

MRI to MGMT: Predicting Drug Efficacy for Glioblastoma Patients

Lichy Han
Stanford University
Biomedical Informatics Training Program
lhan2@stanford.edu

Maulik Kamdar
Stanford University
Biomedical Informatics Training Program
maulikrk@stanford.edu

Abstract

Glioblastoma Multiforme (GBM), a malignant brain tumor, is among the most lethal of all cancer categories. Temozolomide is the primary chemotherapy treatment for patients diagnosed with GBM. The methylation status of the promoter region of the O^6 -methylguanine methyltransferase (MGMT) gene may impact the efficacy and sensitivity of temozolomide, and hence may affect overall patient survival. Currently, invasive procedures are used to perform biopsy on the patient to determine the MGMT methylation status. However, these microscopic genetic changes may manifest as macroscopic morphological changes in the brain tumors that can be detected using magnetic resonance imaging (MRI). In this research, we use a compendium of brain MRI scans of GBM patients collected from The Cancer Imaging Archive (TCIA) to predict the methylation state of the MGMT promoter region of these patients. Our approach relies on a bidirectional convolutional recurrent neural network architecture (CRNN) that leverages the spatial and temporal aspects of these 3-dimensional MRI scans. We also compare and contrast our approach with other neural network architectures: convolutional neural networks (CNN), 3-dimensional CNN, pre-trained CNN with LSTM and GRU recurrent neural network (RNN) units, and unidirectional CRNN. Preliminary results indicate that our approach can easily overfit on the training set with a training accuracy of over 0.9, but can only achieve validation and test accuracy of between 0.55–0.6 even using higher regularization values. As all other architectures seem to demonstrate similar results, we may need better methods for MRI data curation, preprocessing and augmentation.

1. Introduction

Glioblastoma multiforme (GBM) is an aggressive brain cancer, with a median survival of only 15 months [30]. The efficacy of the first-line chemotherapy treatment, temozolomide, is in part dependent on the methylation status of the O^6 -methylguanine methyltransferase (MGMT) pro-

motor. MGMT removes alkyl groups from compounds and is one of the few known proteins in the DNA Direct Reversal Repair pathway in mammals [18]. Loss of the MGMT gene, or silencing of the gene through DNA methylation in the promoter region, may increase the carcinogenic risk after exposure to alkylating agents. Similarly, high levels of MGMT activity in cancer cells create a resistant phenotype by blunting the therapeutic effect of alkylating agents and may be an important determinant of treatment failure [12]. Thus, methylation of MGMT increases efficacy of alkylating agents such as temozolomide [30].

As such, MGMT promoter methylation status has important prognostic implications and can affect therapy selection in GBM. Currently, determining promoter methylation status is done using samples obtained from fine needle aspiration biopsies, which is an invasive procedure. However, several works have demonstrated that some genetic changes can manifest as macroscopic changes, which can be detected using magnetic resonance imaging (MRI) [8, 32].

Convolutional neural networks (CNN) are particularly adept at feature extraction and have excelled at many image classification tasks, including those in the medical domain based on MRI scans [27, 1]. Moreover, as MRI scans are 3-dimensional reconstruction of the human brain, they can be treated as volumetric objects or videos. Volumetric objects and sequences of image frames can be analyzed effectively using a 3-dimensional CNN (3D-CNN) [22] and recurrent neural networks (RNN) [23] respectively.

In this work, **we present an approach using a bidirectional convolutional recurrent neural network (CRNN) architecture on brain MRI scans to predict the methylation status of the MGMT gene.** We use a dataset of 5,235 MRI scan ‘sessions’, consisting of a total 458,951 MRI image frames, from 262 patients diagnosed with GBM from The Cancer Imaging Archive (TCIA) [28, 5]. Genomics data corresponding to these patients is retrieved from The Cancer Genome Atlas (TCGA) [31]. The CNN and RNN modules in the architecture are jointly trained in an end-to-end fashion. We evaluate our model using accuracy, precision, and recall. We compare our approach with other

neural network architectures: (i) CNN, (ii) 3D-CNN, (iii, iv) pretrained CNN with gated recurrent unit (GRU) or long short term memory (LSTM) RNN, and (v) unidirectional CRNN. We also visualize the hidden layers in our network to evaluate the importance of certain visual features.

2. Related Work

There have been some previous approaches using MRI and clinical features to predict MGMT promoter methylation status. Moreover, there has been research to predict disease subtypes or other gene mutations using imaging data of different modalities. For example, Korfiatis et al. extracted co-occurrence and run length texture features from MRI scans, and used random forests and SVM classifiers to predict MGMT methylation status in GBM patients with an AUROC statistic of 0.85 [14]. Levner et al. extract features from MRI scans using space-frequency texture analysis methods and use an L1-regularized neural network to predict MGMT methylation status with an accuracy of 0.87 [16]. Another similar method by Georgiadis et al. also extracted textural features from brain MRI scans, and used a probabilistic neural network for tumor characterization with an accuracy of 0.95 [9]. Chen et al. curate features from CT-scan and MRI images from TCGA-TCIA datasets to predict gene mutations in renal cancer [4]. Most of these approaches have not used advanced deep learning architectures and rely on manually extracted texture features.

An extensive literature review of deep learning methods and research papers in analyzing biomedical data was recently published [20]. The review outlines the use of neural networks, CNNs, RNNs and other emergent architectures to analyze imaging data, *-omics* data and biomedical literature. A recent study by Akkus et al. has used CNNs to extract features from MRI images and predict chromosomal aberrations [1]. Sarraf et al. use the LeNet-5 CNN architecture [15] to determine the most discriminative clinical features, and predict Alzheimer’s disease using brain MRI scans with an accuracy of 0.96 [27]. Roth et al. use multi-level CNNs for pancreas segmentation in abdominal CT-scans, and obtain a dice overlap statistic of 0.71 [26]. Payan et al. use deep learning methods, such as sparse autoencoders and 3D-CNN, to build an algorithm that can predict the disease status of a patient (3 states: healthy, mild-cognitive impairment and Alzheimer’s disease), based on an MRI scan of the brain [22]. They achieve a 3-way classification accuracy of 0.89 using 3D-CNN, as compared to 0.85 using a normal CNN. 3D-CNNs have also demonstrated a higher accuracy over conventional 2D-CNNs for detecting cerebral microbleeds in volumetric brain MRI scans [7].

Recently, recurrent neural networks have shown a lot of promise to analyze image sequences or 3-dimensional volumetric objects, such as videos and MRI scans. For example, Poudel et al. have developed a novel recurrent

fully-connected CNN to learn image representations from cardiac (heart) MRI scans and leverage inter-slice spatial dependences through internal memory units. The architecture combines anatomical detection and segmentation, and is trained end-to-end to reduce computational time [23]. Experimental evaluation for segmentation on two gold standard datasets gave accuracies of 0.95 and 0.97 respectively. Stollenga et al. developed a novel architecture, PyramidLSTM, to easily parallelize multi-dimensional LSTM networks, and leverage the spatial-temporal context in brain MRI scans that is lost by conventional CNNs. They achieved a 0.83 dice overlap statistic on tumor segmentation goals over two gold standard datasets [29]. Hybrid architectures, consisting of CNN filters and recurrent memory units, have been used in the past to classify videos in 400 categories [33]. Chen et al. developed a transferred-RNN, which incorporates a deep hierarchical visual feature extractor and a temporal sequence learning model, to analyze fetal ultrasound videos and detect fetal standard plane. They implement end-to-end training of the architecture and knowledge transfer between layers to deal with limited training data [3]. Kong et al. combined an LSTM-RNN with a CNN, and designed a new loss function, to detect the end-diastole and end-systole frames in cardiac MRI scans [13]. For review of other methods that use deep learning architectures to analyze medical images and videos for various tasks, we refer the reader to the literature review by Litjens et al. that summarizes 300 recent papers [17].

Most of the end-to-end deep learning methods listed here have either not been implemented for radio-genomic analyses, i.e. prediction of gene mutations from MRI scans, or rely on manually extracted features from MRI scans. In this work, we implement few variations of these methods for the goal of predicting MGMT methylation from MRI scans. Our primary approach is inspired from some of the methods that combine RNNs with CNNs [33, 3, 13].

3. Methods

Given that our MRI scans are similar to video objects with a variable number of frames, we implemented a bidirectional convolutional recurrent neural network (CRNN) architecture (**Figure 1**). Each image frame of the MRI scan is first input into a CNN. After multiple convolutional layers, the image then goes through two fully connected layers, so that the output from each image is a vector of length 512. All frames from one MRI scan session are then represented by a series of vectors, which are input into a many-to-one bidirectional RNN. The bidirectional RNN is dynamic and can adjust for variable-length sequences. Padding and bucketing of MRI scan sessions of similar length is carried out for efficient computation. The final output of the RNN is a binary classification of methylation status per MRI session. The advantage of combining CNN with RNN is that we are

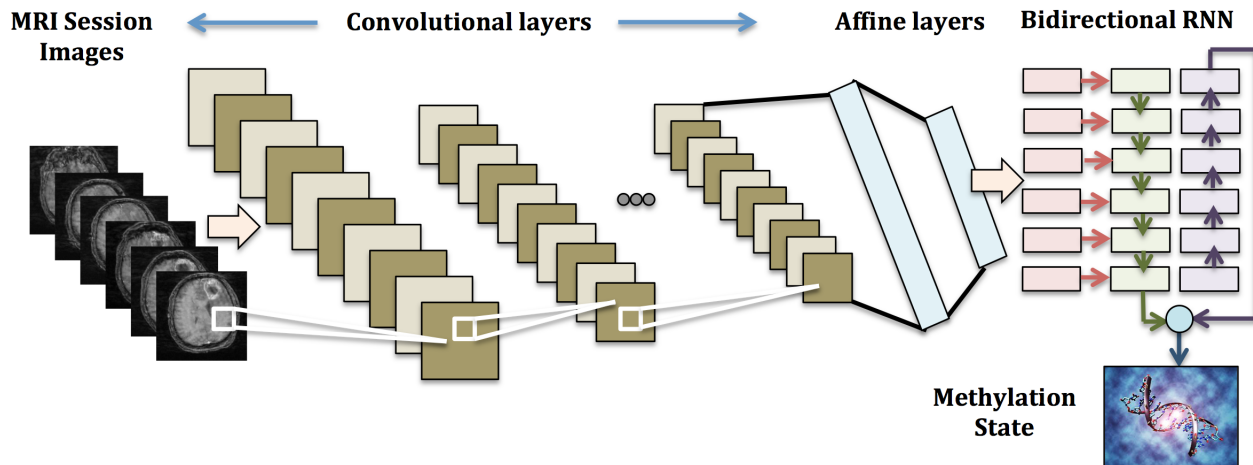


Figure 1. **CRNN Architecture.** Combining CNN and RNN to predict the methylation state from MRI session images.

able to leverage the image feature extraction capabilities of CNN with the variable sized input and temporal feature extraction capabilities of RNN.

We implemented several additional architectures for comparison with our CRNN. These include CNN on each frame, pretrained CNN followed by RNN with gated recurrent units (GRU), pretrained CNN followed by long short term memory (LSTM) RNN, 3-dimensional CNN (3D-CNN), and unidirectional CRNN. CNN was used for each image frame, where the methylation status associated for the video MRI session was attributed to each individual image. The remaining methods used full MRI sessions as input. For 3D-CNN, we limited our data to MRI sessions between 100 and 200 frames, and took the median 150 frames, padding with blank frames as necessary (**Figure 2**). For the pretrained CNN followed by RNN, we first trained and tuned our CNN, and generate an output vector from the fully connected layer with 512 neurons. These layers were then input separately into a many-to-one RNN. This is contrast to our CRNN architectures, in which we performed end-to-

end training of the CNN with the RNN.

Our CNN implementation used CS231N assignment 2 code [6]. The initial preprocessing of the DICOM-formatted MRI scan sessions was done using established guidelines [25]. To make our bidirectional RNN able to handle variable-length sequences efficiently, the bucketing iterator was inspired from an R2RT tutorial [24].

4. Dataset and Features

We are using the brain MRI scans of glioblastoma patients from The Cancer Imaging Archive (TCIA) and the methylation data, for those corresponding patients, from The Cancer Genome Atlas (TCGA).

4.1. Preprocessing of Methylation Data

We downloaded all 450 methylation data files available via TCGA. These samples correspond to 423 unique patients, with 16 patients having duplicate samples. We extracted methylation sites which are located in the minimal promoter and enhancer regions shown to have maximal methylation activity and affect MGMT expression [11, 10, 21]. Specifically, these methylation sites are cg02941816, cg12434587, and cg12981137, which are the same sites used in previous MGMT methylation studies using TCGA data [2]. Similar to Alonso, et al, we considered a methylation beta value of at least 0.2 to be a positive methylation site. As methylation of either the minimal promoter or the enhancer were shown to decrease transcription, we considered a patient to have a positive methylation status if any of the three sites were positive. Of the 450 methylation samples, 224 were negative, and 226 were positive.




Figure 2. **3D-CNN Architecture.** Uniform size MRI sessions are used with several 3D convolutional layers, followed by two fully connected layers with softmax classification.

4.2. Preprocessing of the MRI scans

We downloaded 5,235 MRI scan ‘sessions’ for 262 patients diagnosed with GBM from TCIA. Each brain MRI scan session can be envisioned as a 3-dimensional reconstruction of the brain (**Figure 3**). Each session consists of a set of image frames captured at a specific slice thickness and pixel spacing (based on the MRI machine specifications). The raw dataset contained a total of 458,951 image frames. For this research we selected ‘labeled’ *T1/T2 axial* MRI sessions for those patients for whom we had corresponding methylation data, and we reduced our dataset to 623 positive MRI scan sessions (80 patients, 45,159 images) and 628 negative MRI scan sessions (68 patients, 57,902 images).

These image frames are made available in a DICOM format (Digital Imaging and Communications in Medicine), a non-proprietary data interchange protocol, digital image format, and file structure for biomedical images and image-related information [19]. The image frames are grayscale (1-channel) and the DICOM format allows storage of other patient-related metadata (sex, age, weight, etc.) as well as image-related metadata (slice thickness, pixel spacing etc.). As these image frames may be generated by different MRI machines with varying slice thickness (*range*: 1 to 10) and pixel spacing, we normalize these attributes across different MRI sessions by resampling to a uniform slice thickness of 1.0 and pixel spacing of [1, 1]. After resampling, the




Figure 3. **MRI session.** A visualization of different MRI image frames in one MRI session. Each MRI session is a 3-dimensional reconstruction of the brain, and the GBM tumor (highlighted in red) can only be observed in few slices (Slices 50-90).

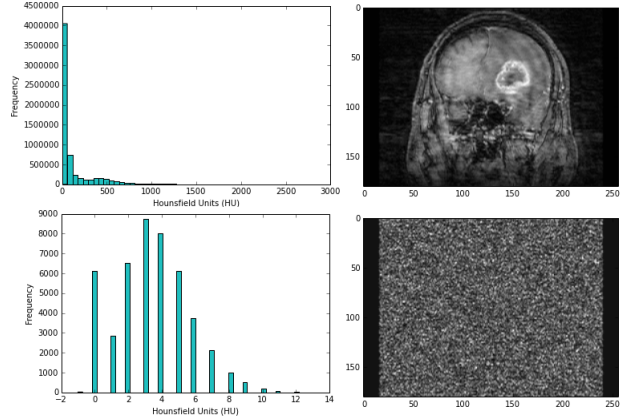


Figure 4. **Removing noisy images.** We use the distributions of Hounsfield units (which vary drastically) to determine if an image is a valid MRI scan, or has only noisy pixels.

number of image frames increased to 164,884 and 150,555 respectively for positive and negative MRI scan sessions.

MRI images are grayscale, and instead of RGB channel values, each pixel is assigned a numerical value termed the Hounsfield Unit (HU), which is a measure of radiodensity. We filter out those image frames that are ‘noisy’ by looking at the distribution of Hounsfield Units in the pixels. The distributions and the images are shown in **Figure 4**. Finally, we resize all images to 256×256 dimensions.

4.3. Data Augmentation and Training

We split our data into a 70% training set and 30% validation set by patient, so all images pertaining to each patient are in the same set. This resulted in 984 MRI sessions (117 patients) in the training set, and 269 (31 patients) in the validation set. From TCIA, we have an additional 22 patients with methylation data, but the MRI sessions are not ‘labeled’ as *T1/T2 axial*. We manually curated these sessions to derive our independent test set, which consists of 42 MRI sessions.

We doubled the amount of data to 1968 training, 538 validation, and 84 test sessions by using the MRI sessions in forward and in reverse order. We used a batch size of 7, which was chosen as the maximum size given our GPU memory constraints. For the larger architectures (CRNN and 3D-CNN), we first applied a pooling layer to reduce the dimensions of our images to 128 by 128, in order to achieve a batch size of 7.

We randomized the order of the training data based on the number of frames, bucketing MRI sessions with similar frame numbers. We padded MRI sessions within each bucket so all MRI sessions in each batch had the same number of frames, while the number of frames differed across batches. We trained using softmax cross entropy as our loss function using the Adam optimizer with learning rates rang-

Table 1. Deep learning architecture implementation details with training and validation set accuracies. Numbers following fully connected (FC) layers denote number of neurons. We used the rectified linear unit (ReLU) as our activation function, followed by batch normalization and dropout layers.

Method	Architecture	Hyperparameters	Training Accuracy	Validation Accuracy
CNN	[5x5 Conv-ReLU-BatchNorm-Dropout-2x2 Max Pool] x 2 [5x5 Conv-ReLU-BatchNorm-Dropout] x 1 [5x5 Conv-ReLU-BatchNorm-Dropout-2x2 Max Pool] x 1 [5x5 Conv-ReLU-BatchNorm-Dropout] x 1 [5x5 Conv-ReLU-BatchNorm-Dropout-2x2 Max Pool] x 1 [FC-ReLU-BatchNorm-Dropout] - 1024 [FC-ReLU-BatchNorm-Dropout] - 512 Softmax	L2 Regularization: 0.001 Learning Rate: 1e-5 Dropout Keep Probability: 0.7 Number of Filters: 32	0.92	0.61
3D-CNN	2x2x2 Max Pool [5x5x5 Conv-ReLU-BatchNorm-2x2x2 Max Pool]x3 [FC-ReLU-BatchNorm] - 1024 [FC-ReLU-BatchNorm] - 512 Softmax	Learning Rate: 1e-3 Number of Filters: 16	0.46	0.50
Pretrained CNN + GRU RNN	CNN Architecture [Uni-directional GRU with ReLU-Dropout] - 256 [FC-ReLU] - 256 Softmax	Learning Rate: 1e-4 Dropout Keep Probability: 0.7	0.97	0.52
Pretrained CNN + LSTM RNN	CNN Architecture [Uni-directional LSTM with ReLU-Dropout] - 256 [FC-ReLU] - 256 Softmax	Learning Rate: 5e-2 Dropout Keep Probability: 0.6	0.53	0.56
Unidirectional CRNN	2x2 Average Pool [5x5 Conv-ReLU-BatchNorm-Dropout-2x2 Max Pool] x 2 [5x5 Conv-ReLU-BatchNorm-Dropout] x 1 [5x5 Conv-ReLU-BatchNorm-Dropout-2x2 Max Pool] x 1 [FC-ReLU-BatchNorm-Dropout] - 512 [Uni-directional LSTM with ReLU-Dropout] - 256 [FC-ReLU] - 256 Softmax	L2 Regularization: 0.01 Learning Rate: 5e-3 Dropout Keep Probability: 0.7 Number of Filters: 16	0.65	0.56
Bidirectional CRNN	2x2 Average Pool [5x5 Conv-ReLU-BatchNorm-Dropout-2x2 Max Pool] x 2 [5x5 Conv-ReLU-BatchNorm-Dropout] x 1 [5x5 Conv-ReLU-BatchNorm-Dropout-2x2 Max Pool] x 1 [FC-ReLU-BatchNorm-Dropout] - 1024 [FC-ReLU-BatchNorm-Dropout] - 512 [Bi-directional GRU with ReLU-Dropout] - 256 [FC-ReLU] - 256 Softmax	L2 Regularization: 0.005 Learning Rate: 5e-5 Dropout Keep Probability: 0.8 Number of Filters: 16	0.99	0.56

ing from 5e-6 to 5e-1. We applied L2 regularization, with coefficients from 0.001 to 0.01 and dropout with keep probabilities ranging from 0.5 to 1.

5. Results and Discussion

The specific architectures we used for each of our learning algorithms, along with the hyperparameters and training and validation accuracies, are detailed in (Table 1). For CRNN, our test set results are shown in (Table 2). We assessed our results using accuracy, precision, and recall. Precision is the number of true positives divided by the total number of predicted positives, and recall is the number of true positives divided by the number of actual positives. In this work, a positive sample refers to a positive methylation status. The test data yielded an accuracy of 0.57, with a precision of 0.67 and recall of 0.71.

We tuned our Adam optimizer using learning rates from

5e-6 to 5e-1 (Figure 5). We noticed a sharp decrease in our ability to train our network with learning rates larger than 5e-4. We proceeded using the learning rate that yielded the highest validation set accuracy, which was 5e-5.

We recognize that our problem is challenging; though it has been tackled before using other machine learning algorithms with manually curated features, methylation status is not discernible by a human radiologist from MRI sessions. Though we observe overfitting in some of our architectures, including our bidirectional CRNN, increasing the dropout or L2 regularization coefficient resulted in training accuracies closer to 0.5 and a decrease in validation accuracy to 0.5 or less. We also experimented with the complexity of our model, assessing the effect of fewer convolutional layers, and using attention-based mechanisms in our RNN, but neither resulted significant gains in performance (data not shown).

Table 2. Contingency table showing classification results from the test set using bidirectional CRNN.

		Actual	
		Positive	Negative
Predicted	Positive	40	20
	Negative	16	8

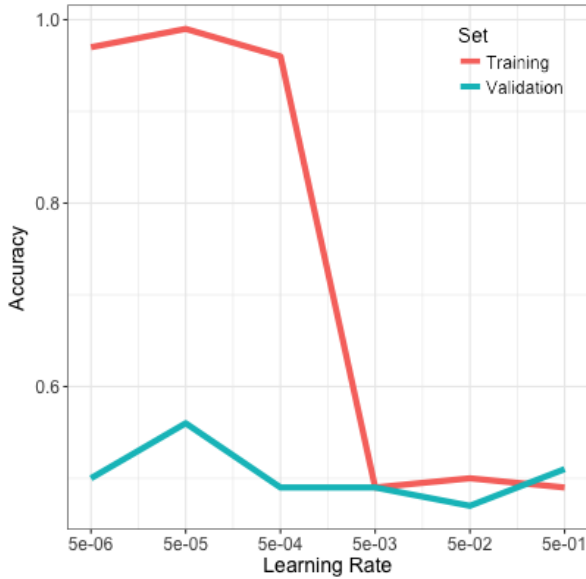


Figure 5. **Tuning the CRNN Learning Rate.** Training and validation set accuracies using learning rates ranging from 5e-6 to 5e-1.

The best results were seen using CNN without connecting the slices of the MRI together into MRI sessions. We believe this is the case due to the high number of training samples that can be used for CNN (103,061 images), versus the 1,968 sessions used to train the remaining architectures. However, we believe that CRNN, with the next best performance of 0.56 on the validation set and 0.57 on the test set, has the most potential for improvement with additional data curation and further data augmentation to increase the size of our training set.

We examined our classifier predictions in the test set, and show representative examples of true and false positives and negatives in **Figure 6**. In particular, it appears that our classifier tends to classify lesions with ring enhancement and high tumor contrast as having a positive methylation status. We see that the true positive and false positive both exhibit this ring enhancement, whereas it is not present in the true and false negatives. The tumors classified as positive also tend to have necrotic cores which further contribute to the high contrast and distinct borders of these tumors.

For the true positive example, we have additionally visualized the output from the filters of the first convolutional layer (**Figure 7**). As with many CNN architectures, the first layer places a heavy emphasis on edge detection, and

we can clearly see the outline of the cranium and the ring-enhanced tumor in each of these filters. Each filter also appears to show the brain slice at different contrasts, which is particularly exemplified by the varying intensity of the lateral ventricles. As specific tissues (white matter, grey matter, cerebrospinal fluid, etc) attenuate signal in different amounts, in some sense these filters are attempting to high-




Figure 6. **Test Set Examples.** Images are labeled with actual label/predicted label, and the tumors are highlighted using the red boxes.




Figure 7. **First Convolutional Layer Filter Output.** Filter outputs shown are from slice 70 of the true positive example from Figure 6

light different tissue types by varying the contrast, akin to a radiologist integrating multiple differently weighted scans to arrive at a clinical diagnosis.

6. Conclusions

In this work, we implemented and compared several neural network architectures to predict MGMT promoter methylation status using axial brain MRI scans from glioblastoma patients. Using bidirectional CRNN, we achieved a test set accuracy of 0.57, with a precision of 0.67 and recall of 0.71. We demonstrate that our classifier predictions highlight macroscopic features of tumor morphology, which may provide additional insight into glioblastoma pathogenesis and prognosis.

7. Future Work

Knowledge that is learned for one domain or a task via a CNN could benefit the training of a CNN for another domain or a task. The demonstration of these knowledge transfer mechanisms was shown by Chen et al. while training different CNNs to detect different fetal ultrasound positions [3]. We plan to implement these mechanisms to use the same MRI scans to detect other types of gene mutations, as well as use of other orientations of MRI scans (sagittal, etc.) for MGMT methylation prediction. Use of other loss functions, different from sparse softmax cross entropy loss, have also demonstrated better results (e.g. temporal structured loss by [13]).

Since, TCIA is an uncurated data source, we have observed that there may be certain MRI scans that are almost ‘black’ and as such hold no value for a trained radiologist. The image frames in these scans may have escaped our noise filtering steps, due to presence of very light out-lines. We plan to manually look at each MRI session to prepare our training and validation sets. Furthermore, we will perform additional data augmentation by using sagittal MRI scans. Essentially, since each MRI scan is a 3-dimensional object, sagittal MRI scans can be rotated along an axis to generate axial MRI scans. We also plan to explore multi-layer RNN architectures and better attention mechanisms.

Acknowledgments

We acknowledge CS231N staff for their guidance and the CNN code samples, as well as Sherlock, XStream and Google cloud computation resources.

References

[1] Z. Akkus, I. Ali, J. Sedlar, T. L. Kline, J. P. Agrawal, I. F. Parney, C. Giannini, and B. J. Erickson. Predicting 1p19q Chromosomal Deletion of Low-Grade Gliomas from MR Images using Deep Learning. nov 2016.

[2] S. Alonso, Y. Dai, K. Yamashita, S. Horiuchi, T. Dai, A. Matsunaga, R. Sánchez-Muñoz, C. Bilbao-Sieyro, J. C. Díaz-Chico, A. V. Chernov, A. Y. Strongin, and M. Perucho. Methylation of MGMT and ADAMTS14 in normal colon MUCOSA: biomarkers of a field defect for cancerization preferentially targeting elder African-Americans. *Oncotarget*, 6(5):3420–3431, 2015.

[3] H. Chen, Q. Dou, D. Ni, J.-Z. Cheng, J. Qin, S. Li, and P.-A. Heng. Automatic Fetal Ultrasound Standard Plane Detection Using Knowledge Transferred Recurrent Neural Networks. pages 507–514. Springer, Cham, 2015.

[4] X. Chen, Z. Zhou, K. Thomas, and J. Wang. MO-DE-207B-05: Predicting Gene Mutations in Renal Cell Carcinoma Based On CT Imaging Features: Validation Using TCGA-TCIA Datasets. *Medical Physics*, 43(6):3705–3705, jun 2016.

[5] K. Clark, B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, S. Moore, S. Phillips, D. Maffitt, M. Pringle, L. Tarbox, and F. Prior. The Cancer Imaging Archive (TCIA): Maintaining and Operating a Public Information Repository. *Journal of Digital Imaging*, 26(6):1045–1057, dec 2013.

[6] CS231N. CS231N Convolutional Neural Networks for Visual Recognition.

[7] Q. Dou, H. Chen, L. Yu, L. Zhao, J. Qin, D. Wang, V. C. Mok, L. Shi, and P.-A. Heng. Automatic Detection of Cerebral Microbleeds From MR Images via 3D Convolutional Neural Networks. *IEEE Transactions on Medical Imaging*, 35(5):1182–1195, may 2016.

[8] B. M. Ellingson. Radiogenomics and Imaging Phenotypes in Glioblastoma: Novel Observations and Correlation with Molecular Characteristics. *Current Neurology and Neuroscience Reports*, 15(1):506, jan 2015.

[9] P. Georgiadis, D. Cavouras, I. Kalatzis, A. Daskalakis, G. C. Kagadis, K. Sifaki, M. Malamas, G. Nikiforidis, and E. Solomou. Improving brain tumor characterization on MRI by probabilistic neural networks and non-linear transformation of textural features. *Computer Methods and Programs in Biomedicine*, 89(1):24–32, jan 2008.

[10] L. C. Harris, P. M. Potter, K. Tano, S. Shiota, S. Mitra, and T. P. Brent. Characterization of the promoter region of the human O6-methylguanine- DNA methyltransferase gene. *Nucleic Acids Research*, 19(0305-1048 SB - M SB - X):6163–6167, 1991.

[11] L. C. Harris, J. S. Remack, and T. P. Brent. Identification of a 59 bp enhancer located at the first exon/intron boundary of the human O6-methylguanine DNA methyltransferase gene. *Nucleic Acids Research*, 22(22):4614–4619, 1994.

[12] M. E. Hegi, A.-C. Diserens, T. Gorlia, M.-F. Hamou, N. de Tribolet, M. Weller, J. M. Kros, J. A. Hainfellner, W. Mason, L. Mariani, J. E. Bromberg, P. Hau, R. O. Mirimanoff, J. G. Cairncross, R. C. Janzer, and R. Stupp. MGMT Gene Silencing and Benefit from Temozolomide in Glioblastoma. *New England Journal of Medicine*, 352(10):997–1003, mar 2005.

[13] B. Kong, Y. Zhan, M. Shin, T. Denny, and S. Zhang. Recognizing End-Diastole and End-Systole Frames via Deep Temporal Regression Network. pages 264–272. Springer, Cham, 2016.

- [14] P. Korfiatis, T. L. Kline, L. Coufalova, D. H. Lachance, I. F. Parney, R. E. Carter, J. C. Buckner, and B. J. Erickson. MRI texture features as biomarkers to predict MGMT methylation status in glioblastomas. *Medical Physics*, 43(6):2835–2844, may 2016.
- [15] Y. Lecun. COMPARISON OF LEARNING ALGORITHMS FOR HANDWRITTEN DIGIT RECOGNITION.
- [16] I. Levner, S. Drabycz, G. Roldan, P. De Robles, J. G. Cairncross, and R. Mitchell. Predicting MGMT Methylation Status of Glioblastomas from MRI Texture. pages 522–530. Springer, Berlin, Heidelberg, 2009.
- [17] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. W. M. van der Laak, B. van Ginneken, and C. I. Sánchez. A Survey on Deep Learning in Medical Image Analysis. feb 2017.
- [18] G. P. Margison, A. C. Povey, B. Kaina, and M. F. Santibáñez Koref. Variability and regulation of O6-alkylguanine-DNA alkyltransferase. *Carcinogenesis*, 24(4):625–35, apr 2003.
- [19] P. Mildemberger, M. Eichelberg, and E. Martin. Introduction to the DICOM standard. *European Radiology*, 12(4):920–927, apr 2002.
- [20] S. Min, B. Lee, and S. Yoon. Deep learning in bioinformatics. *Briefings in Bioinformatics*, page bbw068, jul 2016.
- [21] T. Nakagawachi, H. Soejima, T. Urano, W. Zhao, K. Higashimoto, Y. Satoh, S. Matsukura, S. Kudo, Y. Kitajima, H. Harada, K. Furukawa, H. Matsuzaki, M. Emi, Y. Nakabeppu, K. Miyazaki, M. Sekiguchi, and T. Mukai. Silencing effect of CpG island hypermethylation and histone modifications on O6-methylguanine-DNA methyltransferase (MGMT) gene expression in human cancer. *Oncogene*, 22(55):8835–8844, 2003.
- [22] A. Payan and G. Montana. Predicting Alzheimer’s disease: a neuroimaging study with 3D convolutional neural networks. feb 2015.
- [23] R. P. K. Poudel, P. Lamata, and G. Montana. Recurrent Fully Convolutional Neural Networks for Multi-slice MRI Cardiac Segmentation. aug 2016.
- [24] R2RT. Recurrent Neural Networks in Tensorflow III - Variable Length Sequences - R2RT.
- [25] RADDQ. DICOM Processing and Segmentation in Python Radiology Data Quest.
- [26] H. R. Roth, L. Lu, A. Farag, H.-C. Shin, J. Liu, E. Turkbey, and R. M. Summers. DeepOrgan: Multi-level Deep Convolutional Networks for Automated Pancreas Segmentation. jun 2015.
- [27] S. Sarraf and G. Tofghi. Classification of Alzheimer’s Disease using fMRI Data and Deep Learning Convolutional Neural Networks. mar 2016.
- [28] L. Scarpace, T. Mikkelsen, S. Cha, S. Rao, S. Tekchandani, D. Gutman, J. Saltz, B. J. Erickson, N. Pedano, A. E. Flanders, J. Barnholtz-Sloan, Q. Ostrom, D. Barboriak, and L. J. Pierce. Radiology Data from The Cancer Genome Atlas Glioblastoma Multiforme [TCGA-GBM] collection. jan 2016.
- [29] M. F. Stollenga, W. Byeon, M. Liwicki, and J. Schmidhuber. Parallel Multi-Dimensional LSTM, With Application to Fast Biomedical Volumetric Image Segmentation. jun 2015.
- [30] J. P. Thakkar, T. A. Dolecek, C. Horbinski, Q. T. Ostrom, D. D. Lightner, J. S. Barnholtz-Sloan, and J. L. Villano. Epidemiologic and Molecular Prognostic Review of Glioblastoma. *Cancer Epidemiology and Prevention Biomarkers*, 23(10), 2014.
- [31] J. N. Weinstein, E. A. Collisson, G. B. Mills, K. R. Mills Shaw, B. A. Ozenberger, K. Ellrott, I. Shmulevich, C. Sander, and J. M. Stuart. The Cancer Genome Atlas Pan-Cancer analysis project. *Nature Publishing Group*, 45, 2013.
- [32] S. Yamamoto, D. D. Maki, R. L. Korn, and M. D. Kuo. Radiogenomic Analysis of Breast Cancer Using MRI: A Preliminary Study to Define the Landscape. *American Journal of Roentgenology*, 199(3):654–663, sep 2012.
- [33] J. Yue, H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond Short Snippets: Deep Networks for Video Classification.