Deep Learning and Transfer Learning for Brain Tumor Detection and Classification

Faris Rustom[1], Pedram Parva[2,3] , Haluk Ogmen[4], Arash Yazdanbakhsh[5*]

1. *Computational Neuroscience and Vision Lab, Neuroscience Program, Boston University*
2. *Department of Radiology, Veterans Affairs Boston Health System, Boston*
3. *Boston University School of Medicine, Boston*
4. *Department of Electrical & Computer Engineering, Laboratory of Perceptual & Cognitive Dynamics, University of Denver, Denver, CO, USA*
5. *Department of Psychological and Brain Sciences, Computational Neuroscience and Vision Lab, Center for Systems Neuroscience, Boston University, Boston, MA, USA*
   **Corresponding author*

Original Research

Summary Statement

Convolutional neural networks can be trained with transfer learning to accurately detect and classify brain tumor MRIs while simultaneously demonstrating an approach to imaging data resembling that of trained radiologists.

Key Results

- Transfer learning from camouflage animal recognition neural networks, and between T1-T2 MRI sequence, improved brain tumor detection and classification accuracy in retrospective study.
- Trained neural networks correctly identified oligoastrocytomas as an intermediate of astrocytomas and oligoastrocytomas, as visualized by feature space mapping.
- Image sensitivity maps of networks suggested a high sensitivity to midline shifts, tissue compressions, and symmetrical comparisons of brain regions for accurate tumor classification.

List of Abbreviations

CNN: Convolutional Neural Network
FC: Fully Connected
MATLAB: Matrix **Lab**oratory
MRI: Magnetic Resonance Imaging
ROC: Receiver Operating Characteristic
T1: MRI T1 Weighted Image, which enhances the signal of the fatty tissue and suppresses the signal of the water
T2: MRI T2 Weighted Image, which enhances the signal of the water
TCIA: The Cancer Imaging Archive of NIH

Abstract

Convolutional neural networks (CNNs) are powerful tools that can be trained on image classification tasks and share many structural and functional similarities with biological visual systems and mechanisms of learning. In addition to serving as a model of biological systems, CNNs possess the convenient feature of transfer learning where a network trained on one task may be repurposed for training on another, potentially unrelated, task. In this retrospective study of public domain MRI data, we investigate the ability of neural networks to be trained on brain cancer imaging data while introducing a unique camouflage animal detection transfer learning step as a means of enhancing the network's tumor detection ability. Training on glioma, meningioma, and healthy brain MRI data, both T1- and T2-weighted, we demonstrate the potential success of this training strategy for improving neural network classification accuracy both quantitatively with accuracy metrics and qualitatively with feature space analysis of the internal states of trained networks. In addition to animal transfer learning, similar improvements were noted as a result of transfer learning between MRI sequences, specifically from T1 to T2 data. Image sensitivity functions further this investigation by allowing us to visualize the most salient image regions from a network's perspective while learning. Such methods demonstrate that the networks not only 'look' at the tumor itself when deciding, but also at the impact on the surrounding tissue in terms of compressions and midline shifts. These results suggest an approach to brain tumor MRIs that is comparatively similar to that of trained radiologists while also exhibiting a high sensitivity to subtle structural changes resulting from the presence of a tumor.

Introduction

Convolutional Neural Networks and Transfer Learning

Convolutional neural networks (CNNs) are powerful tools trained on image recognition tasks and perform at a high level of accuracy. CNNs learn to extract salient features that are unique to trained object classes and associate these features with the specific category labels [1]. They also have the capacity for "transfer learning", where a model trained on one task can be repurposed for a new one [2-4]. Image-based tumor classification involves a difficult pattern recognition problem and can therefore benefit from a transfer learning approach.

Figure 1 – Data Flowchart

Gliomas and Meningiomas

Gliomas and meningiomas are among the most common types of brain tumors, accounting for two-thirds of cases [5]. Gliomas can be further divided into astrocytomas and oligodendrogliomas. Oligoastrocytomas are no longer recognized as a glioma class as they are considered mixed cell tumors, however, as part of our study it is used to evaluate CNN performance [6]. Gliomas and meningiomas can be distinguished from one another through MRI data alone in the majority of cases [7, 8], while glioma subtypes are virtually indistinguishable from MRI data alone [9].

Figure 2 - Glioma vs Meningioma Examples

Transfer Learning from Camouflage Animal Detection Task

Although camouflage animal detection and brain tumor classification tasks involve different images, there might be a parallel between an animal hiding through natural camouflage and a bundle of cancerous cells blending in with the surrounding healthy tissue. The learned process of *generalization* - the grouping of different appearances under the same object identity - is essential to a network's ability to detect camouflaged objects and is why such training could prove advantageous for a tumor detection task [10]. A trained CNN's ability to generalize would be inherited through transfer learning from camouflage animal detection to improve its performance on tumor detection and classification [11].

Figure 3 - Sample Camouflage Animal Images

Feature Spaces as a Metric for Neural Networks' Internal Representation

Testing accuracy does not describe the changes in internal representation of data by the network in response to training, nor does it directly demonstrate a network's generalization ability. The activity of thousands of neurons (referred to as *nodes)* within a CNN can be extracted after training to generate a *feature space* showing the internal representation of stimuli by the network [12]. Feature spaces present the network's distribution of every test image relative to one another, providing valuable insight into the network's organization of data following training. The relationship between individual data points in such a space can be analyzed through the *Universal Law of Generalization*, which states that the perceived likeliness of two images increases as the distance between them in a feature space decreases [13].

Methods

Building Datasets for Neural Network Training

CNNs were trained on different brain tumor classification tasks using MRIs obtained from public online repositories. The Cancer Imaging Archive (TCIA) [14] provided low-grade glioma subtypes, T1- and T2-weighted (serving as different inputs to the network). Manual pre-processing of the dataset resulted in 1,680 axial glioma MRIs (247 astrocytomas, 708 oligodendrogliomas, and 725 mixed tumors formerly described as oligoastrocytomas). This was later augmented with healthy brain MRIs from the Human Connectome Project database [15], which provided 186 healthy axial brain MRIs of different age groups (45-55, 65-75). Meningioma MRIs were obtained from a dataset published on Kaggle [16, 17], which provided 208 axial meningioma MRIs and 493 additional glioma MRIs (these were used instead for *TumorNet* to minimize formatting differences between datasets).

The camouflage animal datasets used here were imported from a previous project using neural networks for the detection of camouflage animals [11]. It consisted of almost 3,000 images of clear and camouflaged animals divided into 15 categories.

Table 1 – Subject Demographics

Neural Network Training Parameters

AlexNet is a convolutional neural network consisting of a 25-layer series architecture, imported via MATLAB version 2019b (https://www.mathworks.com/products/matlab.html). AlexNet is pretrained on over a million images in up to 1000 object categories, and was used as the basis of the trained networks described below. Images were randomly split into training (60%) and testing (40%) sets, and resized to

227x227 pixels. All training was performed with the same parameters on the same computer. Modeling and data analysis code can be accessed here: https://github.com/frustom/Brain-Tumor-Classification (ID: d8cf63a)

Training Neural Networks on Datasets

GliomaNet was the first neural network trained in this project on the preliminary glioma dataset to differentiate glioma subtypes, classifying MRIs as astrocytomas, oligoastrocytomas, or oligodendrogliomas. Due to the outdated oligoastrocytoma label [6], two versions of GliomaNet were separately trained to include or exclude oligoastrocytoma data, *GliomaNet_V1* and *GliomaNet_V2*. *TumorNet* was trained on the general brain tumor classification task of labeling MRIs as either gliomas, meningiomas, or healthy.

The previously mentioned camouflage animal detection networks were used as seeds for the tumor classification networks *ExpCamoGliomaNet* and *ExpCamoTumorNet*, which both used the ExpCamoNet layer weights as a starting point for training on glioma subtypes and general tumor classification, respectively. The T1 and T2 MRI sequences present in the dataset described above were also used as the basis for additional transfer learning for one another.

Dimensionality Reduction

To map the feature spaces (examples in Figs 4&5), the vector of activation values of each category node was extracted from the Fully Connected (FC) layer of a CNN. This vector contains the activation values of every testing image across all category nodes. Principal Component Analysis (PCA) was used as a dimensionality-reduction method to simplify the larger activation matrices to three key dimensions that are representative of the overall activation values. The first three Principal Components were then plotted to model the networks' feature space in three dimensions. In these feature spaces, each point represents the activation of a single testing image on the trained FC layer nodes and its position within the feature space, as part of a color-coded *region*, is a visualization of the networks' internal representation of the data.

Image Sensitivity Maps

The *occlusionSensitivity* MATLAB function provides a heat map of a given image for its most salient features on which the trained network bases its decision, with a color bar showing the scale of sensitivity across the image. This function iteratively occludes small areas of the image input and calculates the change in probability score for a given class as the occluding object moves across the image. The *occlusionSensitivity* function was applied to randomly selected test images of each category for both cancerous and healthy subjects. Glioma subtype sensitivity maps were generated with GliomaNet_V2, meningioma and healthy brain sensitivity maps were generated with TumorNet.

Receiver Operating Characteristic (ROC) Curves

To generate the ROC curve of a given network, we retrieve the activity level of Healthy (H) and Cancer (C) category nodes and normalize them by the maximum and minimum levels of the H and C nodes, setting H_{norm} and C_{norm} within a 0 - 1 range. For 100 test images equally divided into Healthy and Cancer classes, we pass them to our trained CNN and measure H_{norm} and C_{norm} for each image. With both category-node activities of H and C normalized, we choose a threshold α for both H_{norm} and C_{norm},

representing the 'criterion level' for the decision. If a cancer testing image generates a normalized activity level greater than the set threshold when passed to the network, $C_{norm} > \alpha$, then it will be registered as a hit. On the other hand, if a healthy image generates an activity level in the cancer node $C_{norm} > \alpha$, then it will be registered as a false alarm. By continuously changing the criterion threshold α for a decision from 0 to 1 we can generate the full ROC curve.

Results

GliomaNet Performance and Feature Space

The first task the convolutional neural networks were trained on was the classification of glioma subtypes from axial brain MRIs. GliomaNet_V1 was trained on the two most up-to-date glioma subtypes, astrocytomas and oligodendrogliomas, to assess the network's performance on a clinically relevant task to establish a metric with which to potentially compare with human radiologists' performance. GliomaNet_V2 was additionally trained on the now outdated oligoastrocytoma category (Fig 4) to demonstrate if the network's classification of oligoastrocytomas is consistent with the current medical literature. Both GliomaNet versions showed an immediate strong performance on their respective tasks, and after several trials in which the learning parameters were fine tuned to maximize performance, the highest accuracy achieved by V1 was 97.22% closely followed by V2's 94.26% (Table 2). Looking at the GliomaNet feature spaces, both versions showed distinct separation of the glioma types across the first two dimensions. With both V1 and V2, the astrocytoma and oligodendroglioma category clusters each inhabited a defined region with minimal overlap between them (Fig 4b/c). In the case of V2, the additional oligoastrocytoma category occupied the space between the astrocytoma and oligodendroglioma regions (Fig 4b).

Table 2 - Neural Network Names and Accuracies

Figure 4 - GliomaNet_V2/TumorNet Confusion Matrix and FS

TumorNet Performance and Feature Space

In an effort to build a more comprehensive tumor detection and classification network that could potentially serve a beneficial clinical purpose, both healthy brain and meningioma patient MRIs were added to the glioma dataset. TumorNet was trained on three categories: gliomas (now represented as a single category), meningiomas, and healthy brains. The network's initial performance, while not as strong as GliomaNet's on the glioma classification task, demonstrated a good foundation for this general tumor classification task with an accuracy of 86.14% (Table 2). The network achieved a perfect accuracy on the healthy brain MRIs, indicating an ability to correctly recognize whether or not an MRI brain section contained a tumor (Fig 4d).

Examining the feature space of TumorNet revealed the internal representation of the tumor and healthy brain MRIs by the network across the first three principal components (Fig 4e). Healthy brain MRIs were contained in a well-defined, contained region of the feature space with ample separation from the two tumor regions and no overlap with either of them, consistent with the network's perfect performance on the healthy brain testing images. The glioma and meningioma regions showed some mixing across early principal components. These regions both contained segments extending from the overlapping area in which the generalization of the tumor type was more effective.

Figure 5 – Before/After transfer learning feature spaces

Camouflage Animal Transfer Learning for Tumor Detection

The previously trained camouflage animal detection CNN, ExpCamoNet [17], was used as the seed network for transfer learning prior to glioma MRI training, resulting in ExpCamoGliomaNet_V1 which achieved a higher accuracy than its non-transfer learned counterpart GliomaNet_V1. This was also the case with ExpCamoGliomaNet_V2 (which includes the oligoastrocytoma category) compared to GliomaNet_V2 (Table 2). In both versions of the transfer-learned glioma networks' feature spaces the glioma regions expanded in response to transfer learning (Fig 5a/d, 5b/e).

The camouflage animal transfer learning step was additionally implemented to train ExpCamoTumorNet. As before, this training paradigm increased the tumor classification accuracy relative to TumorNet (Table 2). Similar transformations in the feature space of ExpCamoTumorNet (see Methods) were observed following transfer learning. As before, the healthy brain MRI region was well-separated from the tumor regions in a manner consistent with the network's perfect accuracy on this subset of images. Again, the category regions appear to have expanded after transfer learning from the camouflage animal networks (Fig 5c/f).

Transfer Learning MRI Sequences

Following the accuracy increase observed with camouflage animal network transfer learning, a separate transfer learning training pathway was designed using the two MRI sequences (T1 & T2) present in the glioma dataset. Previously, GliomaNet was indiscriminately trained on a mixture of T1 and T2 MRI training data, with the testing same composed of a similar mixture. With the introduction of the transfer learning step, T1_2Net was first trained exclusively on T1 sequence MRI data followed by a secondary independent training on T2 sequence MRIs. Again, the use of transfer learning from T1 to T2 data while training the neural network increased the accuracy (94.48%) compared to T1Net (90.14%) or T2Net (92.38%) on the glioma classification task. However, when the order of the transfer learning was reversed in T2_1Net - that is, training on T2 data before T1 - the accuracy suffered (85.92%) relative to the baseline of either T1Net or T2Net (Table 2).

ROC Curves for Trained Neural Networks

The confusion matrix of TumorNet performance in Fig 4d shows the perfect accuracy of the network when predicting healthy brains. This result indicates that there are no false alarms for this category, and for all healthy brain test images when passed to the network the activity level generated at the cancer node is less than the set decision threshold (i.e., $C_{norm} < \alpha$). Therefore, the ROC curve is trivial similar to the one shown in Fig 4f (red line). To illustrate a non-trivial ROC curve emerging from the trained networks, we considered ExpCamoTumorNet. Fig 4f shows the Glioma hit rate against the Glioma false alarms for the Meningioma images (blue line) and the Meningioma hit rate against Meningioma False alarms for the Glioma images (green line). The ROC curve of the Meningioma hit rate demonstrates that the trained CNN correctly classifies Meningiomas at a higher hit rate and lower false alarm rate than the Glioma counterparts.

Figure 6 – Image sensitivity maps and ROC curves

Discussion

GliomaNet Accurately Distinguishes Glioma Subtypes

The inclusion of oligoastrocytomas in GliomaNet_V2's dataset contributed to its lower accuracy as a source of confusion for the network (Fig 4a), compared to V1. The high accuracy with which these networks distinguished astrocytomas and oligodendrogliomas was reflected in their feature spaces (Fig 4b/c), where it translated to well-generalized category regions with distinct internal representations.

In GliomaNet_V2's feature space, the oligoastrocytoma category cluster occupied the space in between the astrocytoma and oligodendroglioma regions (Fig 4b). This is also a reflection of the mixed cell nature of oligoastrocytomas, represented in GliomaNet_V2's internal representation of the data. This internal representation of the network is consistent with the updated classification of gliomas described above.

TumorNet's Similar Performance with Perfect Accuracy on Healthy MRIs

The introduction of meningioma and healthy image data, coupled with grouping gliomas into one category, resulted in a lower accuracy. TumorNet's confusion matrix suggests its mistakes were the result of confusion between certain instances of gliomas and meningiomas, indicating that the network struggled with the intra- versus extra-axial positional differences of these tumors (Fig 4d). TumorNet correctly identified every instance of the healthy images, and its successful generalization of healthy images reflected in the feature space as a well-defined and separated category cluster (Fig 4c). It is worth noting that the healthy image dataset was smaller than other categories due to the lack of data available, contributing to the 'perfect' accuracy on this particular category's classification.

Transfer Learning from Camouflage Animals Improved Tumor Detection

ExpCamoGliomaNet and ExpCamoTumorNet showed an increased classification accuracy following transfer learning. This training paradigm appears to have successfully transferred an enhanced generalization ability that was abstracted from camouflage animal detection training to the classification of brain tumors. The increase in accuracy caused by transfer learning here, while modest, is important as it is the only training modification that increased the accuracy beyond the previous maximum.

The changes in internal representation after transfer learning observed also indicate an enhanced generalization ability not conveyed by the accuracy metric. This is better demonstrated by ExpCamoGliomaNet_V2 (Fig 5e) as the presence of the additional oligoastrocytoma category cluster fleshes out the remaining regions. After transfer learning, the glioma regions expanded and the separation between regions increased, minimizing the overlap between adjacent regions previously noted with GliomaNet. The increased size of each glioma region is indicative of improved generalization by the network, as it represents a larger surface area in which image activations can still belong to a category's domain. Such changes are also observed with ExpCamoGliomaNet_V1, although less pronounced (Fig 5d).

With ExpCamoTumorNet, the glioma and meningioma regions both expanded following transfer learning and migrated away from the previously overlapping region observed with TumorNet, indicating better generalization by the network. The adjacent tumor regions did not completely separate after transfer learning, reflecting the lingering confusion of the network on certain glioma and meningioma images (Fig 5f).

Transfer Learning Between MRI Sequences Offered Mixed Results

T1_2Net performed better than its seed network T1Net, while T2_1Net's accuracy suffered relative to the baseline of T2Net. While the reason behind the decreased accuracy in this particular order (where the reverse resulted in an increase) isn't clear, this suggests that transfer learning isn't a commutative process. The networks were shown the same sets of data in opposite orders with conflicting results, implying that the order in which a network is trained on data affects how well the information is learned.

Trained Neural Networks Study MRIs Like Radiologists

Consistently across glioma and meningioma MRIs, the network did not highlight the tumor as the most salient feature. Rather, it focused on the surrounding tissue directly adjacent to the tumor, and even in more distal regions in some cases. This is further evidenced by the color bar values associated with each image, which measure the contribution of features to the score assigned for each category label. This approach of examining the distortion of neighboring tissue caused by the growing tumor is similar to that of radiologists when studying similar cancerous MRIs [18], which the network determined without prior suggestion. Specifically, the perceived saliency to brain regions directly opposite the tumor itself in some cases (Fig 6b) suggests an ability of the network to compare symmetrical regions for structural differences secondary to tumors, representing a highly sensitive feature of the network. Through repetitive training on these images, the network learned to look at the areas affected by the tumor more than the tumor itself to differentiate tumors (Fig 6).

In the case of healthy brain MRIs, the network nonsensically highlights specific points as being highly salient. The color bar reveals that the degree of sensitivity of these points is much lower than the sensitivity scale associated with tumor-containing MRIs, implying that no specific feature within a healthy MRI registers as particularly salient (Fig 6e). Since the network can't rely on the tissue distortions associated with tumors previously demonstrated, it bases its classification of healthy MRIs on the lack of salient features within the image.

Acknowledgments: The authors would like to thank the Undergraduate Research Opportunity Program (UROP) in Boston University (BU) for three semesters of support, the members of Computational Neuroscience and Vision Lab at BU for the helpful discussion and comments during completion of the manuscript.

References

[1] S. Albawi, T. A. Mohammed and S. Al-Zawi, Understanding of a convolutional neural network, 2017 International Conference on Engineering and Technology (ICET), 2017, pp. 1-6, doi: 10.1109/ICEngTechnol.2017.8308186.

[2] R. Ribani and M. Marengoni, A Survey of Transfer Learning for Convolutional Neural Networks, 2019 32nd SIBGRAPI Conference on Graphics, Patterns and Images Tutorials (SIBGRAPI-T), 2019, pp. 47-57, doi: 10.1109/SIBGRAPI-T.2019.00010.

[3] S. J. Pan and Q. Yang, A Survey on Transfer Learning, in IEEE Transactions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345-1359, Oct. 2010, doi: 10.1109/TKDE.2009.191.

[4] Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., & Liu, C. (2018). A Survey on Deep Transfer Learning. ArXiv, abs/1808.01974.

[5] McKinney PABrain tumours: incidence, survival, and aetiology. Journal of Neurology, Neurosurgery & Psychiatry 2004;75:ii12-ii17.

[6] van den Bent, M. J., Weller, M., Wen, P. Y., Kros, J. M., Aldape, K., & Chang, S. (2017). A clinical perspective on the 2016 WHO brain tumor classification and routine molecular diagnostics. Neuro-oncology, 19(5), 614–624. https://doi.org/10.1093/neuonc/now277

[7] Piper, R. J., Mikhael, S., Wardlaw, J. M., Laidlaw, D. H., Whittle, I. R., & Bastin, M. E. (2016). Imaging signatures of meningioma and low-grade glioma: a diffusion tensor, magnetization transfer and quantitative longitudinal relaxation time MRI study. Magnetic Resonance Imaging, 34(4), 596–602. https://doi.org/10.1016/J.MRI.2015.12.006

[8] Kondziolka, D., Lunsford, L. D., & Martinez, A. J. (1993). Unreliability of contemporary neurodiagnostic imaging in evaluating suspected adult supratentorial (low-grade) astrocytoma, Journal of Neurosurgery, 79(4), 533-536. doi:https://doi.org/10.3171/jns.1993.79.4.0533

[9] Djuno I van Lent, Kirsten M van Baarsen, Tom J Snijders, Pierre A J T Robe, Radiological differences between subtypes of WHO 2016 grade II–III gliomas: a systematic review and meta-analysis, Neuro-Oncology Advances, Volume 2, Issue 1, January-December 2020, vdaa044, https://doi.org/10.1093/noajnl/vdaa044

[10] Hughes J.W. (2011) Stimulus Generalization. In: Kreutzer J.S., DeLuca J., Caplan B. (eds) Encyclopedia of Clinical Neuropsychology. Springer, New York, NY. https://doi.org/10.1007/978-0-387-79948-3_2118

[11] Rustom FB, Ogmen H, Yazdanbakhsh A. Object Detection, Recognition, Deep Learning, and the Universal Law of Generalization. IEEE PAMI (under review, pre-print at arXiv). 2022 June 10. Available from: https://doi.org/10.48550/arXiv.2206.05365 DOI: https://doi.org/10.48550/arXiv.2206.05365

[12] T. S. Lee, "The Visual System's Internal Model of the World," in Proceedings of the IEEE, vol. 103, no. 8, pp. 1359-1378, Aug. 2015, doi: 10.1109/JPROC.2015.2434601.

[13] Shepard, R. N. Toward a Universal Law of Generalization for Psychological Science. Science (80-.). 237, 1317–1323 (1987).

[14] Zeynettin Akkus, Issa Ali, Jiří Sedlář, Jay P. Agrawal, Ian F. Parney, Caterina Giannini, and Bradley J. Erickson (2017). Predicting Deletion of Chromosomal Arms 1p/19q in Low-Grade Gliomas from MR Images Using Machine Intelligence. J Digit Imaging 30(4): 469–476. DOI: https://doi.org/10.1007/s10278-017-9984-3.

[15] David C. Van Essen, Stephen M. Smith, Deanna M. Barch, Timothy E.J. Behrens, Essa Yacoub, Kamil Ugurbil, for the WU-Minn HCP Consortium. (2013). The WU-Minn Human Connectome Project: An overview. NeuroImage 80(2013):62-79.

[16] Cheng, Jun, et al. "Enhanced Performance of Brain Tumor Classification via Tumor Region Augmentation and Partition." PloS one 10.10 (2015).

[17] Cheng, Jun, et al. "Retrieval of Brain Tumors by Adaptive Spatial Pooling and Fisher Vector Representation." PloS one 11.6 (2016).

[18] Grossman, Robert I., and David M. Yousem. Neuroradiology: The Requisites. Mosby (2003)

Dataset	Image Type	Age	Sex
TCIA_LGG-1p19q	OA, Astrocytoma, OD	N/A	N/A
WU-Minn HCP LS	Healthy		
Patient ID: 5007	Healthy	45-55	F
5038	Healthy	45-55	M
5040	Healthy	45-55	M
5041	Healthy	45-55	F
5049	Healthy	45-55	M
6003	Healthy	65-75	M
6006	Healthy	65-75	M
6009	Healthy	65-75	F
6038	Healthy	65-75	F
6046	Healthy	65-75	M
Kaggle	Meningioma, Glioma	N/A	N/A

Table 1. Demographic information of subjects used from public online image databases. TCIA_LGG-1p19q: The Cancer Imaging Low Grade Gliomas 1p19q Deletion dataset, demographic information not available. OA: oligoastrocytoma. OD: oligodendroglioma. WU-Minn HCP LS: Human Connectome Project Lifespan Pilot dataset. Kaggle dataset from BrainTumor.npy project, demographics information not available. Race/ethnicity not reported for subjects in any of the projects used.

Network Name	Trained On	Tested On	Accuracy (%)
Seed Networks			
GliomaNet_V1	AS, OD MRIs	AS, OD MRIs	97.22
GliomaNet_V2	AS, OD, OA MRIs	AS, OD, OA MRIs	94.26
TumorNet	Glioma, Meningioma, Healthy MRIs	Glioma, Meningioma, Healthy MRIs	86.14
T1Net	T1 sequence AS, OD, OA MRIs	T1 sequence AS, OD, OA MRIs	90.14
T2Net	T2 sequence AS, OD, OA MRIs	T2 sequence AS, OD, OA MRIs	92.38
Transfer Learned Networks			
ExpCamoGliomaNet_V1	Clear animals → Camo animals → AS, OD MRIs	AS, OD MRIs	98.61
ExpCamoGliomaNet_V2	Clear animals → Camo animals → AS, OD, OA MRIs	AS, OD, OA MRIs	94.88
ExpCamoTumorNet	Clear animals → Camo animals → Glioma, Meningioma, Healthy MRIs	Glioma, Meningioma, Healthy MRIs	88.76
T1_2Net	T1 sequence → T2 sequence AS, OD, OA MRIs	T2 sequence AS, OD, OA MRIs	94.48
T2_1Net	T2 sequence → T1 sequence AS, OD, OA MRIs	T1 sequence AS, OD, OA MRIs	85.92

Table 2. Trained neural networks and their accuracies. AS: astrocytoma, OD: oligodendroglioma, OA: oligoastrocytoma

Figure Legends

Figure 1 - Flowchart of data distribution from different sources. TCIA: The Cancer Imaging Archive. HCP: Human Connectome Project. BrainTumor.npy project dataset obtained from Kaggle. *M/G: Meningiomas/Gliomas. Glioma MRIs from Kaggle dataset were not differentiated into subtypes, used instead as single Glioma category to compare with meningiomas for imaging consistency from same source.

Figure 2 - Axial MRI cross-section examples of the different brain tumor categories used for training neural networks, both T1- (top row) and T2-weighted (bottom row). These include the three glioma types - astrocytoma, oligoastrocytoma, and oligodendroglioma – as well as healthy brain images.

Figure 3 - Examples of natural animal camouflage used in dataset to train ExpCamoNet prior to brain tumor transfer learning.

Figure 4 – Confusion matrices **(a, d)** and feature spaces **(b, c, e)** of GliomaNet (top row), both V1 and V2, and TumorNet (bottom row). Each image point belongs to a larger color-coded region of same category points. Example of T1 Meningioma MRI **(f)** used in TumorNet training also shown.

Figure 5 – Feature spaces of direct networks **(top row)** and corresponding networks after camouflage animal transfer learning **(bottom row)** across first two principal components.

Figure 6 – GliomaNet_V2 sensitivity maps of sample astrocytoma **(a)**, oligoastrocytoma **(b)**, and oligodendroglioma **(c)** images. TumorNet sensitivity maps of sample meningioma **(d)** and healthy **(e)** images. ExpCamoTumorNet ROC curves for glioma, meningioma, and healthy image hit rates vs false detections.

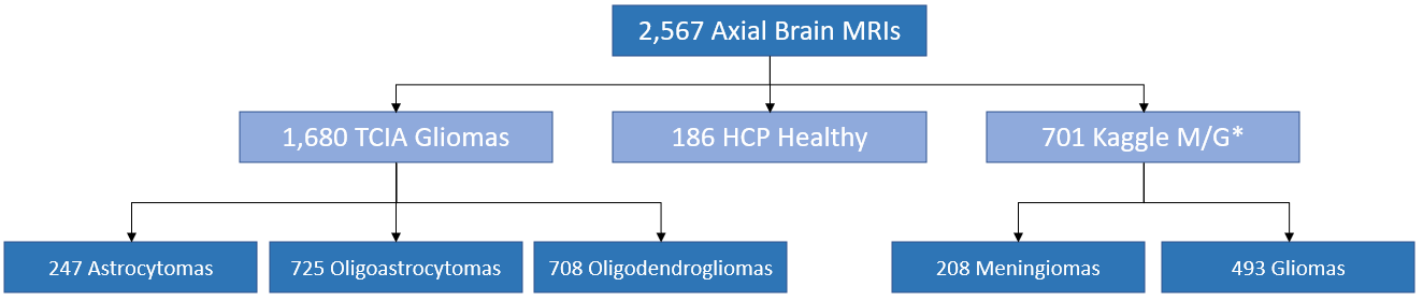

Figure 1 - Flowchart of data distribution from different sources. TCIA: The Cancer Imaging Archive. HCP: Human Connectome Project. BrainTumor.npy project dataset obtained from Kaggle. *M/G: Meningiomas/Gliomas. Glioma MRIs from Kaggle dataset were not differentiated into subtypes, used instead as single Glioma category to compare with meningiomas for imaging consistency from same source.

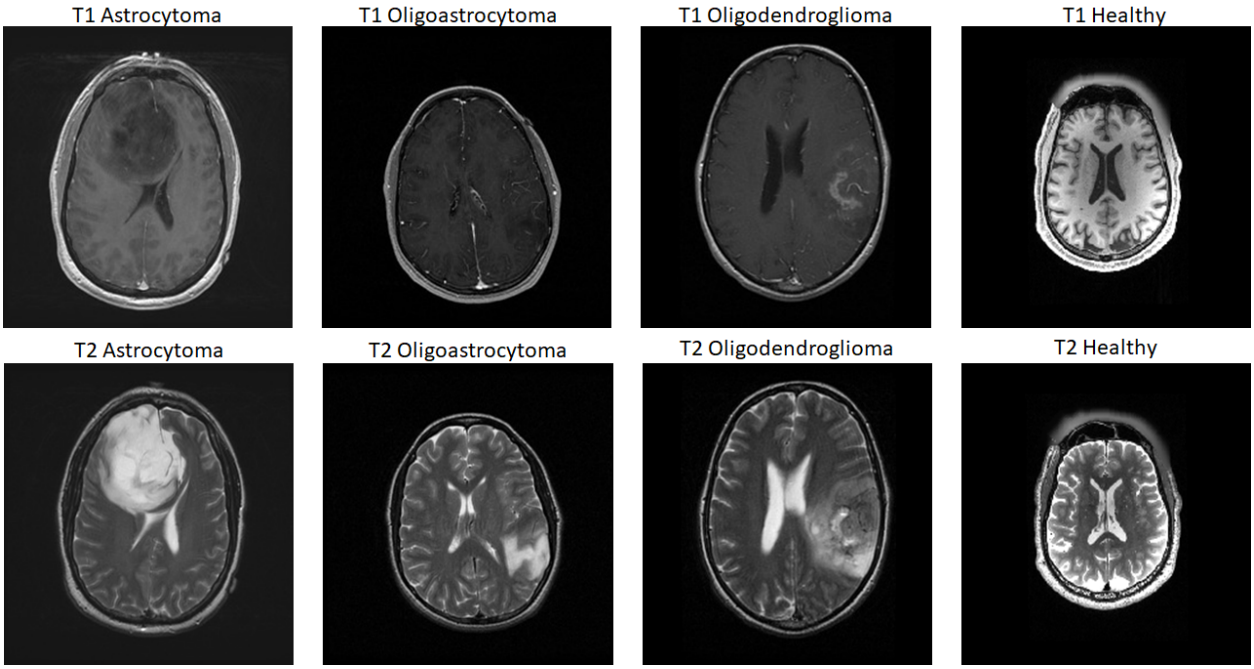

Figure 2 - Axial MRI cross-section examples of the different brain tumor categories used for training neural networks, both T1- (**top row**) and T2-weighted (**bottom row**). These include the three glioma types - astrocytoma, oligoastrocytoma, and oligodendroglioma – **as well as healthy brain images.**

Figure 3 - Examples of natural animal camouflage used in dataset to train ExpCamoNet prior to brain tumor transfer learning.

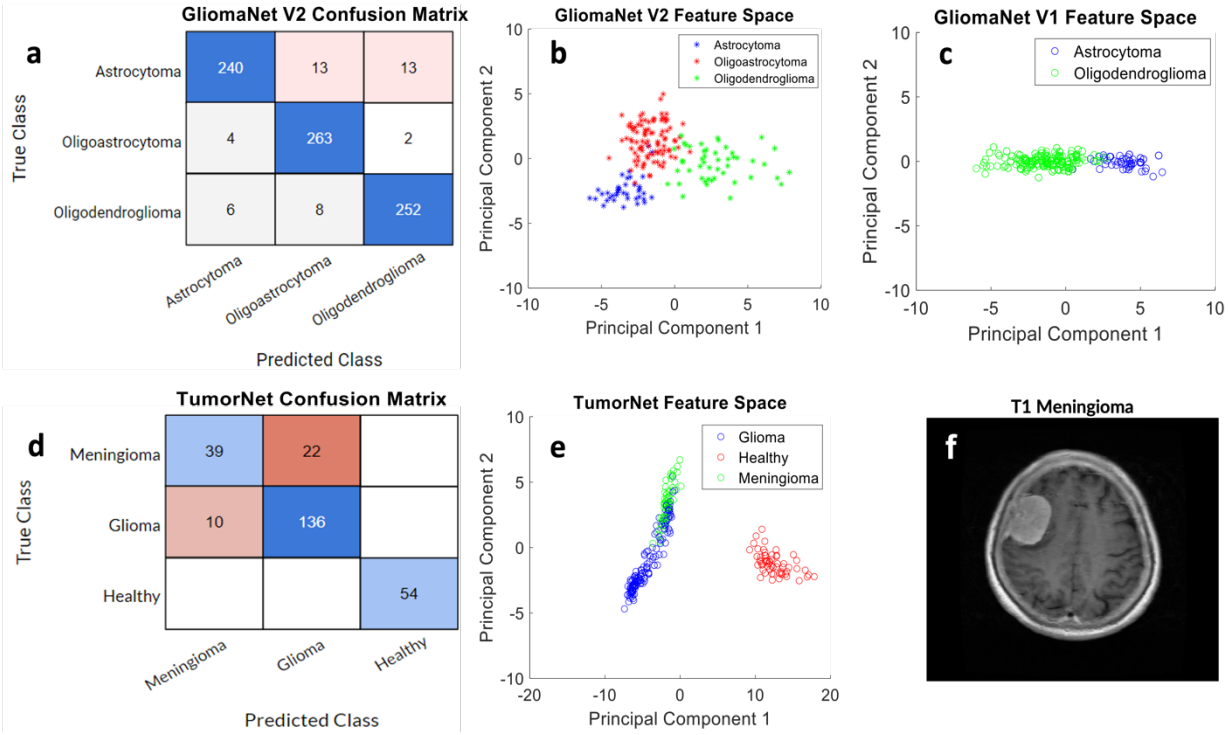

Figure 4 – Confusion matrices (**a, d**) and feature spaces (**b, c, e**) of GliomaNet **(top row**), both V1 and V2, and TumorNet (**bottom row**). Each image point belongs to a larger color-coded *region* of same category points. Example of T1 Meningioma MRI (**f**) used in TumorNet training also shown.

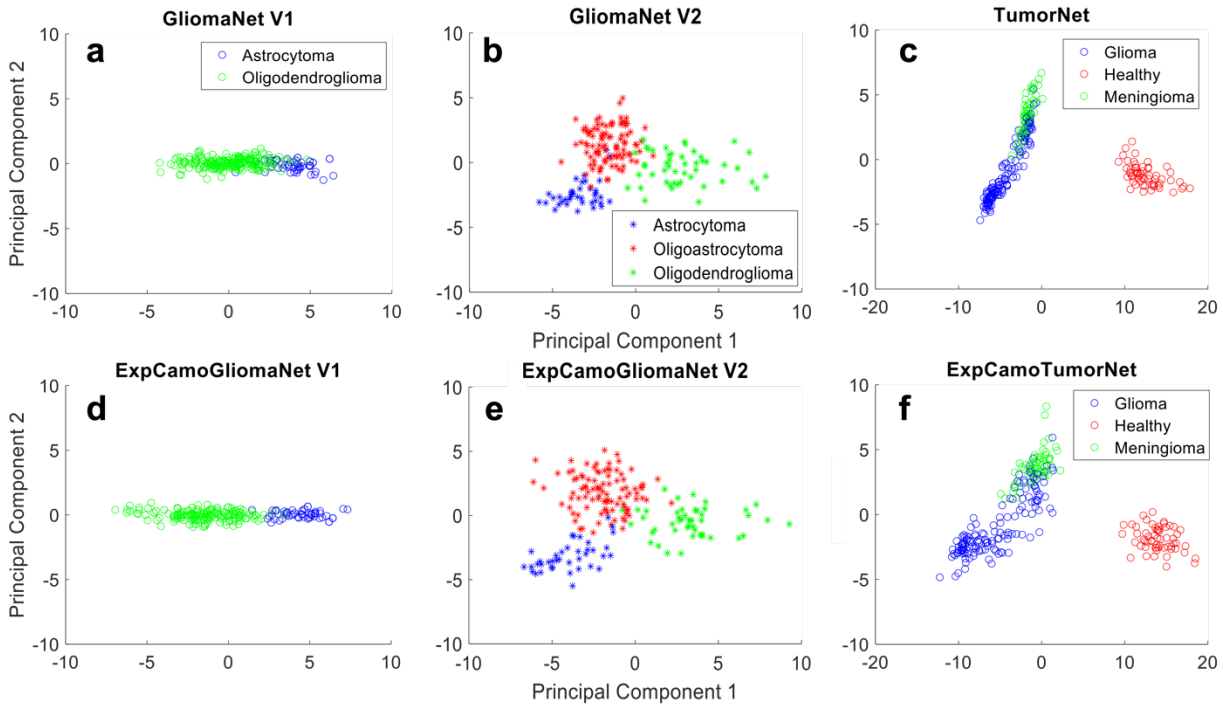

Figure 5 – Feature spaces of direct networks (**top row**) and corresponding networks after camouflage animal transfer learning (**bottom row**) across first two principal components.

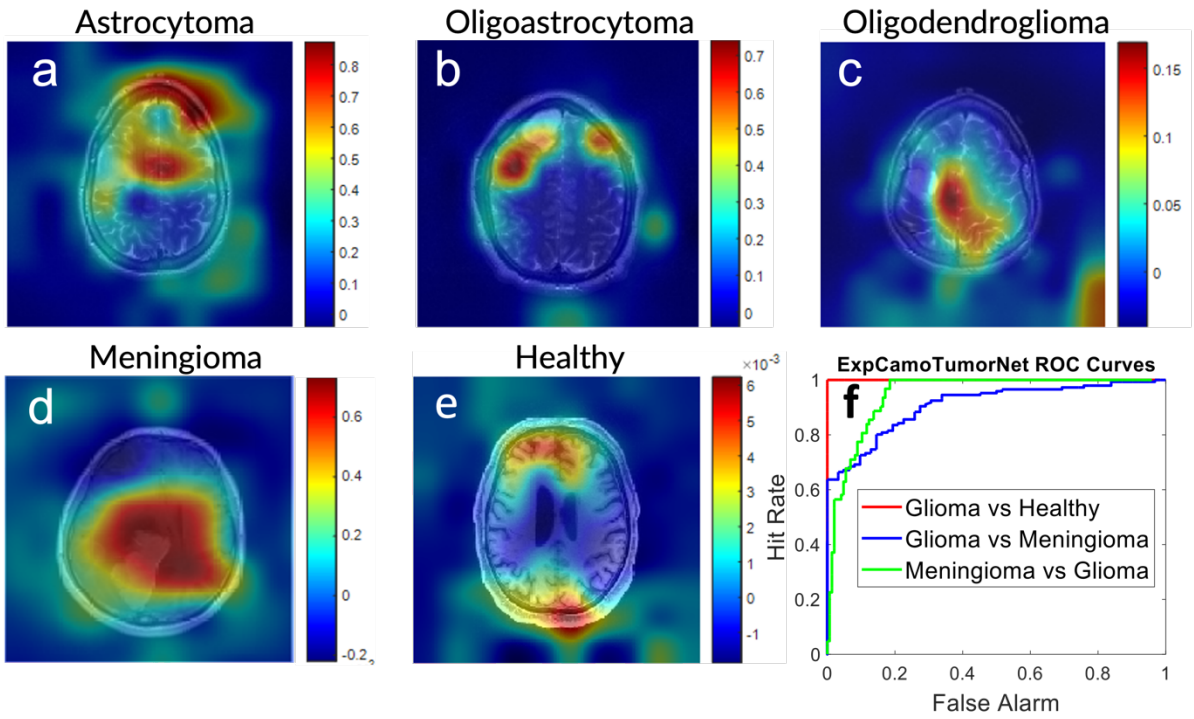

Figure 6 – GliomaNet_V2 sensitivity maps of sample astrocytoma **(a),** oligoastrocytoma **(b),** and oligodendroglioma **(c)** images. TumorNet sensitivity maps of sample meningioma **(d)** and healthy **(e)** images. ExpCamoTumorNet ROC curves for glioma, meningioma, and healthy image hit rates vs false detections.