI traits, see Supp. Fig. S27. In building the multivariable Cox models, three clinical traits, six MRI traits, and 54 radiogenomic traits were considered, see Supp. Fig. S29.

Figure 9 show patients had significant differences when dichotomized by their radiogenomic traits. Patients with neural genes among the most salient genes for predicting larger nCET had *better* PFS compared to those who did not. Likewise, patients had significantly *worse* OS when chr3p21 genes were important in predicting nCET proportion was $\geq 1/3$, and when chr19p13 or chr1p35 genes were important in predicting infiltration. In contrast, dichotomizing patients

Figure 9: Overall survival (OS) and progression-free survival (PFS) dichomotized by (a) imaging traits compared to (b) radiogenomic traits. Patients split based on the association between the nCET trait and neural (NL) subtype genes had a median PFS of 0.96 years vs. 0.52 years (161 day difference). The median OS was 1.19 years vs. 0.91 years (101 day difference) when split by the nCET and chr3p21 trait, 1.18 years vs. 1.14 years (15 day difference) split by the infiltrative and chr19p13 trait, and 1.19 years vs. 0.85 years (125 day difference) split by the infiltrative and chr1p35 trait.




Table 4: Cox regression analysis of traits associated with overall survival and progression-free survival. Both OS and PFS multivariable models had $p < 0.001$ in the likelihood ratio test.

	univariate HR	adjusted HR	
	(95%CI)	(95%CI)	p-value
Overall survival (n=127, deaths=107)			
clinical			
gender is male	0.95 (0.64–1.41)	0.80 (0.53–1.20)	0.280
race is white	1.69 (0.85–3.37)	2.30 (1.09–4.85)	0.029*
diagnosis age is below median	0.82 (0.56–1.21)	0.79 (0.53–1.20)	0.273
radiogenomic			
infiltrative + chr1p35	2.06 (0.98–4.31)	2.06 (0.95–4.43)	0.066
edema + endothelial [†]	2.07 (0.76–5.67)	4.36 (1.47–12.89)	0.008*
necrosis + GBM core astrocytes [†]	0.46 (0.14–1.49)	0.11 (0.03–0.45)	0.002*
necrosis + epithelial mesenchymal transition	1.40 (0.80–2.43)	3.45 (1.75–6.82)	<0.001*
nCET + myogenesis	2.85 (0.89–9.07)	10.73 (2.48–46.51)	0.002*
necrosis + MYC targets (v2)	0.76 (0.31–1.87)	0.31 (0.10–0.98)	0.045*
infiltrative + mTORC1 signaling	1.87 (0.76–4.61)	2.38 (0.93–6.10)	0.071
Progression-free survival (n=127, progressions=88)			
clinical			
gender is male	1.03 (0.65–1.62)	0.80 (0.48–1.34)	0.394
race is white	1.28 (0.62–2.65)	1.96 (0.89–4.30)	0.094
diagnosis age is below median	0.98 (0.64–1.50)	1.25 (0.77–2.04)	0.373
imaging			
tumor was infiltrative	0.92 (0.58–1.45)	1.61 (0.96–2.71)	0.072
radiogenomic			
infiltrative + chr1p35	2.06 (0.98–4.31)	4.20 (1.89–9.33)	<0.001*
infiltrative + epithelial mesenchymal transition	2.57 (1.23–5.39)	2.24 (1.27–3.96)	0.006*
necrosis + MYC targets (v2)	0.21 (0.03–1.49)	0.06 (0.01–0.47)	0.008*
edema + fetal neurons replicating ^{††}	1.48 (0.93–2.35)	2.50 (1.44–4.33)	0.001*
infiltrative + TGF- β signaling	1.51 (0.80–2.86)	1.91 (0.94–3.92)	0.076
edema + chr18p11	2.22 (0.80–6.15)	5.79 (1.87–17.99)	0.002*
edema + G2M checkpoint	0.75 (0.36–1.56)	0.46 (0.20–1.05)	0.066
nCET + chr22q13	0.35 (0.09–1.42)	0.36 (0.08–1.60)	0.180
infiltrative + chr6q27	1.33 (0.76–2.33)	1.79 (0.98–3.30)	0.060
necrosis + p53 pathway	1.06 (0.39–2.91)	2.44 (0.81–7.39)	0.114

* $p < 0.05$, [†]gene set from [39], ^{††}gene set [38]

based solely on individual MRI traits had no OS or PFS differences, except in the counterintuitive case of expansive vs. infiltrative, see Supp. Fig. S28. Infiltrative tumors had a univariate HR of 0.92 when estimating PFS, see Table 4. However, after adjusting for patient covariates and radiogenomic traits, infiltrative tumors had an adjusted HR of 1.61 and correctly follows the intuition that infiltrative tumors would have a higher probability of progression than expansive tumors.

Males and non-white races had better OS and PFS, while patients diagnosed below the median age had better OS, but worse PFS. The final Cox model consisted of six significant traits, five radiogenomic traits and race when estimating OS, see Table 4. The final PFS model had five significant radiogenomic traits. In comparison, a Cox model with clinical and imaging traits had no significant factors in estimating OS or PFS. The survival analyses suggest that radiogenomic traits extracted from neural network models have prognostic value.

4 Discussion

We demonstrate how deep neural networks can be used to discover radiogenomic associations. *First*, we predict imaging traits using gene expression profiles, showing that our neural network-based approaches outperforms other classifiers. We also illustrate the benefit of transfer learning to train radiogenomic neural networks using a transcriptomic autoencoder modeled on a much larger cohort to address the impedance of relatively small radiogenomic datasets. *Second*, we present methods based on input masking and class saliency that facilitate interpretation of radiogenomic associations, providing a way to understand the results of an otherwise “black box” method, which is the main criticism against neural networks. *Third*, using our network analysis techniques, we identify pertinent gene expressions that may act as transcriptomic drivers for each imaging trait. We put forth a set of potential imaging surrogates that provide a clearer biological basis of commonly assessed imaging phenotypes in GBM and relate them to trends in patients’ overall and progression-free survival.

Gene masking identifies cohort-level radiogenomic associations, where the strength of association was measured by the model’s classification when only a subset or one gene was used. Radiogenomic associations have common themes related to major GBM candidate driver genes and terms, e.g., cell growth and vasculature. However, each MRI trait also show specificity towards components of these general themes, e.g., different functionalities of *EGFR* or cell death by autophagy in the edema model compared to apoptosis in the necrosis model. We identify unique associations between imaging traits and different themes: edema is associated with cell invasion and differentiation; enhancement is associated with immune system processes and hormones; nCET is associated with cell cycle and UV response; necrosis is associated with apoptosis; and focal is associated with cell transport and response to certain compounds.

Prior radiogenomic studies have mainly reported cohort-level associations. In reality, multiple gene expression profiles, when influenced by different environmental factors, may lead to the same observed imaging trait. Towards this end, gene saliency was used to identify patient-level radiogenomic traits.

With gene saliency, each patient has his/her own list of relevant genes for each imaging trait; it is then determined if the patient’s salient genes are significantly associated with a gene set. We describe subsets of patients with common radiogenomic associations that are not apparent in gene masking, such as the association between infiltrative traits and epithelial-

mesenchymal transition genes or larger nCET proportions and drug metabolism. Some of these radiogenomic traits are significant factors in predicting patient survival.

Furthermore, we validate our modeling approach by training a neural network to predict molecular subtypes. We report an experiment that evaluates model’s ability to learn meaningful relationships. Not only does the subtype model achieve near-perfect classification, the model is able to select genes relevant to each subtype among 12,042 genes. In the radiogenomic models, we validate our radiogenomic associations with prior GBM studies in radiogenomics and genomics and found corresponding relationships. We also identify new findings that have not been widely reported in radiogenomics due to the ability of gene saliency to provide patient-specific radiogenomic traits and the inclusion of the entire gene profile in our models. These results support the neural network’s abilities to identify associations with existing domain knowledge and to suggest potential starting points for further investigation.

We recognize the limits of radiogenomic analysis, particularly in terms of small sample size, limited tissue sampling of a heterogeneous tumor, and limited follow-up information. Sample size is an inherent challenge in radiogenomics. TCGA tends to have the most radiogenomic data, but lacks detailed clinical data. While larger cohorts do exist, tumors are across multiple grades [8] and do not use molecular profiling [5]. These limitations may be addressed as the cost of high-throughput platforms decreases and multiple tumor regions are sampled [40]. With 528 gene expression profiles and a radiogenomic subset of 175, we show that neural networks can model transcriptomic heterogeneity to reflect phenotypic differences in imaging. The VASARI feature set is also limited in that it provides a gross categorization of imaging features and only one experienced reader’s annotations of the imaging data is used. A more comprehensive analysis that includes quantitative (radiomic) traits may be warranted. Finally, the radiogenomic associations are only hypothesized and not proven through experimentation, though we attempt to compare our discovered associations with those that have been previously reported in literature. To validate the identified associations, cell and animal studies would allow controlled experiments between genes and imaging phenotypes [53].

5 Conclusion

Using a neural network-based approach to radiogenomic mapping, we highlight the representational and discriminative capacity of neural networks to model the high-dimensional, non-

linear, and correlative nature of gene expressions to predict typical GBM imaging traits. We demonstrate the use of neural network interpretation techniques, e.g., input masking and class saliency to understand what the model has learned and to extract relevant radiogenomic associations. The learned radiogenomic associations may point to potential transcriptomic drivers of imaging traits and could further clarify the understanding of the relationship between two often disjoint datasets, gene expression profiling and medical imaging. As such, prognostication and treatment options may be further individualized, where a targeted pathway could be considered in the selection of an appropriately tailored chemotherapeutic agent.

6 Acknowledgements

We are thankful for the feedback from the faculty and students of the Medical & Imaging Informatics group, discussions about performance differences with biostatisticians, Drs. James Sayre and Audrey Winter, and correspondences with TCIA members to better understand the data.

7 Funding

This work was supported in part by the National Institutes of Health [F31CA221061 and T32EB016640 to N.F.S., R01CA157553 to N.F.S., W.H., S.E.S.]; the National Science Foundation [#1722516 to W.H.]; Amazon Web Services and UCLA Department of Computational Medicine partnership to W.H.; and the Integrated Diagnostics Program, jointly funded by the Department of Radiological Sciences and Pathology & Laboratory Medicine to N.F.S., W.H.. *Conflict of Interest:* none declared.

References

- [1] Segal, E. *et al.* Decoding global gene expression programs in liver cancer by noninvasive imaging. *Nature Biotechnology* **25**, 675–80 (2007).
- [2] Diehn, M. *et al.* Identification of noninvasive imaging surrogates for brain tumor gene-expression modules. *Proceedings of the National Academy of Sciences of the United States of America* **105**, 5213–8 (2008).
- [3] Pope, W. B. *et al.* Relationship between gene expression and enhancement in glioblastoma multiforme: Exploratory DNA microarray analysis. *Radiology* **249**, 268–277 (2008).
- [4] Naeini, K. M. *et al.* Identifying the mesenchymal molecular subtype of glioblastoma using quantitative volumetric analysis of anatomic magnetic resonance images. *Neuro-Oncology* **15**, 626–34 (2013).
- [5] Chang, K. *et al.* Residual convolutional neural network for the determination of IDH status in low- and high-grade gliomas from MR imaging. *Clinical Cancer Research* **24**, 1073–1081 (2018).
- [6] Zinn, P. O. *et al.* Radiogenomic mapping of edema/cellular invasion MRI-phenotypes in glioblastoma multiforme. *PLoS ONE* **6**, e25451 (2011).

- [7] Gutman, D. A. *et al.* MR imaging predictors of molecular profile and survival: multi-institutional study of the TCGA glioblastoma data set. *Radiology* **267**, 560–569 (2013).
- [8] Aerts, H. J. *et al.* Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nature Communications* **5**, 4006 (2014).
- [9] Gevaert, O. *et al.* Glioblastoma multiforme: Exploratory radiogenomic analysis by using quantitative image features. *Radiology* **273**, 168–174 (2014).
- [10] Jamshidi, N., Diehn, M., Bredel, M. & Kuo, M. D. Illuminating radiogenomic characteristics of glioblastoma multiforme through Integration of MR imaging, messenger RNA expression, and DNA copy number variation. *Radiology* **270**, 1–2 (2014).
- [11] Grossmann, P. *et al.* Defining the biological basis of radiomic phenotypes in lung cancer. *eLife* **6** (2017).
- [12] Hu, L. S. *et al.* Radiogenomics to characterize regional genetic heterogeneity in glioblastoma. *Neuro-Oncology* **19**, 128–137 (2016).
- [13] Kickingeder, P. *et al.* Radiogenomics of glioblastoma: Machine learning-based classification of molecular characteristics by using multiparametric and multiregional MR imaging features. *Radiology* **281**, 907–918 (2016).
- [14] Zhang, B. *et al.* Multimodal MRI features predict isocitrate dehydrogenase genotype in high-grade gliomas. *Neuro-Oncology* **19**, 109–117 (2017).
- [15] Gevaert, O. *et al.* Predictive radiogenomics modeling of EGFR mutation status in lung cancer. *Scientific Reports* **7**, 41674 (2017).
- [16] Gevaert, O. *et al.* Non-small cell lung cancer: Identifying prognostic imaging biomarkers by leveraging public gene expression microarray data-methods and preliminary results. *Radiology* **264**, 387–396 (2012).
- [17] Zhu, Y. *et al.* Deciphering genomic underpinnings of quantitative MRI-based radiomic phenotypes of invasive breast carcinoma. *Scientific Reports* **5**, 17787 (2015).
- [18] Guo, W. *et al.* Prediction of clinical phenotypes in invasive breast carcinomas from the integration of radiomics and genomics data. *Journal of Medical Imaging* **2**, 041007 (2015).
- [19] Yamashita, K. *et al.* MR imaging-based analysis of glioblastoma multiforme: Estimation of IDH1 mutation status. *American Journal of Neuroradiology* **37**, 58–65 (2016).
- [20] Lecun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
- [21] Goodfellow, I., Bengio, Y. & Courville, A. *Deep learning* (MIT Press, 2016).
- [22] Korfiatis, P. *et al.* Residual deep convolutional neural network predicts MGMT methylation status. *Journal of Digital Imaging* **30**, 622–628 (2017).
- [23] Li, S., Han, H., Sui, D., Hao, A. & Qin, H. A novel radiogenomics framework for genomic and image feature correlation using deep learning. In *2018 IEEE International Conference on Bioinformatics and Biomedicine*, 1, 1–8 (IEEE, 2018).
- [24] Ha, R. *et al.* Predicting breast cancer molecular subtype with MRI dataset utilizing convolutional neural network algorithm. *Journal of Digital Imaging* **32**, 276–282 (2019).

- [25] Chen, Y., Li, Y., Narayan, R., Subramanian, A. & Xie, X. Gene expression inference with deep learning. *Bioinformatics* **32**, 1832–1839 (2016).
- [26] Zeiler, M. D. & Fergus, R. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, 818–833 (Springer, 2014).
- [27] Simonyan, K., Vedaldi, A. & Zisserman, A. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034* (2014).
- [28] Verhaak, R. G. *et al.* Integrated genomic analysis identifies clinically relevant subtypes of glioblastoma characterized by abnormalities in PDGFRA, IDH1, EGFR, and NF1. *Cancer Cell* **17**, 98–110 (2010).
- [29] Network, C. G. A. R. Comprehensive genomic characterization defines human glioblastoma genes and core pathways. *Nature* **455**, 1061–1068 (2008).
- [30] Scarpance, L. *et al.* Radiology Data from The Cancer Genome Atlas Glioblastoma Multiforme [TCGA-GBM] collection (2016). URL <http://doi.org/10.7937/K9/TCIA.2016.RNYFUYE9>.
- [31] Group, T. G. P. R. VASARI (2015). URL <https://wiki.cancerimagingarchive.net/display/Public/VASARI+Research+Project>.
- [32] Chollet, F. & Others. Keras (2015). URL <https://github.com/fchollet/keras>.
- [33] Chen, T. & Guestrin, C. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794 (San Francisco, California, USA, 2016).
- [34] Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011).
- [35] Subramanian, A. *et al.* Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences of the United States of America* **102**, 15545–50 (2005).
- [36] Liberzon, A. *et al.* Molecular signatures database (MSigDB) 3.0. *Bioinformatics* **27**, 1739–1740 (2011).
- [37] Patel, A. P. *et al.* Single-cell RNA-seq highlights intratumoral heterogeneity in primary glioblastoma. *Science* **344**, 1396–401 (2014).
- [38] Darmanis, S. *et al.* A survey of human brain transcriptome diversity at the single cell level. *Proceedings of the National Academy of Sciences* **112**, 7285–7290 (2015).
- [39] Zhang, Y. *et al.* Purification and characterization of progenitor and mature human astrocytes reveals transcriptional and functional differences with mouse. *Neuron* **89**, 37–53 (2016).
- [40] Puchalski, R. B. *et al.* An anatomic transcriptional atlas of human glioblastoma. *Science* **360**, 660–663 (2018).
- [41] Parsons, D. W. *et al.* An integrated genomic analysis of human glioblastoma multiforme. *Science* **321**, 1807–12 (2008).

- [42] Kotikalapudi, R. & Others. keras-vis (2017). URL <https://github.com/raghakot/keras-vis>.
- [43] Sergushichev, A. A. An algorithm for fast preranked gene set enrichment analysis using cumulative statistic calculation. *bioRxiv* 060012 (2016).
- [44] Naba, A. *et al.* The matrisome: In silico definition and in vivo characterization by proteomics of normal and tumor extracellular matrices. *Molecular & Cellular Proteomics* **11**, M111.014647 (2012).
- [45] Liberzon, A. *et al.* The molecular signatures database hallmark gene set collection. *Cell Systems* **1**, 417–425 (2015).
- [46] Mongiardi, M. P. *et al.* c-MYC inhibition impairs hypoxia response in glioblastoma multiforme. *Oncotarget* **7**, 33257–71 (2016).
- [47] Bourgonje, A. M. *et al.* Intracellular and extracellular domains of protein tyrosine phosphatase PTPRZ-B differentially regulate glioma cell growth and motility. *Oncotarget* **5**, 8690–702 (2014).
- [48] Colen, R. R. *et al.* Imaging genomic mapping of an invasive MRI phenotype predicts patient outcome and metabolic dysfunction: A TCGA glioma phenotype research group project. *BMC Medical Genomics* **7**, 30 (2014).
- [49] Agarwala, R. *et al.* Database resources of the National Center for Biotechnology Information. *Nucleic Acids Research* **46**, D8–D13 (2018).
- [50] Omuro, A. & DeAngelis, L. M. Glioblastoma and other malignant gliomas. *JAMA* **310**, 1842 (2013).
- [51] Pitter, K. L. *et al.* Corticosteroids compromise survival in glioblastoma. *Brain* **139**, 1458–1471 (2016).
- [52] Ohgaki, H. & Kleihues, P. Genetic pathways to primary and secondary glioblastoma. *The American Journal of Pathology* **170**, 1445–53 (2007).
- [53] Zinn, P. O. *et al.* A coclinical radiogenomic validation study: Conserved magnetic resonance radiomic appearance of periostin-expressing glioblastoma in patients and xenograft models. *Clinical Cancer Research* **24**, 6288–6299 (2018).

Supplemental data

Dataset

Supplemental Materials contains the following data used in this study:

- `gene_expression.txt` - Affymetrix transcriptomes
- `gene_expression_ids.txt` - patients' TCGA barcodes
- `TCGA_clinical_data` - TCGA-GBM clinical data
- `TCGA_unified_CORE_ClaNC840.txt` - molecular subtype and gene sets from Verhaak et al.
- `vasari_annotations.csv` - MRI traits

Generated data and models can be found in the paper's source code on Github.

Table S1: Patient characteristics with transcriptomes. Transcriptomes were analyzed with Affymetrix HT Human Genome U133 Arrays by the Broad Institute. The quantile normalized and background corrected (GenePattern platform; Level 3) data were downloaded from the Genomic Data Commons using the legacy data portal at <https://gdc.cancer.gov/> . The radiogenomic subset is based on patients with any pre-op MRI study. The Verhaak et al. study was based on unified gene expressions from multiple platforms. Here, we took gene expression profiles measured on the same an Affymetrix platform as the radiogenomic models to create the subtype subset.

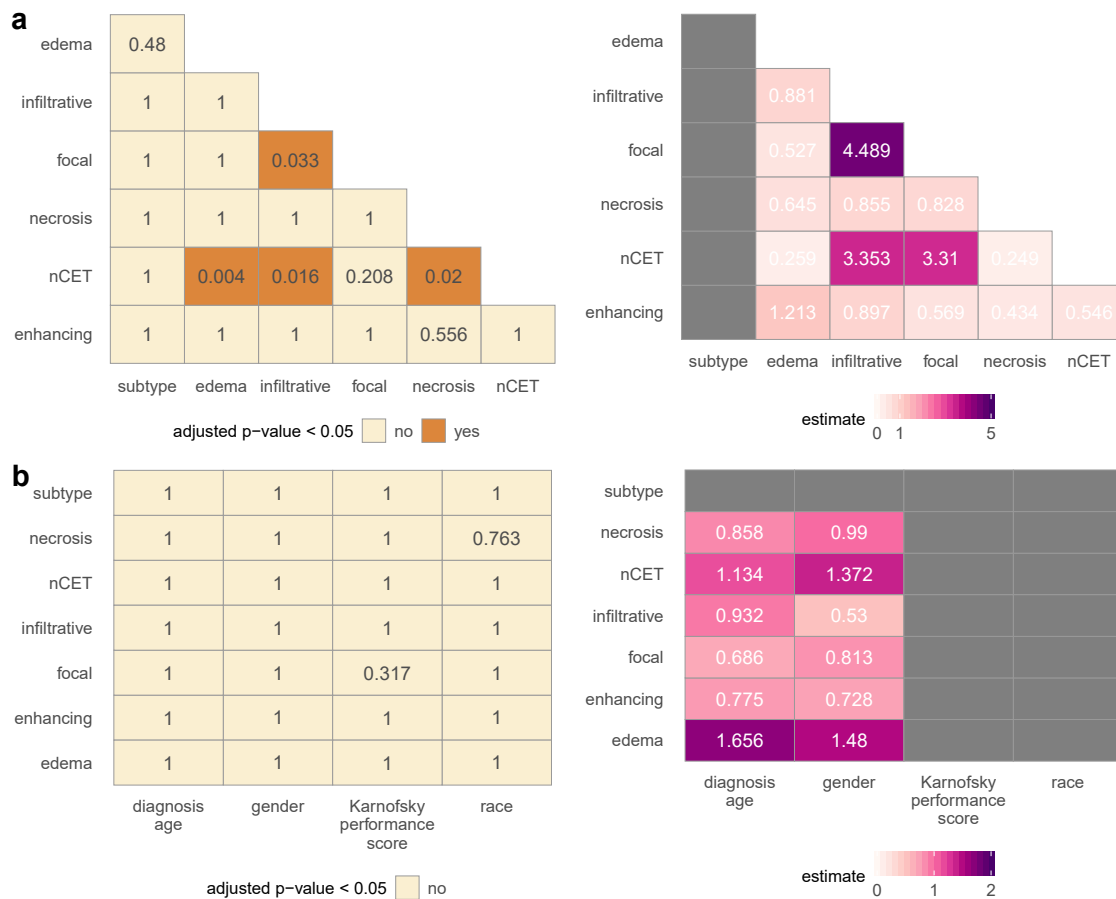
			modeling datasets			
			total	autoencoder	radiogenomic	subtype
number of patients			528	353	175	171
diagnosis age (years)	min.	10	10	14	17	
	mean	57.9	57.1	59.4	56.3	
	med.	59	58	60.5	57.5	
	max.	89	89	86	86	
vital status	deceased	388	250	138	146	
	alive	113	85	28	8	
	n/a	27	18	9	17	
gender	female	196	134	62	56	
	male	306	202	104	98	
	n/a	26	17	9	17	
race	white	443	302	141	135	
	african	29	17	12	6	
	asian	11	7	4	5	
	n/a	45	27	18	25	
ethnicity	hispanic	12	9	3	5	
	not hispanic	415	278	137	141	
	n/a	101	66	35	25	
KPS	min.	20	20	40	40	
	mean	77	76	78	82	
	med.	80	80	80	80	
	max.	100	100	100	100	
diagnosis method	biopsy	63	46	17	6	
	resection	435	288	147	147	
	other	2	1	1	0	
	n/a	28	18	10	18	
molecular subtype	classical	38	23	15	38	
	mesenchymal	54	26	28	54	
	neural	26	12	14	26	
	proneural	53	30	23	53	
	n/a	357	262	95	0	

KPS (Karnofsky Performance Score), n/a (not available)

Table S2: Classification labels in the radiogenomic and subtype datasets. Based on personal communication with TCIA, the most recent imaging studies without signs of surgery or biopsy were estimated to be images closely associated with time of tissue sampling, as pre-op imaging was necessary for surgical planning. (GBM patients usually undergo surgical resection to remove the bulk of the tumor.) Proportion labels were binarized using $\geq 1/3$ as the threshold due to small numbers in other categories. Likewise, subcategories of multifocal or multicentric (24 patients) and gliomatosis (2 patients) were combined into one class, non-focal.

n	trait	description	values	# (%)
175	surgical	evidence of prior surgery or biopsy in the earliest imaging study	pre-op post-op	175 (100%) 0 (0%)
166	f5	proportion of tumor estimated to be enhancing	$< 1/3$ $\geq 1/3$	95 (57%) 71 (43%)
156	f6	proportion of tumor estimated to be non-contrast enhancing and not edema	$< 1/3$ $\geq 1/3$	99 (63%) 57 (37%)
167	f7	proportion of tumor estimated to be necrosis	$< 1/3$ $\geq 1/3$	116 (70%) 51 (30%)
161	f9	lesions outside of main tumor and its edema: a) none b) spread via dissemination or in majority of a hemisphere	focal non-focal	135 (84%) 26 (16%)
159	f10	ratio of abnormality sizes in T1 and FLAIR: a) $T1 \approx FLAIR$ b) $T1 < FLAIR$ or $T1 \ll FLAIR$	expansive infiltrative	103 (65%) 56 (35%)
162	f14	proportion of tumor estimated to be edema	$< 1/3$ $\geq 1/3$	86 (53%) 76 (47%)
171	subtype	molecular subtypes defined by Verhaak et al.	classical mesenchymal neural proneural	38 (22%) 54 (32%) 26 (15%) 53 (31%)

Figure S1: Label association testing with Fisher's exact test (a) within labels and (b) with clinical traits in R using `fisher.test`. P-values were adjusted using Bonferroni correction in `p.adjust`. Multi-class labels do not have estimates, i.e, odds ratios. For binary labels, labels "< 1/3", "focal", and "expansive" were class 1 and "≥ 1/3", "non-focal", "infiltrative" were class 2. Rows are the explanatory variable. E.g., the odds of nCET < 1/3 and an expansive tumor were 3.353 times higher than nCET > 1/3 and an expansive tumor. Estimates above 1 indicated positive correlations, and vice versa. The traits infiltrative was positively correlated with the traits focal and nCET proportions. nCET was negatively correlated with edema and necrosis. No other associations were significant. Diagnosis age was dichotomized based on the mean diagnosis age of 58, a value calculated from all patients whose diagnosis age was known. Patients with MRI traits had tissue samples obtained from eight different sites (site codes: 02, 06, 08, 14, 19, 27, and 76). Karnofsky performance scores included 40, 60, 80, and 100; the higher the better.



Models

Radiogenomic neural networks were fully connected, had a batch size of ten. Regularization with dropout was used, where the same dropout rate was applied across the input and all hidden layers. For binary classes, 0 and 1 labels, binary cross-entropy loss and sigmoid activation in the prediction layer was used. Since subtypes were multi-classes, the subtype neural networks instead used categorical cross-entropy loss and softmax activation in the prediction layer.

Transcriptomic autoencoders considered five different architectures of three encoding layers: the first hidden layer was either 1000, 2000, 3000, or 4000, where subsequent encoding layers decreased by half each time. The autoencoder used mean absolute error as the loss function, a batch size of 50, a patience of 200 epochs while monitoring validation R^2 , and no dropout.

All neural networks used batch normalization in each hidden layer and a maximum of 500 epochs was set. Due to the inherent class imbalance of imaging traits, sample weighting based on class size and stratified fold splitting were used.

Table S3: Radiogenomic models and hyperparameters.

model type	no. models	hyperparameter	values
logistic regression	4000	penalty type C penalty solver	L1, L2 $\log(3) - \log(1)$ liblinear, newton-cg, lbfgs, sag, saga
support vector machines	4000	kernel C penalty	linear, poly, rbf, sigmoid $\log(-6) - \log(1)$
random forest	1200	trees split criterion max. features max. depth	[50: 50: 2000] Gini, entropy $G, \sqrt{G}, \log_2(G)$ None, [1-4]
gradient boosted trees	760	trees max. depth learning rate	[50: 50: 1000] [1-4] [0.01: 0.05: 0.50]
neural network, autoencoder	40	hidden layers hidden nodes architectures optimizer activation dropout loss epochs patience batch weight initializer no. layers frozen	3 [4000-250] 4 Nadam, Adadelata sigmoid, tanh, relu [0.0:0.2:0.6] binary cross-entropy, mean absolute error 200, 500 200 epochs 10, 50 Glorot normal, autoencoder 0, 1, 2

G: number of genes


Table S4: Subtype neural network hyperparameters.

model	no. models	hyperparameter	values
neural network	90	hidden layers	3
		hidden nodes	[4000-125]
		architectures	5
		optimizer	Nadam, Adadelta
		activation	sigmoid, tanh, relu
		dropout	[0.4:0.2:0.8]
		loss	categorical cross-entropy
		epochs	200
		batch	10
		weight initializer	Glorot normal

Model performance

Autoencoder

Figure S2: Performance of transcriptomic autoencoder in **(top)** 10-fold cross-validation grid search, and **(bottom)** R^2 distribution after retraining. Architecture refers to the hidden nodes in the three encoding layers.



Radiogenomic models

Table S5: Radiogenomic model performances. For each model, the best performing hyperparameters were selected based on cross-validation AUC, and their results are shown here.

	label	name	nn	gbt	rf	svm	logit
AUC	f5	enhancing	0.722	0.511	0.542	0.571	0.538
	f6	nCET	0.826	0.722	0.721	0.667	0.617
	f7	necrosis	0.751	0.638	0.627	0.607	0.553
	f14	edema	0.784	0.673	0.647	0.613	0.517
	f10	infiltrative	0.780	0.492	0.545	0.617	0.573
	f9	focal	0.849	0.728	0.697	0.694	0.650
Δ AUC (nn - another model)	f5	enhancing	-	0.211	0.180	0.151	0.184
	f6	nCET	-	0.104	0.105	0.159	0.209
	f7	necrosis	-	0.113	0.124	0.144	0.198
	f14	edema	-	0.111	0.137	0.171	0.267
	f10	infiltrative	-	0.288	0.235	0.163	0.207
	f9	focal	-	0.121	0.152	0.155	0.199

nCET (non-contrast enhancing tumor), neural network (nn), gradient boosted trees (gbt), random forest (rf), support vector machines (svm), logistic regression (logit)

Subtype neural network

Table S6: Cross-validation results of selected hyperparameters for subtype neural network. Values are in AUC and averaged over 10 folds. Individual subtype AUCs were calculated based on one-versus-others.

	classical	mesenchymal	neural	proneural	micro-averaged	macro-averaged
training	0.9965	0.9980	1	0.9974	0.9964	0.9983
validation	0.9841	0.9974	1	0.9910	0.9938	0.9956

Table S7: Performance scores of subtype model in gene masking with various gene sets, corresponds to Fig. 5 in main text. Subtype gene sets were defined by [28]. Random gene sets were obtained by random sampling of genes, excluding subtypes genes. The fully trained subtype model had perfect classification, as expected. The model was able to retain high performance scores when only using subtype genes.

gene set	gene size	AUC	f1-score	average precision
random	100	0.741	0.485	0.528
random	200	0.829	0.643	0.676
mesenchymal	216	0.897	0.725	0.789
all subtypes	840	0.994	0.930	0.984

Bootstrapped performances

Figure S3: Comparison of neural network performance compared to other models in 100 bootstrapped datasets. Points centered below the diagonal line indicate cases where neural networks had better 10-fold cross-validation performance, and was the case for all bootstrapped datasets. For each bootstrap, the difference was equal to the neural network performance minus another model's performance. Notation - gbt: gradient boosted trees, rf: random forest, svm: support vector machines, logit: logistic regression.




Figure S4: Distribution of 10-fold cross-validation performances in 100 bootstrapped datasets across models. Dashed vertical lines represent 95% confidence intervals (CI). Solid vertical lines indicate 0.5 AUC - classification was as good as random. Confidence intervals did not overlap between neural networks and any other model and suggested neural networks were better models.





Figure S5: Distribution of training and validation of neural networks in the 100 bootstrapped datasets. The 10-fold cross-validation means: (a) number of epochs (b) loss, (c) AUC, (d) average precision. Values were recorded at the epoch where highest performance was reached in early stopping. Grey intervals show the standard deviation of a value within a 10-fold cross-validation. Bootstraps were sorted by validation values for visualization purposes.



Gene masking

As gene set size increased, both AP and AUC increased. Although the full input, i.e., 12,042 gene expressions, was much larger, the models mapped multiple subsets of the input without losing a large proportion of classification performance. This was likely reflected the correlative nature between gene expressions, e.g., previous work with neural networks from Chen et al. was able to use landmark genes to predict another 9520 target genes [25].

Subtype neural network

Figure S6: Subtype gene set masking with subtype model. **(top)** Model's probabilities and, **(bottom)** model performances. Subtypes and their gene sets were taken from [28]. Random genes were randomly sampled and did not overlap with subtype genes.




Figure S7: Perturbation of the subtype model with gene sets describing cell types and phenotypes [39, 38, 37], downloaded from Puchalski et al. [40]. (**top**) Model performance in gene set masking and (**bottom**) gene set enrichment in genes ranked by single gene masking. For comparison, see the paper’s Fig. S6. Neither rows nor columns were clustered in order to match the order in the paper. Cell type enrichment for subtypes are similar to the gene masking scores for neural, proneural and mesenchymal findings in the paper. A major difference includes the subtype neural network’s ability to predict with high AUC and precision the mesenchymal subtype with the GBM core astrocyte gene set, but this was not shown in the paper. The study found neural and proneural subtypes are often enriched by astrocyte gene sets; mesenchymal associated with endothelial cells; and proneural with oligodendrocytes and quiescent fetal neurons. However, the associations between the study’s and the neural network’s gene masking do not completely agree. This disagreement was likely due to different goals: the study measures gene set associations based on single sample gene set enrichment analysis and gene masking was based on classification performance. Although, the overlap does show consistency between GBM subtypes and brain cell types, suggested that cell types were both enriched and predictive of the subtypes.





Figure S8: Perturbation of the subtype model with hallmark gene sets [45]. Model performance in gene set masking. Shown are the top 20 hallmarks for each subtype ranked by average precision, totaling 32 gene sets.



Radiogenomic neural networks

The 22 queried GBM genes from [29, 41] were *AKT3*, *CCND2*, *CDK4*, *CDK6*, *CDKN2A*, *CDKN2B*, *CDKN2C*, *EGFR*, *ERBB2*, *IDH1*, *MDM2*, *MDM4*, *MET*, *MYCN*, *NF1*, *PARK2*, *PDGFRA*, *PIK3CA*, *PIK3R1*, *PTEN*, *RB1*, *TP53*.

Figure S9: Hallmark gene set masking in radiogenomic models ranked by AUC with corresponding average precision. Shown are the top 10 genes per label.




Figure S10: Single gene masking in radiogenomic models ranked by (**top**) AUC with corresponding average precision values on the right and (**bottom**) average precision with corresponding AUC values on the right. Shown are the top 10 genes per label.




Figure S11: Gene masking using gene sets associated with GBM genomic abnormalities [29, 41]. Shown are the top 15 gene sets ranked by average precision (AP) for (a) enhancing, (b) edema, (c) nCET, (d) necrosis, (e) focal, and (f) infiltrative neural networks. Only the collections from GO, motif, and canonical pathways were considered. There were 22 genes queried, where gene sets may involve more than one of the queried genes.




Figure S12: Comparison of enhancement associations found in previous studies. Radiogenomic models were masked with gene sets queried from the Molecular Signature Database (MSigDB) based on key words from published findings. **(a)** Enhancement was found to be associated with hypoxia, ECM, and angiogenesis gene modules from Diehn et al. (n=22) [2]. There were 15 gene sets returned from querying hypoxia, ECM, and angiogenesis. **(b)** Jamshidi et al. found associations between 17 Biocarta pathways and enhancement [n=23] [10]. **(c)** Jamshidi et al. also stated *C1orf172*, *CAMSAP2*, *KCNK3*, *LTBP1* genes were related to enhancement. There were 114 gene sets containing the four genes **(d)** Gutman et al. found a non-significant correlation between *EGFR* copy number amplification and enhancement (n=75) [7]. Querying MSigDB for *EGFR* resulted in 257 gene sets. The top ten gene sets ranked by the average precision (AP) in predicting enhancement were kept for (b-c).




Figure S13: Comparison of edema associations found in Zinn et al. [n=78] [6]. Radiogenomic models were masked with gene sets queried from the MSigDB based on key words from published findings. Zinn et al. found concordant changes in mRNA and microRNA fold changes for patients with high edema or invasive traits: **(a)** The top five upregulated microRNA and their associated genes and **(b)** top five upregulated genes in patients with high FLAIR volumes, see Table 4 of the authors' paper. This resulted in 343 and 160 related MSigDB gene sets, respectively. Shown are the top ten gene sets ranked by average precision (AP) in predicting edema in our study.

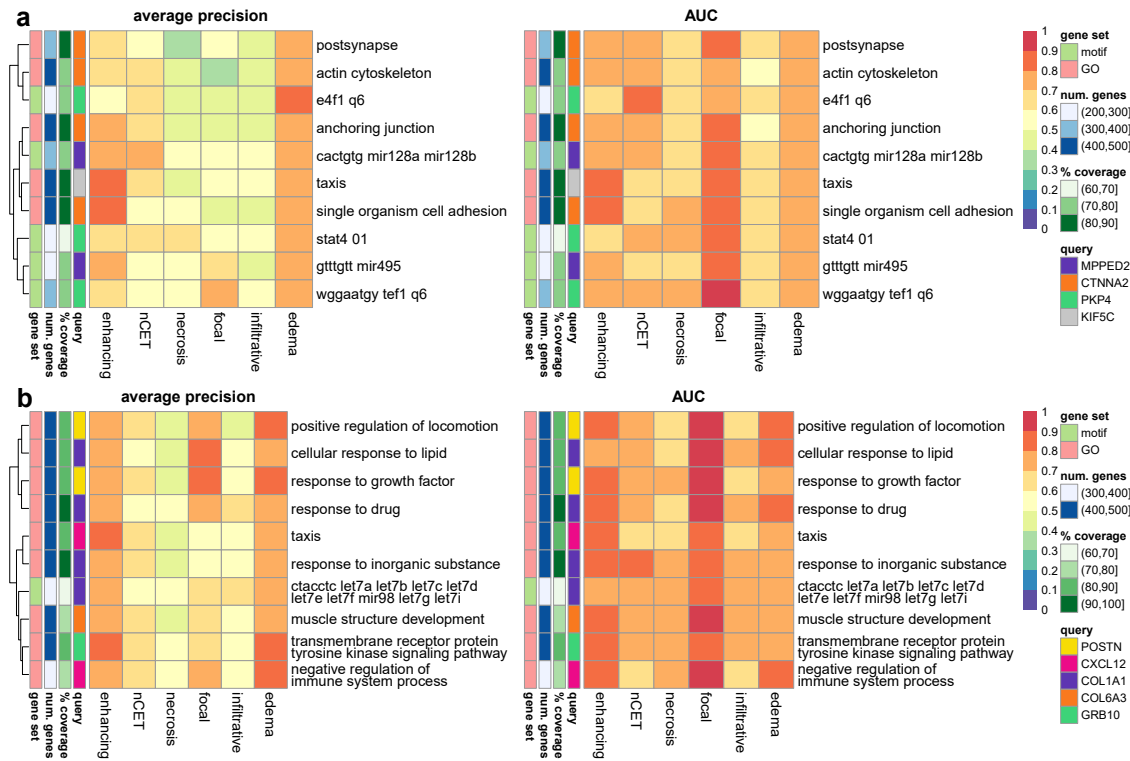


Figure S14: Comparison of necrosis associations found in previous studies. Radiogenomic models were masked with gene sets queried from the MSigDB based on key words from published findings. (a) Gevaert et al. reported the boundary sharpness of necrosis region was associated with *GAP43*, *WWTR1*, IL4 pathway, and cell membrane genes [n=55][9]. This returned 378 gene sets from MSigDB. (b) Jamshidi et al. found four Biocarta pathways (named gene sets) and (c) two genes associated with necrosis [10]. (d) Gutman et al. found necrosis was correlated with the deletion of *CDKN2A*, but was not significant [7]. An MSigDB query for the gene returned 183 gene sets. The top ten gene sets ranked by the average precision in predicting necrosis were kept in (a,b,d).




Figure S15: Comparison of expansiveness versus infiltrative associations found in previous studies. Radiogenomic models were masked with gene sets queried from the MSigDB based on key words from published findings. **(a)** Colen et al. found invasive tumors to be associated with *MYC*, leading to *NFKBIA* inhibition [n=104] [48]. Querying MSigDB for *MYC* and *NFKBIA* resulted in 316 gene sets. **(b)** Diehn et al. looked at infiltrative vs. edematous T2 abnormality and found an association with an immune cell gene module [2]. An MSigDB query for ‘immune’ returned 87 gene sets. Shown are the top ten gene sets ranked by average precision (AP) in predicting infiltrative tumors in our study.

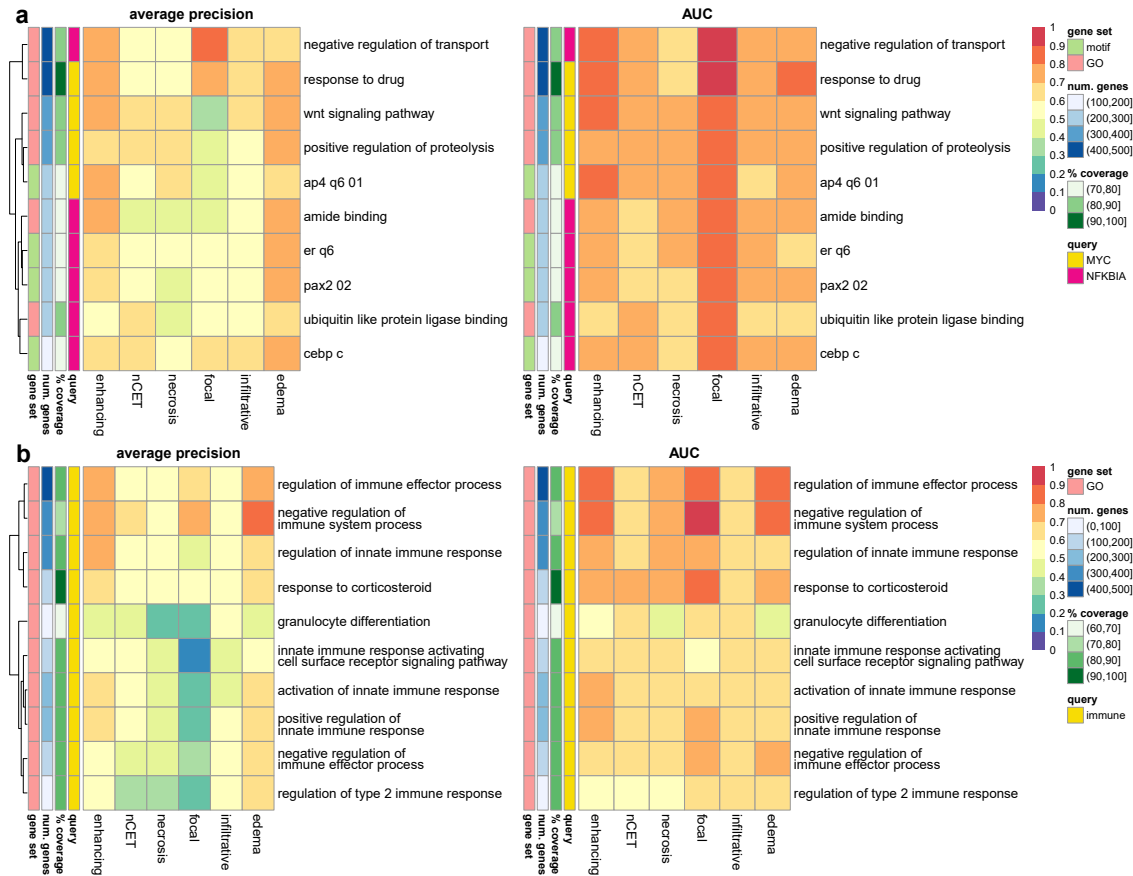


Figure S16: Perturbation of radiogenomic models with subtype gene sets from [28]. **(top)** Model performance in gene set masking, and **(bottom)** gene set enrichment in ranked genes after single gene masking. The authors found mesenchymal tumors had less nCET and proneural tumors had less enhancement. However, the radiogenomic models found neural genes were more predictive of nCET than mesenchymal genes (0.64 vs. 0.55 AUC, 0.52 vs. 0.46 AP) and other subtype genes were more predictive than proneural genes in predicting enhancement.




Figure S17: Perturbation of radiogenomic models with gene sets from brain cell types and phenotypes [39, 38, 37], downloaded from Puchalski et al. [40]. **(top)** Gene set masking, where the top 20 ranked by average precision were kept; shown are 21 gene sets. **(bottom)** GSEA of genes ranked in single gene masking. Shown are gene sets with at least one significant enrichment.




Figure S18: Gene masking in radiogenomic models with canonical gene sets from MSigDB. **(top)** Gene set masking, where the top 5 ranked by average precision were kept; shown are 21 gene sets. **(bottom)** GSEA of genes ranked in single gene masking. Shown are gene sets with at least one significant enrichment.




Figure S19: Gene masking in radiogenomic models with motif gene sets from MSigDB. (top) Gene set masking, where the top 5 ranked by average precision were kept; shown are 29 gene sets. **(bottom)** GSEA of genes ranked in single gene masking. Shown are gene sets with at least one significant enrichment.




Figure S20: Gene masking in radiogenomic models with chromosome gene sets from MSigDB. **(top)** Gene set masking, where the top 5 ranked by average precision were kept; shown are 25 gene sets. **(bottom)** GSEA of genes ranked in single gene masking. Shown are gene sets with at least one significant enrichment.




Figure S21: Gene masking in radiogenomic models with oncogenic signatures gene sets from MSigDB. **(top)** Gene set masking, where the top 5 ranked by average precision were kept; shown are 25 gene sets. **(bottom)** GSEA of genes ranked in single gene masking. Shown are gene sets with at least one significant enrichment.




Figure S22: Gene masking performance compared to **(top)** gene set size and **(bottom)** coverage. Coverage was the percent of genes in the gene set that was in input gene expressions. Average precision (left column) and AUC (right column) were used.




Figure S23: The trend between gene masking performance and gene set size broken down by gene set category: **(top)** average precision and **(bottom)** AUC.






Figure S24: The trend between gene masking performance and gene set coverage broken down by gene set category: **(top)** average precision and **(bottom)** AUC.



Gene saliency

Subtype neural network

Figure S25: Gene saliency in the subtype neural network. **(a)** Clustering of the enrichment scores between each patients (columns) and each subtype gene set (rows). Enrichments were significant at an adjusted p-value < 0.05 . **(b)** The number of patients with significant enrichment in each subtype gene set based on the bottom heatmap in (a) and broken down by their true subtypes. Of the patients with the mesenchymal subtype, more patients (33 patients) were enriched with mesenchymal genes than any other subtype genes. To a lesser degree, the same was seen among patients with the classical subtype, where 32 patients were enriched with classical genes. In contrast, 49 patients in the proneural subtype were enriched with classical genes while 48 were enriched with proneural genes. These aforementioned associations between gene saliency and subtype overlap with the associations from single gene masking in Fig. 6b. As was seen in gene masking, genes from other subtypes were influential in predicting any individual subtype. In particular, neural patients' salient genes were enriched more often with mesenchymal genes (22 patients) than with neural genes (14 patients). Likewise, proneural patients were similarly enriched with both classical and proneural genes.



Radiogenomic neural networks

Figure S26: Patients with significant enrichment radiogenomic models based on gene saliency. Shown are other MSigDB gene set collections with at least 10 enriched patients. The chemical and genetic perturbations and computational collections were thresholded at 20 patients; all others were thresholded at 10 patients.



Survival

Ethnicity was removed since 98% of patients with radiogenomic data were non-Hispanic.

Overall survival (OS) was defined in 'patient' file. For patients without death events, days-to-last-followup were obtained by finding the maximum day among any reported days-to-event in the 'patient', 'radiation', 'drug', and 'followup' files.

Progression-free survival (PFS) outcomes were defined by the 'followup' file, which included the event types: locoregional disease, metastatic, progression of disease, or recurrence. Unknown event types with days-to-event data were also removed unless other files stated patients had therapy regimes treated for progression, e.g., an unknown progression event type on day 554 and received radiation on day 576 to treat progression. For these cases (n=3), days-to-event data was kept, while event types were set to progression.

Figure S27: Survival estimates between TCGA-GBM (patients with transcriptome data), VASARI subset (patients with at least one MRI trait), and radiogenomic subset (patients with all MRI traits) in (a) overall survival (OS) and (b) progression-free survival (PFS). There were no differences in OS or PFS, see Supp. Table S8. Note: 27 of the overall cohort of 528 patients had missing survival and/or progression data, see Supp. Table S1. Similarly, 166 of 175 patient with any MRI trait also had survival information. Of the 166 patients, only 127 had all six MRI traits.




Table S8: Summary measures of survival curves corresponding to cohorts in Supp. Figure S27.

cohort	n	events	median (years)	95% CI (years)
Overall survival				
TCGA-GBM	501	388	1.16	1.05, 1.24
VASARI subset	166	138	1.16	0.94, 1.26
radiogenomic	127	107	1.05	0.90, 1.27
Progression-free survival				
TCGA-GBM	502	331	0.70	0.64, 0.77
VASARI subset	166	115	0.66	0.52, 0.85
radiogenomic	127	88	0.54	0.46, 0.73

Figure S28: Survival estimates between patients split by their MRI traits. (a) Overall survival (OS) and (b) progression-free survival (PFS) differences.




Figure S29: Frequency of enriched patients in each of the 54 radiogenomic traits. n = number of enriched patients

