

Artificial Intelligence in Glioma Imaging: Challenges and Advances

Weina Jin¹, Mostafa Fatehi², Kumar Abhishek¹,
Mayur Mallya¹, Brian Toyota³, and Ghassan Hamarneh¹

¹School of Computing Science, Simon Fraser University, Burnaby, Canada

²Division of Neurosurgery, The University of British Columbia, Vancouver, Canada

³Department of Surgery, Kingston General Hospital, Queen's University, Kingston, Canada

E-mail: weinaj@sfu.ca, fatehi@alumni.ubc.ca, kabhishe@sfu.ca,
mmallya@sfu.ca, hamarneh@sfu.ca

Abstract.

Primary brain tumors including gliomas continue to pose significant management challenges to clinicians. While the presentation, the pathology, and the clinical course of these lesions are variable, the initial investigations are usually similar. Patients who are suspected to have a brain tumor will be assessed with computed tomography (CT) and magnetic resonance imaging (MRI). The imaging findings are used by neurosurgeons to determine the feasibility of surgical resection and plan such an undertaking. Imaging studies are also an indispensable tool in tracking tumor progression or its response to treatment. As these imaging studies are non-invasive, relatively cheap and accessible to patients, there have been many efforts over the past two decades to increase the amount of clinically-relevant information that can be extracted from brain imaging. Most recently, artificial intelligence (AI) techniques have been employed to segment and characterize brain tumors, as well as to detect progression or treatment-response. However, the clinical utility of such endeavours remains limited due to challenges in data collection and annotation, model training, and the reliability of AI-generated information.

We provide a review of recent advances in addressing the above challenges. First, to overcome the challenge of data paucity, different image imputation and synthesis techniques along with annotation collection efforts are summarized. Next, various training strategies are presented to meet multiple desiderata, such as model performance, generalization ability, data privacy protection, and learning with sparse annotations. Finally, standardized performance evaluation and model interpretability methods have been reviewed. We believe that these technical approaches will facilitate the development of a fully-functional AI tool in the clinical care of patients with gliomas.

Keywords: Glioma imaging, Brain radiomics, Machine learning, Deep learning

1. Introduction

Gliomas are the most common primary tumor of the brain and spine, representing 80% of malignant brain tumors. These lesions arise from astrocytes, oligodendrocytes or ependymal cells and cause significant morbidity and mortality. The most aggressive subtype, glioblastoma accounts for $\sim 45\%$ of all gliomas and affected patients have a 5-year overall survival of $\sim 5\%$ [1]. These tumors exhibit an incredible degree of inter and intra-tumoral heterogeneity which leads to variable disease presentation and prognosis (see Figure 1 as an example). This variability was considered by Bailey and Cushing when they initially described these lesions ninety years ago [2]. The advent of molecular genetics and improvements in techniques like immunohistochemistry and DNA sequencing have led to novel diagnosis classification systems that more accurately predict disease progression and overall survival. However, the personalization of treatment with targeted therapies remains an ongoing challenge for most gliomas.

Astrocytomas and oligodendrogliomas constitute the great majority of brain gliomas in adults [3, 4]. These tumors were classically described and graded based upon histopathological features. However, these diagnostic criteria led to significant inter-observer differences and did not correspond to prognosis very well [5]. Seminal studies by Cairncross et al. highlighted the importance of chromosome 1p/19q co-deletion for diagnosing oligodendrogliomas [6]. Later studies found that diffuse gliomas with a mutation in the isocitrate dehydrogenase (IDH, a key rate-limiting enzyme in the energy metabolism cycle [7]) gene had a much better overall survival than unmutated cases (IDH wild-type) [8, 9]. These findings led scientists to incorporate genetic characteristics in the new classification of gliomas in 2016 [3, 4].

The recent primacy of molecular genetics in diagnosing gliomas has coincided with improving knowledge of the aberrant pathways involved in disease progression and susceptibility to chemotherapy and radiation therapy. Consequently, clinicians have attempted to protocolize the treatment of low- and high-grade gliomas. Stupp et al. found that concurrent radiation and temozolamide (TMZ), an alkylating chemotherapeutic agent, improved overall survival in patients with glioblastomas [10]. Moreover, the greatest benefits from TMZ are seen in patients with methylguanine methyl transferase (MGMT) promoter methylation [11, 12]. Adjuvant chemotherapy and radiation are also beneficial for some patients with low grade gliomas. The European EORTC-NCIC 26351 trial showed that chemotherapy (procarbazine, lomustine and vincristine) after radiation increased progression free survival (PFS) and survival [13].

In addition to studying glioma genetics which focuses on specific gene mutations, there is growing interest in understanding these lesions at higher levels. Specifically, genomic studies aim to understand the role of the whole genome on developing gliomas while proteomic studies have investigated the complex interaction of various proteins (normal and aberrant) in affected cells [14]. These fields produce large quantities of data, and hence, rely upon novel computational methods to analyze and confirm findings. As mentioned, brain tumor imaging techniques are non-invasive and relatively cheap. Thus,




Figure 1. There is considerable variability in MR images obtained from patients with gliomas. Panels A and B were obtained from patients with glioblastoma (grade IV astrocytoma). Panels C and D from patients with a low grade astrocytoma (IDH mutated). Panels E and F from patients with oligodendroglioma (IDH mutated and 1p/19q codeleted). All images are axial T1-weighted with gadolinium.

they are integral in diagnosis, surgical planning and prognostication. The prospect of obtaining genomic information from brain imaging led to the establishment of a new field: radiomics [15, 16, 17].

Glioma radiomics aim at extracting clinically relevant features from glioma imaging as a non-invasive image biomarker, to provide information for glioma diagnosis, treatment, or prognosis. The earlier studies used statistical methods to identify the significant radiographic features associated with the clinical outcome. In recent years, artificial intelligence (AI), represented by machine learning (ML) and deep learning (DL), has enhanced glioma radiomics research. This is largely because of the increasing volume of brain imaging data, growing computational power, and advances in AI algorithms. Compared to traditional statistical analysis methods, machine and deep learning are able to leverage larger amounts of clinical data with high dimensionality, yet they also require more training data, longer training time and more computational resources, and sacrifice interpretability (and in turn trustworthiness) for predictive power. DL is a subset of ML which is more prominent for the above problems, but it has the advantage of automatically learning features from raw input data, instead of manually engineering features as in ML. Because of the




Figure 2. Growth of publications on glioma imaging and AI in the past 20 years by PubMed searching (the search keywords can be found at: https://bit.ly/AI_in_glioma_imaging).

superior predictive power, ML and DL have been widely applied to a variety of glioma imaging-related tasks including: (1) **brain tumor segmentation**: quantifying tumor volume or segmenting the tumor region for downstream analysis tasks, and (2) **predictive tasks (classification or regression)**: identifying the tumor type (e.g., distinguishing oligodendrogliomas from astrocytomas), grade, molecular subtypes, or genetic mutation [17, 18], and predicting patients’ treatment response, length of survival, prognosis, or recurrence (e.g., differentiating glioma recurrence from radiation necrosis and pseudo-progression [18]). These advances have been summarized in previous surveys on DL in glioma imaging [19, 20, 21], ML in glioma imaging applications [22, 23], brain tumor radiomics [24, 25] and neuroimaging biomarkers [26].

The introduction of AI to neuro-oncology imaging has been met with great enthusiasm from clinicians. This is partially evidenced by the exponential growth in the number of articles reporting findings. Figure 2 shows the increasing number of publications involving ML/DL on glioma imaging in the past two decades. Despite significant breakthroughs in this field, AI has not yet been applied to the clinical care of patients with gliomas, due to practical challenges in data collection, model training, and the lack of trust in performance and generalizability. To bridge the clinical utility gap in applying AI to neuro-oncology, our review focuses on recent technical advances addressing various challenges.

2. Challenges and technical approaches of applying AI in glioma imaging

Although machine and deep learning techniques have exhibited great potential in analyzing glioma images, their implementation in clinical care remains an elusive goal.




Figure 3. The three major challenges and technical approaches to applying AI in glioma imaging.

Several recent reviews have broadly summarized the challenges in applying AI to clinical medicine (some of which also apply to neuro-oncology) [27, 28, 29]. These challenges involve the full life-cycle of developing an AI model, from **1) obtaining the training data**, to **2) training the AI models**, to **3) evaluating and deploying the AI model to clinical settings**. Non-technical challenges also pose practical constraints in developing AI techniques, including: patient data safety and privacy issue; ethical, legal and financial barriers to developing and distributing tools that may impact a patient’s treatment course; medical authority regulation, usability evaluation, clinical acceptance, and medical education around the implementation of the AI-assistive tools. Next, we review the recent advances in glioma imaging research addressing the three main challenges. Figure 3 illustrates the relationship of the challenges and their corresponding approaches.

2.1. Obtaining data and annotations

Machine and deep learning models require large amounts of training data to learn from, so that their performances can be generalizable in real-world applications. One prominent example is the ImageNet dataset (containing over 14 million images of natural scenes) [30] that has greatly advanced AI analysis of images. However, the field of glioma imaging is lacking such comparable large-scale, consolidated public datasets. The prevalence of gliomas is relatively low; which, combined with the acquisition cost and time, hinders the collection of large datasets. Additionally, medical data is protected by patient privacy and safety laws, making it difficult to share these images. Furthermore, the cost of data labelling can be quite high since it requires annotation by medical experts [26]. Figure 4 shows a comparison of several popular natural, medical and brain imaging datasets with respect to these image acquisition factors. As shown in the figure, natural images (ImageNet) and certain medical imaging modalities, such




Figure 4. The comparison of natural, medical and brain image datasets with respect to their dataset size (y-axis), the image acquisition time (x-axis), and the size per image (shown by the dot size). **Brain** imaging datasets (blue dots in the figure): ADNI (Alzheimer’s Disease Neuroimaging Initiative) [31], ABIDE (Autism Brain Imaging Data Exchange) [32, 33], BraTS (Brain Tumor Segmentation) [34], and ISLES (Ischemic Stroke Lesion Segmentation) [35]. Other **medical** imaging datasets (red dots): OCT (retinal optical coherence tomography images) [36], CheXpert (chest radiographs) [37], ISIC (International Skin Image Collaboration Melanoma Project) [38], and fastMRI (knee MRI) [39]. **Natural** image dataset (green dot): ImageNet [30].

as dermoscopic images (ISIC) are quick to acquire and generate small file sizes, whereas brain imaging datasets take longer to acquire and are large in size.

Several strategies have been employed to meet the challenge of limited data and annotations: (i) data sharing platforms and initiatives to build a collective glioma imaging database; (ii) data imputation and synthesis techniques to address issues of missing data and lack of sufficient data, and (iii) active learning and crowd-sourcing to facilitate data annotation. Next, we discuss recent advances in these approaches.

2.1.1. Clinical data collection efforts

International and multi-institutional efforts have been made in the past decades towards collecting and sharing glioma imaging data. The TCIA/TCGA (The Cancer Imaging Archive/The Cancer Genome Atlas) service provides large publicly-available datasets, and is a part of a collective effort by the National Institutes of Health (NIH), the National Human Genome Research Institute (NHGRI), and multiple hospitals to encourage the sharing of clinical, imaging (including radiographic and pathological images) and genomic datasets for cancer research [40]. There are about 1,200 raw

brain images related to glioma in the TCIA dataset. The Brain Tumor Segmentation (BraTS) challenge dataset [34] derived from TCIA contains 461 expert-assessed scans (including low- and high-grade gliomas) from 19 institutions. As the largest publicly-available glioma imaging dataset, BraTS is critical for benchmarking and has helped advance AI applications in glioma image analysis.

2.1.2. Data imputation and synthesis methods

As mentioned, data heterogeneity and data paucity are two prominent issues in the clinical adoption of AI in glioma imaging. As such, data imputation approaches have been proposed to handle scenarios such as missing MRI sequences, missing voxel intensities, low-resolution or artifact-degraded scans. Similarly, data synthesis and augmentation techniques have been proposed to expand the available imaging data.

- **Data imputation methods to deal with missing data**

Several conventional machine learning methods have been proposed for MR image imputation. Cruz et al. [41] simulated two types of missing morphological measures in T_1 -weighted brain MRI scans: missing completely at random (MCAR) and missing not at random (MNAR). They also evaluated four data imputation techniques representative of distinct statistical approaches: substituting the missing value with the mean of all subjects (univariate approach), k-nearest neighbors (inferential approach), low-rank matrix approximation (analytical approach), and random forest (statistical learning approach). The random forest-based approach was shown to be the most accurate in recovering the gold standard results. Jog et al. [42] proposed a random forest-based supervised image imputation approach, which performs a non-linear regression in the feature space to predict the intensities in various contrast-enhanced MR sequences. They evaluated their proposed approach by synthesizing 1) T_2 -weighted images from skull-stripped T_1 -weighted images 2) FLuid Attenuated Inversion Recovery (FLAIR) images from T_1 -weighted, T_2 -weighted, and P_D -weighted images and 3) whole head T_2 -weighted images (non-skull-stripped) from T_1 -weighted images. Their method demonstrated an improvement in the synthetic image quality over state-of-the-art image synthesis algorithms, and the downstream segmentation task performed similarly for real and imputed images. Dalca et al. [43] leveraged the shared intrinsic local fine scale structural similarity in a collection of medical images. They used a generative model, Gaussian mixture model (GMM), to perform brain MRI restoration in a sparsely sampled scan, and demonstrated superior performance as compared to state-of-the-art up-sampling super-resolution methods. They cautioned, however, that imputed data should not be used in clinical evaluation; instead the brain mask obtained from the restored scans can be applied to the original scans and improve subsequent analyses.

Recently, deep learning-based approaches have been proposed for data imputation. Chartsias et al. [44] proposed a multi-input multi-output deep model trained end-to-end, with encoders and decoders for each MRI sequence. The latent representations

from multiple encoders were fused, and a combination of loss functions were used to ensure that the multiple latent representations from multiple modalities are similar. Another strategy for MRI imputation relies upon generative adversarial networks (GANs) [45]. Welander et al. [46] compared the performances of two image-to-image translation methods: CycleGAN [47] and UNIT [48] for the generation of T_1 and T_2 MRI sequences; the generated T_2 sequences had higher mean absolute error (MAE) and lower peak signal-to-noise ratio (PSNR) than the generated T_1 sequences. Several other GAN-based methods for imputing missing MRI sequences have been proposed [49, 50, 51, 52, 53], including CollaGAN [54] which investigated the reconstruction of various sequences. They concluded that while T_1 , T_2 , and T_2 -FLAIR can be imputed from other sequences, contrast-enhanced T_1 -weighted sequences cannot be efficiently imputed.

- **Data synthesis methods for data augmentation**

Similar to data imputation techniques using GANs, several data synthesis-based augmentation approaches using GANs have been proposed. To synthesize new images, Han et al. [55] trained two generative models, a deep convolutional GAN (DCGAN) [56] and a Wasserstein GAN (WGAN) [57], on the BraTS dataset to synthesize new brain MR images. To evaluate their performance, the authors conducted a ‘visual Turing test’: doctors were asked to differentiate real MRI from the generated images. Images generated by WGAN (classification accuracy ranged from 53% to 64%) were mis-classified by doctors more frequently than those generated by DCGAN (classification accuracy ranged from 54% to 70%); thus inferring that WGAN produced more visually realistic images. A similar visual assessment was also conducted by Kazuhiro et al. [58] where five radiologists (including two experienced neuroradiologists) attempted to classify a random selection of 50 real and synthesized images. The classification accuracy for the neuroradiologists were 55% and 30%; indicating that many synthesized images were practically indistinguishable from real MRIs. Shin et al. [59] used an image-to-image translation framework [60] to predict tumor segmentation masks from brain MR scans and to generate brain MR scans given tumor masks. The latter was then used to augment the segmentation training dataset. Using only the synthetic MR images to train the model followed by fine-tuning on 10% of the real images, the segmentation performance was superior to using only the real images. Jin et al. [61] proposed an image-to-image translation framework for synthesizing brain MR images from brain CT images. They leveraged a combination of adversarial loss, dual cycle-consistent loss, and voxel-wise loss strategies to train. Models trained using the paired data (i.e., CT and corresponding MRI scans from the same patient) and the unpaired data (two independent collections of CT and MRI images) together achieved higher PSNR and lower MAE compared to models trained using only either paired or unpaired data.

2.1.3. Obtaining data annotations

Classically, MRI annotation has required the time and expertise of medical experts, which imposes higher costs than for natural images. As such, there has been interest in 1) making the best use of expert input by annotating in an active learning setting, or 2) bypassing physician expertise and asking non-experts to annotate. We will expound upon each of these strategies below.

Active learning (AL) is a possible workaround to training models with limited annotated data. AL is an ML technique that reduces the annotation workload by automatically deciding which representative samples should be annotated in order to train a model as quickly and efficiently as possible. Damangir et al. [62] proposed using a cascade of support vector machines (SVMs) with an AL scheme to segment white matter lesions from brain MR scans (a problem compounded by the dataset being biased towards normal tissue). They observed that the proposed method yielded accurate results while also achieving the highest sensitivity and specificity scores amongst all methods compared. Konyushkova et al. [63] proposed leveraging the feature and geometric uncertainties from images to increase the efficiency of AL for tumor segmentation from brain MR images. In a later study, they modeled AL as a regression task [64], essentially predicting the expected generalizability of a trained classifier when annotating a specific data point. They used features extracted from a random forest regressor trained on synthetic 2D data to perform tumor segmentation from brain MR scans, and obtained superior results as compared to other active learning methods (such as uncertainty sampling [65] and active learning by learning [66]) with considerably fewer annotated samples.

Another popular approach of obtaining a large number of annotations quickly is to crowd-source the annotation process; a strategy which has led to rapid advances in computer vision [67]. The aforementioned ImageNet dataset [30] has been annotated using a crowd-sourcing platform (Amazon Mechanical Turks). Although there may be few medical experts within the labelling group, crowd-sourcing has been shown to be effective in creating large quantities of annotated data; and, it is faster and cheaper than annotation by medical experts [68]. In neuroscience research, crowd-sourcing and gamification have helped neuroscientists to explore brain networks by identifying neurons and their synaptic connections [69, 70]. Prior to annotating, the crowd will be trained by learning from a few examples labelled by medical experts. The labelling results will need to be aggregated from multiple workers to ensure the labelling quality. To date, crowd-sourcing tasks in medical imaging have included classification, localization and segmentation of organs or lesions [68]. Due to the image variability and the requirement of domain knowledge, crowd-sourcing in neuroimaging, especially in glioma research, is still in its infancy, with only a few studies in this area. For example, to analyze diffuse intrinsic pontine glioma, 823 non-expert participants annotated 5,152 brain MR images to segment brain stem or tumor [71]. Rajchl et al. recruited 12 non-expert to annotate super-pixels for fetal brain segmentation. The fully convolutional network trained on crowd-sourcing annotations reached similar performance as the one

annotated by experts [72]. Keshavan et al. used crowd-sourcing to amplify expert-labelled data for quality control of brain MRI images, and a deep learning model trained on the amplified data reached AUC of 0.99 [73].

2.2. Training the model

Successful training requires the model to meet multiple desiderata such as satisfactory performance, generalizability, data privacy protection, and training with sparsely annotated data.

2.2.1. Choosing and training models

Some practical challenges are related to model training, and we list the three most important ones: choosing the optimal model, model generalization, and learning under data privacy constraint.

- **Choosing the optimal model**

The proliferation of deep learning models has led to a considerable number of layer designs and model architectures, loss functions, and optimizers to choose from when designing a network. With an infinite space of possible computation graphs, a plethora of architectures were proposed (some of the popular model architectures include AlexNet [74], VGG16 and VGG19 [75], GoogLeNet [76], ResNet [77], DenseNet [78], inception architectures [79, 80], U-Net [81], etc.). Moreover, even randomly connected graphs [82], as well as strategies to perform neural architecture search (NAS) [83, 84] have been developed in order to find the optimal model architecture. We list the model architectures used by various works in glioma imaging in Table 1.

- **Model generalization**

Models trained on images from one hospital may not be generalizable or perform equally well on new data from another hospital or image scanner, due to domain distribution shift and the confounding information models exploited [85]. Regularization strategies are utilized to improve model generalization and prevent over-fitting the training dataset. These include dropout [86], early stopping [87], data augmentation [88], or gather sufficient data from different scanners and/or hospitals [89]. To encourage model generalization among multiple scanners, Aslani et al. proposed to use an auxiliary network and corresponding regularization loss term to learn the domain-specific knowledge [89]. This auxiliary network learned to predict the category of the input scanner’s domain, thus encouraging the backbone segmentation network to ignore domain-specific information. The experiments on brain lesion segmentation from 56 different scanning sites showed the proposed method had better generalization performance to data from new sites than other baseline networks. **Transfer learning** [90] is also a common technique to deploy models at new site where data distribution may be different from the original training set. To achieve equivalent performance on new data, transfer learning

uses models pre-trained on datasets from previous hospitals to re-train the model on new data. Ghafoorian et al. conducted experiments on transfer learning on brain lesion segmentation tasks; and demonstrated that, without transfer learning, the model completely failed on new data with the same follow-up patients and when images were acquired with different acquisition protocols. With transfer learning on a small set of new training examples, the model substantially outperformed a model trained from scratch with the same size of training data [91]. The authors also observed as the new training data becomes available, fine-tuning can be done from the last full-connected layer to the last few convolutional layers.

Transfer learning, however, can suffer from catastrophic forgetting issues, where the knowledge about the old task may not be maintained when adapting parameters to a new dataset or task. To avoid this pitfall and enable **continual learning** [92], several approaches were proposed in the realm of brain segmentation. Garderen et al. applied a regularization called elastic weight consolidation (EWC) during transfer learning [93]. It penalizes large changes in model parameters weighted based on their importance to the old dataset. Research on segmenting low- and high-grade gliomas showed that EWC improved performance on the old domain after transfer learning on the new domain. Conversely, it also restricted the adaptation capacity to the new domain. Karani et al. suggested that learning batch normalization parameters for each scanner and sharing the convolutional filters between all scanners addressed the distribution shift among scanners [94]. The experiment showed this strategy can be adapted to new scanners or protocols with only a few (≈ 4) labelled images and without degrading performance on the previous scanners.

- **Learning under data privacy constraint**

Training models on data from multiple sites can improve model generalization. However, due to data privacy and regulation issues, data is usually retained within host hospital servers, and is not easily shared. Data privacy protection methods can range from data anonymization, obfuscation (modifying data with noise to provide protection, e.g.: differential privacy) [95], to federated learning (sharing model parameters rather than raw data) and learning over encrypted data [96]. **Federated learning** is a private distributed and decentralized machine learning method that trains the shared model locally with private data without exchanging the raw patient data [97, 98]. Sheller et al. demonstrated federated learning on the clinically-acquired BraTS data for institution-level segmentation tasks [99]. Their study showed that the performance of federated semantic segmentation models (Dice=0.852) was similar to that of models trained by sharing data (Dice=0.862), and outperformed two alternative collaborative learning methods.

2.2.2. Handling different levels of annotation quality

Machine and deep learning models can be trained with varying levels of supervision. Fully annotated ground truth labels for medical imaging data are expensive and




Figure 5. Different levels of supervision for a tumor segmentation task. (a) Fully supervised learning uses pixel-level annotations, whereas (d) unsupervised learning does not rely upon labels. Semi-supervised learning uses weak supervision varying from (b) bounding box annotations to (c) just image-level labels.

relatively scarce; however, images with weaker levels of annotation such as bounding box annotations or image level annotations are relatively easier to acquire. Figure 5 shows the different levels of supervision for training a tumor segmentation model. A fully supervised system would require manually annotated tumor segmentation regions (pixel-wise labeling). Weaker levels of annotation can vary from bounding box annotations of tumor region or image-level labels indicating whether a tumor is present (semi-supervised learning), or no labels at all (unsupervised learning). For detailed description of semi-supervised medical image analysis techniques, we refer the interested readers to Cheplygina et al. [100].

Conventional machine learning methods have been deployed to learn from sparsely annotated data. Azmi et al. [101] used stationary wavelet transform, edge features, and fractal features in a semi-supervised framework with an ensemble of three algorithms for segmenting brain tissues. Their approach achieved a higher accuracy and precision than fully-supervised approaches, and they noted that supervised methods such as k-nearest neighbor (KNN), SVM, and Bayesian classifiers were not able to perform well with limited data. Blessy et al. [102] proposed an unsupervised approach for MRI brain tumor segmentation based on optimal fuzzy clustering. They evaluated their approach on MRI scans of 150 patients with low grade gliomas from the BraTS dataset and outperformed four other clustering-based methods. Grande-Barreto et al. [103] used priors about the brain structure and features extracted from the 3D gray level co-occurrence matrix (GLCM) for segmenting brain tissue, and observed superior performance when evaluated on synthetic brain MR scans as compared to other state-of-the-art unsupervised approaches.

More recently, there have been several deep learning-based semi-supervised and unsupervised approaches applied to brain MR images. Chen et al. [104] trained variational autoencoder [105] and adversarial autoencoder [106] models to learn the latent space representation of healthy brains and introduced a representation consistency

constraint in the latent space, leading to a better lesion detection accuracy. Alex et al. [107] proposed using stacked denoising autoencoders (DAE) [108] for segmenting high grade gliomas, transfer learning for segmenting low grade gliomas, and a single layer DAE for constructing error maps to improve lesion detection performance. Similarly, Enguehard et al. [109] proposed a semi-supervised algorithm for infant brain MRI tissue segmentation based on deep embedded clustering, and demonstrated the robustness of their proposed approach by using the same network, optimizer, and learning rate for two additional natural image classification tasks and achieving superior performance to state-of-the-art semi-supervised methods. Dalca et al. [110] proposed a probabilistic model and a novel loss function for unsupervised segmentation of brain MRI, and built upon the recent advances in deep learning-based registration methods to jointly estimate the atlas deformation and the image intensity statistics, allowing them to train a deep model on a new dataset without any manually annotated images.

In scenarios with multiple granularities of annotated ground truth, a mixed supervision training approach can be adopted in order to utilize all the available annotations. Zheng et al. [111] proposed a multi-atlas label fusion based method to segment the hippocampus in brain MR images by integrating random forest classification models with semi-supervised label propagation. Baur et al. [112] trained an autoencoder to learn the representation of healthy brain scans, and then used it to detect and segment anomalies in unlabeled scans. A U-Net [81] was then trained for supervised segmentation with images paired with their labeled ground-truth data and unseen images paired with the autoencoder output, thereby combining unsupervised and supervised techniques. Mlynarski et al. [113] extended this architecture [81] by jointly performing classification and segmentation, and utilized weakly annotated and fully annotated data to train a deep model for tumor segmentation from brain MRI scans.

2.3. Clinical deployment

In this section, we describe approaches which address the main challenges to translating the advances in machine and deep learning into clinical deployment: namely, standardized evaluation metrics and model interpretability.

2.3.1. Standardized model evaluation

The existing literature reports the performances of AI models using a wide range of evaluating metrics, e.g., accuracy, sensitivity, specificity, area under the curve (AUC), and F_1 scores for classification tasks, overlap-based measures such as the Dice score coefficient and the intersection over union measure (also known as the Jaccard index) for segmentation tasks. Table 1 summarizes the evaluation metrics used in the literature. The metrics variability can lead to models being evaluated based on the reported best metric, making horizontal comparisons among the models a difficult task. It is therefore necessary to have a standardized set of evaluation metrics to provide a uniform platform for comparison of different machine and deep learning models. One possible approach

Table 1. List of deep learning model architectures and evaluation metrics used in glioma imaging.

Task	Model architecture	Evaluation metrics
Classification	VGG16 [114]	Accuracy [114, 115, 116, 117]
	2D [115, 114] and 3D ResNet [116]	Sensitivity or recall [114, 116, 117]
	CNN+LSTM ^a [118]	Specificity [117]
	3D CNN with gated multimodal unit [117]	Precision or PPV ^b [114, 116, 117, 118]
		F1 score [114, 116, 118] AUC ^c [116, 118]
Segmentation	2D [49, 94, 99, 113] and 3D U-Net [89, 93]	Dice score [49, 89, 91, 93, 94, 99, 113, 119, 120]
	Fully convolutional network [91]	Specificity [120]
	Densely connected CNN [119]	Sensitivity [89, 120]
	TwoPathCNN [120]	Precision [89]
		False positive rate [49, 89] False negative rate [49]
Synthesis	GAN [46, 49, 50, 51, 52, 53, 54, 121]	MSE ^d [51, 121]
		MAE ^e [46, 49]
		PSNR ^f [46, 49, 50, 51, 52, 121]
		SSIM ^g [50, 51, 52, 53, 54, 121]

^aLSTM: long short-term memory; ^bPPV: positive predictive value; ^cAUC: Area under the receiver operating characteristics; ^dMSE: mean squared error; ^eMAE: mean absolute error; ^fPSNR: peak signal-to-noise ratio; ^gSSIM: structural similarity index

to deal with this would be to report the model performance using all the conventional evaluation metrics, thereby allowing for assessing the performance of newer models with respect to the established benchmarks [122].

Moreover, high values for a reported metric of a model do not necessarily translate into superior clinical performance, mainly due to the fact that individual metrics don't capture all the desirable properties for a model [123]. For example, a viable model needs to be accurate, fast, and relatively small; invariably, there are trade-offs between model performance, generalizability, and interpretability.

2.3.2. Model interpretability

Compared to simple statistical models such as linear models, most machine and deep learning approaches suffer from the notorious problem of using opaque models (black-box models) rather than transparent ones (white-box models). Even though the model parameters and architecture are known, tracing the relationship between input-output pairs is a difficult challenge. This is partially due to the large number of model parameters; deep neural networks usually have millions of parameters (for example, the popular VGG16 architecture has over 138 million parameters). Furthermore, there are complex representations in high-dimensional space, and multi-layered non-linear mappings from input space to output predictions. In machine and deep learning literature, this challenge is referred to as the model interpretability or explainable AI

(XAI) problem, i.e., to open the black box models and reveal how the model makes the predictions in terms that human users can understand [124].

Model interpretability is especially important in deploying AI techniques in clinical settings. With a black-box model, the clinical users will only receive a prediction without an explanation or justification. Thus, it is difficult for clinical users to trust such models, adopt the model results as part of the formal and legal medical report, and learn from models to improve their expertise. XAI is therefore regarded as one of the “missing pieces” of AI in medicine [28, 125]. Although XAI is still an emerging field, there have been some recent attempts to explain AI models on glioma imaging predictive tasks. In the subsequent sections, we introduce different approaches to explaining AI decisions to clinical users, and summarize progress in applying XAI to glioma imaging tasks.

Explanations of the model’s prediction can be produced at different levels and using different approaches. **Local** and **global explanations** are explanations at different levels of granularity [126]. Local explanations explain how the model makes the decision for one data point, usually when doctors are seeking explanations for a particular clinical case, and global explanations explain how the model makes decisions in general. Similarly, **post-hoc** and **intrinsic (ante-hoc) explanations** are two distinct approaches to generate explanations [127]. Post-hoc explanation builds a proxy explanatory model to explain the original black-box model. It is especially suitable in explaining black-box models that are already put in use or those that use a specific architecture. In contrast, for models that have not been put into practice or do not have their architecture specified, we can build intrinsic (ante-hoc) explainable models that both perform predictive tasks and explain the predictions.

A prerequisite of useful XAI is that the generated explanations need to be easily understandable by clinical users. Creating explanations for clinical end-users is more challenging since they usually do not have prior knowledge in ML or AI models. Jin et al. surveyed the existing XAI literature and identified three explanatory forms that are accessible to non-technical users [128]:

- Explaining using **feature attributes** to identify the most important features that have the highest influence on the predicted outcome.
- Explaining using **examples** to show users instances that the model regards to be similar to the query image, or are typical to the prediction.
- Explaining using **decision rules** to show users the decision flows that the model follows.

All the features, examples and rules should be presented in user-understandable forms. As such, we organize the following summary of XAI advances on glioma imaging according to the aforementioned clinical-user-friendly explanatory forms.

Applying XAI methods to glioma imaging tasks is still at a nascent stage. The purposes of applying XAI in glioma imaging include model quality assessment, resolving users’ perceived anomalies, gaining users’ trust, improving physician expertise and assisting in scientific discovery (such as identifying imaging biomarkers).

(i) **Explaining using feature attributes**

Feature attribution is the most common form of explanation in the XAI literature. In an interview with clinicians about their requirements for XAI, doctors expressed the need to understand the important features that align with accepted medical paradigms [129]. Features are shown with their types of pathological changes, their importance scores, and their locations on the image. Such information can be presented in forms of text descriptions, color maps (or saliency maps) or segmentation maps overlaid unto input images. Next, we introduce three approaches to generate post-hoc explanations based on “activation”, “gradient”, and “input perturbation”.

(a) **Activation-based methods**

For deep learning-based image tasks using convolutional neural networks (CNNs), a usual approach to reveal what the network has learned is to utilize the information from the internal activation maps. For example, Ahmad et al. applied class activation map (CAM) method on 2D ResNet CNN to predict IDH mutation status in high-grade gliomas [114]. They demonstrated that the generated saliency maps correspond to the tumor area in most IDH1 mutant cases, while in most IDH1 wild-type cases, the peri-tumoral edema is also involved. The CAM method reveals where the intermediate CNN layer is “looking at” to make the prediction. It has become popular in the computer vision literature since 2015 [130]. Given an input image, it acquires the internal activation maps in a particular layer, and aggregates them using their respective weights from the global average pooling layer to generate a color map with the same size as the input image. Since the CAM method gives explanations on a single image, it is a local explanation method.

“Attention mechanism” is another approach to reveal the significant features in input data. It learns to assign “importance weights” to activation maps (in the case of CNN) or hidden states (in the case of recurrent neural network). The weighted sum of the attention map can be overlaid on input data to show how the model attends to important regions in input data. For instance, Choi et al. built an attention-based long short-term memory (LSTM) neural network to identify IDH status from dynamic susceptibility contrast (DSC) perfusion MRI time-series data [131]. The model “paid more attention” to the combination of the end of the pre-contrast baseline, up/downslopes of signal drops, and/or post-bolus plateaus for the curves used to predict IDH genotype.

(b) **Gradient-based methods**

Since gradient reflects the magnitude of output change in accordance with input change, it can provide users with clues of the important input features for the results. For example, to check if the model’s predictions align with accepted medical knowledge, Meier et al. applied Grad-CAM and guided-backpropagation on 3D MRI CNN for classification of brain tumor grades [116]. By visually inspecting some of the saliency maps that do not contain the tumor

regions, users may conclude that the prediction may not be reliable. Similarly, Huang et al. also applied 3D-based Grad-CAM on classifying high- vs. low-grade glioma to compare model performances on synthesized MR modalities with ground-truth complete modalities [121]. The resultant saliency map highlighted the tumor regions as the important image features that contribute to the tumor grading.

Both grad-CAM and guided-backpropagation methods depend on the gradient of the output with respect to the input to compute how the change of input will influence the change of output. Specifically, Grad-CAM, or gradient-weighted class activation map, can be regarded as a modification of the earlier CAM methods. Both utilize the weighted combination of activation maps to produce the saliency map; and the difference is in how they apply weighting to the activation maps. Unlike CAM that uses weights from global average pooling layer, grad-CAM takes weights from the gradients of the target class output with respect to the activation maps. Thus, it aggregates information both from activation maps and gradients, and can create a more localized saliency map than CAM [132]. Guided-backpropagation adds an additional guidance signal (gradient from top layer) to usual backpropagation which computes gradients from the bottom layer. This can create clearer and less noisy saliency maps [133].

(c) **Perturbation-based methods**

The above-mentioned methods apply to cases where we have access to model parameters, activations, and architecture information. Sometimes such knowledge is unknown to its model users, due to safety, privacy, or intellectual property issues. Perturbation-based method is used to reveal how the model makes predictions by probing the model with different inputs and analyzing the input-output pairs. It constitutes a post-hoc method and is model-agnostic (i.e., can be applied to any black-box models). In a work to segment gliomas from MRI, Pereira et al. applied the local interpretable model-agnostic explanations (LIME) [134] to the ML model for local interpretability [135]. Their ML model is a restricted Boltzmann machine for unsupervised representation learning, and a random forest classifier for voxel-level segmentation. They perturbed the features to generate synthetic neighbours that are close to the original data in the feature representation space, and also acquired the predicted output of these neighbours. With the input-output pairs, they trained a ridge regressor and used its weights as the importance-measure of each feature, shown as saliency maps on the MRI images. The saliency maps showed that FLAIR sequences were the most important for segmenting the complete tumor while segmentation of edema was mainly based on FLAIR and T1c sequence. They also showed that the T2 and T1c sequences were important for segmenting necrosis, and T1c was important in segmenting the enhancing rim region. These findings correspond well with

accepted radiologic principles; thus, confirming that the model learned the correct relations in the data.

(ii) **Explaining using examples**

Compared to feature attribution, explaining using examples can provide more contextual information about how the model learned, and is more intuitive as clinicians are used to learning from examples. In a user study with ICU clinicians, the researchers observed the example-based explanations facilitated clinicians' analogical reasoning [136]. Another study involving pathologists showed that examples with similar visual appearance and different diagnoses can help physicians broaden their differential diagnosis [137].

For example-based explanations applied to glioma imaging tasks, Chang et al. identified the prototypical images to predict IDH mutation, 1p/19q codeletion and MGMT methylation status [115]. They built CNN models on 2D MRI slices from the TCIA dataset and identified the prototypical images that caused the most activation of the units in the final fully-connected decision layer. The typical images for IDH mutation demonstrated absent or minimal enhancement with well-defined tumor margins in T1c, or central areas of cysts with FLAIR suppression. Conversely, IDH wild-type tumors exhibit thick and irregular enhancement or thin, irregular peripheral enhancement in T1c, and infiltrative patterns of edema on FLAIR. Such prototypical images serve as global explanations of the model's overall prediction.

XAI evaluating metrics To evaluate interpretability methods, Doshi-Velez and Kim proposed an evaluation framework at three levels: functionally-grounded evaluation (using function as a proxy to measure interpretability), human-grounded evaluation (evaluating with lay-person's cognition and understanding), and domain expert-grounded evaluation (evaluating the interpretability on domain-specific tasks) [124]. Current XAI techniques on glioma imaging have little to no evaluation. This is probably due to the fact that the current XAI explorations are more focused on revealing the model's predictions and checking how well the learned model aligns with doctors' prior knowledge, rather than considering the practical challenges associated with the clinical implementation of XAI. The research of clinical utility of XAI on glioma imaging is at its inception and has many unsolved questions waiting to be explored.

3. Discussion

In this manuscript, we have summarized recent efforts toward increasing the clinical utility of AI in neuro-oncology by addressing the current challenges in data collection, model training and clinical deployment. The current research, however, focuses more on novel algorithm design, rather than actually applying them in real-world patient-care settings. As the performance of machine and deep learning models improves and the challenges enumerated are addressed, we foresee greater interest in translational

research that aims to apply novel technologies to the bedside. To this end, we next highlight further advances which are necessary for the full implementation of AI in glioma imaging; to develop clinical decision support systems (CDSS) and facilitate the personalization of care.

Clinical workflow integration Implementing AI in patient-care settings requires the AI technologies to be integrated into the existing clinical workflow. Several steps are needed to achieve clinical workflow integration: First, before the design of an AI system, clinical requirements should be ascertained and a “needs assessment” should be performed. Furthermore, a user-centric design approach should be applied during the AI system design and iteration. This step solicits clinician feedback in the design of user-friendly interfaces and reduces the cognitive-load for physicians while supporting their clinical tasks. Moreover, the implemented system needs to take advantage of newly acquired patient data using online machine learning, to continually improve the performance of the predictive models. The systems can also use “adaptive learning” to adapt to user’s behaviours and preferences, i.e. “hybrid learning” or “human-in-the-loop” machine learning [138]. Achieving the above steps will need close collaboration between physicians, AI experts and human-computer interaction developers.

Clinical evaluation Appropriately evaluating the functionality of AI systems is necessary prior to its clinical implementation. We can make an analogy of evaluating the clinical utility of AI to the conventional four-phase clinical trials for medications or medical devices:




Figure 6. The four phases of evaluating the clinical utility of AI in glioma imaging.

- *Phase I* is to primarily test the performance of the AI model. This phase only involves the AI model and test data, and does not involve human users. Most current AI system evaluation is at this stage and tests AI performance on unseen data with model-based evaluation metrics, such as classification accuracy, sensitivity, specificity, AUC (area under the receiver operating characteristics),

or mean squared error, as listed in Table 1. For example, Li et al. created a clinical application called GliomaPredict that used an unsupervised ML principle component analysis (PCA) to classify patients with gliomas into 6 subtypes according to their glioma transcriptomic profiles [139]. It generated a visual representation of the analyses, quantified the confidence of the underlying subtype assessment, and presented results as a printable PDF file. The application was evaluated based on the model performance on new patient data and yielded 75-96% of prediction accuracy.

- *Phase II* involves clinical users in experimental settings using simulated tasks (usually on held-out test data) rather than in real clinical settings. The evaluation metrics are usually model-based metrics. Earlier works focused on the head-to-head comparison of doctors' performance with AI models, mainly to indicate that AI matches doctor-level performance [140, 141]. More recently, studies have involved doctor+AI as a third arm in the AI vs. doctor comparisons. For example, Bien et al. found that physicians who used AI predictions had significantly improved specificity in identifying anterior cruciate ligament tears from knee MRI compared to doctors who didn't have these predictions [142]. This evaluation trend reflects a paradigm shift from AI replacing doctors, to AI augmenting doctors.
- *Phase III* involves clinical users in real-world settings using randomized controlled trials. In addition to using model-based evaluation metrics, other clinical outcome metrics can be evaluated, such as patients' outcome, physicians' performance or efficiency, clinical usability and health economics. So far, very few studies have entered this phase. In a randomized controlled trial, Wang et al. enrolled a total of 1,058 patient to evaluate an AI-assisted polyp detection system during colonoscopy [29]. Compared with standard colonoscopy, the AI-assisted colonoscopy increased the adenoma detection rates by 50%, from 20% to 30%.
- *Phase IV* is for post-marketing software support and surveillance. This phase will follow the launch and clinical implementation of AI systems.

4. Conclusions

Because of their significant clinical burden, gliomas are one of the most highly studied cancers. Machine and deep learning techniques have the potential to increase the clinical yield of non-invasive imaging studies, from segmenting tumors for quantitative measures, to identifying molecular subtypes and grades; from evaluating treatment response to predicting patient prognosis. Despite recent advances, AI technologies have not yet become fully functional in the management of patients with gliomas. In our review, we have summarized recent efforts to address practical challenges in applying AI to clinical settings. Multiple image and label collection efforts, together with data imputation and synthesis methods, were presented. We also discussed different model training strategies, and methods of learning under limited supervision. Furthermore, we have presented various standardized model evaluation metrics and model interpretability

methods. Finally, we have proposed necessary future steps towards workflow integration and provided a framework for evaluating the clinical utility of AI-assisted systems. Due to the scope of the paper, we have focused on the technical barriers to implementation; and, approaches to overcome these challenges. Specifically, we have not discussed non-technical, financial and legal issues which will, no doubt, need to be considered in the future. Ultimately, the full implementation of AI-based tools in neuro-oncology will significantly improve the care of patients with gliomas.

Acknowledgments

Partial funding for this project is provided by the Natural Sciences and Engineering Research Council of Canada (NSERC) and Simon Fraser University Big Data The Next Big Question Fund.

References

- [1] Q. T. Ostrom, L. Bauchet, F. G. Davis, I. Deltour, J. L. Fisher, C. E. Langer, M. Pekmezci, J. A. Schwartzbaum, M. C. Turner, K. M. Walsh, M. R. Wrensch, and J. S. Barnholtz-Sloan. The epidemiology of glioma in adults: a "state of the science" review. *Neuro-Oncology*, 16(7):896–913, May 2014.
- [2] Percival Bailey and Paul C. Bucy. Oligodendrogliomas of the brain. *The Journal of Pathology and Bacteriology*, 32(4):735–751, October 1929.
- [3] Jeanette E. Eckel-Passow, Daniel H. Lachance, Annette M. Molinaro, Kyle M. Walsh, Paul A. Decker, Hugues Sicotte, Melike Pekmezci, Terri Rice, Matt L. Kosel, Ivan V. Smirnov, Gobinda Sarkar, Alissa A. Caron, Thomas M. Kollmeyer, Corinne E. Praska, Anisha R. Chada, Chandralekha Halder, Helen M. Hansen, Lucie S. McCoy, Paige M. Bracci, Roxanne Marshall, Shichun Zheng, Gerald F. Reis, Alexander R. Pico, Brian P. O'Neill, Jan C. Buckner, Caterina Giannini, Jason T. Huse, Arie Perry, Tarik Tihan, Mitchell S. Berger, Susan M. Chang, Michael D. Prados, Joseph Wiemels, John K. Wiencke, Margaret R. Wrensch, and Robert B. Jenkins. Glioma groups based on 1p/19q, IDH, and TERT promoter mutations in tumors. *New England Journal of Medicine*, 372(26):2499–2508, June 2015.
- [4] David N. Louis, Arie Perry, Guido Reifenberger, Andreas von Deimling, Dominique Figarella-Branger, Webster K. Cavenee, Hiroko Ohgaki, Otmar D. Wiestler, Paul Kleihues, and David W. Ellison. The 2016 world health organization classification of tumors of the central nervous system: a summary. *Acta Neuropathologica*, 131(6):803–820, May 2016.
- [5] Michael Weller, Ruthild G. Weber, Edith Willscher, Vera Riehmer, Bettina Hentschel, Markus Kreuz, Jörg Felsberg, Ulrike Beyer, Henry Löffler-Wirth, Kerstin Kaulich, Joachim P. Steinbach, Christian Hartmann, Dorothee Gramatzki, Johannes Schramm, Manfred Westphal, Gabriele Schackert, Matthias Simon, Tobias Martens, Jan Boström, Christian Hagel, Michael Sabel, Dietmar Krex, Jörg C. Tonn, Wolfgang Wick, Susan Noell, Uwe Schlegel, Bernhard Radlwimmer, Torsten Pietsch, Markus Loeffler, Andreas von Deimling, Hans Binder, and Guido Reifenberger. Molecular classification of diffuse cerebral WHO grade II/III gliomas using genome- and transcriptome-wide profiling improves stratification of prognostically distinct patient groups. *Acta Neuropathologica*, 129(5):679–693, May 2015.
- [6] Gregory Cairncross and Robert Jenkins. Gliomas with 1p/19q codeletion:a.k.a. oligodendroglioma. *The Cancer Journal*, 14(6):352–357, November 2008.
- [7] J. Huang, J. Yu, L. Tu, N. Huang, H. Li, and Y. Luo. Isocitrate Dehydrogenase Mutations in Glioma: From Basic Discovery to Therapeutics Development. *Front Oncol*, 9:506, 2019.

- [8] Hairui Sun, Lianhu Yin, Showwei Li, Song Han, Guangrong Song, Ning Liu, and Changxiang Yan. Prognostic significance of IDH mutation in adult low-grade gliomas: a meta-analysis. *Journal of Neuro-Oncology*, 113(2):277–284, March 2013.
- [9] Kensuke Tateishi, Hiroaki Wakimoto, and Daniel P. Cahill. IDH1 mutation and world health organization 2016 diagnostic criteria for adult diffuse gliomas: Advances in surgical strategy. *Neurosurgery*, 64(CN_suppl.1):134–138, August 2017.
- [10] Roger Stupp, Warren P. Mason, Martin J. van den Bent, Michael Weller, Barbara Fisher, Martin J.B. Taphoorn, Karl Belanger, Alba A. Brandes, Christine Marosi, Ulrich Bogdahn, Jürgen Curschmann, Robert C. Janzer, Samuel K. Ludwin, Thierry Gorlia, Anouk Allgeier, Denis Lacombe, J. Gregory Cairncross, Elizabeth Eisenhauer, and René O. Mirimanoff. Radiotherapy plus Concomitant and Adjuvant Temozolomide for Glioblastoma. *New England Journal of Medicine*, 352(10):987–996, March 2005.
- [11] Monika E. Hegi, Annie-Claire Diserens, Thierry Gorlia, Marie-France Hamou, Nicolas de Tribolet, Michael Weller, Johan M. Kros, Johannes A. Hainfellner, Warren Mason, Luigi Mariani, Jacoline E.C. Bromberg, Peter Hau, René O. Mirimanoff, J. Gregory Cairncross, Robert C. Janzer, and Roger Stupp. MGMT gene silencing and benefit from temozolomide in glioblastoma. *New England Journal of Medicine*, 352(10):997–1003, March 2005.
- [12] Prerana Jha, Vaishali Suri, Ayushi Jain, Mehar Chand Sharma, Pankaj Pathak, Pankaj Jha, Arti Srivastava, Ashish Suri, Deepak Gupta, Kunzang Chosdol, Parthoprasad Chattopadhyay, and Chitra Sarkar. O6-Methylguanine DNA Methyltransferase Gene Promoter Methylation Status in Gliomas and Its Correlation With Other Molecular Alterations: First Indian Report With Review of Challenges for Use in Customized Treatment. *Neurosurgery*, 67(6):1681–1691, December 2010.
- [13] Jan C. Buckner, Edward G. Shaw, Stephanie L. Pugh, Arnab Chakravarti, Mark R. Gilbert, Geoffrey R. Barger, Stephen Coons, Peter Ricci, Dennis Bullard, Paul D. Brown, Keith Stelzer, David Brachman, John H. Suh, Christopher J. Schultz, Jean-Paul Bahary, Barbara J. Fisher, Harold Kim, Albert D. Murtha, Erica H. Bell, Minhee Won, Minesh P. Mehta, and Walter J. Curran. Radiation plus Procarbazine, CCNU, and Vincristine in Low-Grade Glioma. *New England Journal of Medicine*, 374(14):1344–1355, April 2016.
- [14] Gavin P. Dunn, Ovidiu C. Andronesi, and Daniel P. Cahill. From genomics to the clinic: Biological and translational insights of mutant IDH1/2 in glioma. *Neurosurgical Focus*, 34(2), February 2013.
- [15] David A. Gutman, Lee A.D. Cooper, Scott N. Hwang, Chad A. Holder, Jing Jing Gao, Tarun D. Aurora, William D. Dunn, Lisa Scarpace, Tom Mikkelsen, Rajan Jain, Max Wintermark, Manal Jilwan, Prashant Raghavan, Erich Huang, Robert J. Clifford, Pattanasak Mongkolwat, Vladimir Kleper, John Freymann, Justin Kirby, Pascal O. Zinn, Carlos S. Moreno, Carl Jaffe, Rivka Colen, Daniel L. Rubin, Joel Saltz, Adam Flanders, and Daniel J. Brat. MR imaging predictors of molecular profile and survival: Multi-institutional study of the TCGA glioblastoma data set. *Radiology*, 267(2):560–569, May 2013.
- [16] Luke Macyszyn, Hamed Akbari, Jared M. Pisapia, Xiao Da, Mark Attiah, Vadim Pigrish, Yingtao Bi, Sharmistha Pal, Ramana V. Davuluri, Laura Roccograndi, Nadia Dahmane, Maria Martinez-Lage, George Biros, Ronald L. Wolf, Michel Bilello, Donald M. O’Rourke, and Christos Davatzikos. Imaging patterns predict patient survival and molecular subtype in glioblastoma via machine learning techniques. *Neuro-Oncology*, 18(3):417–425, March 2016.
- [17] Pascal O. Zinn, Sanjay K. Singh, Aikaterini Kotrotsou, Srishti Abrol, Ginu Thomas, Jennifer Mosley, Ahmed Elakkad, Islam Hassan, Ashok Kumar, and Rivka R. Colen. Distinct radiomic phenotypes define glioblastoma TP53-PTEN-EGFR mutational landscape. *Neurosurgery*, 64(CN_suppl.1):203–210, August 2017.
- [18] Saima Rathore, Spyridon Bakas, Hamed Akbari, MacLean P. Nasrallah, Stephen Bagley, and Christos Davatzikos. Abstract 1392: Machine Learning Radiomic Biomarkers Non-invasively Assess Genetic Characteristics of Glioma Patients. In *Clinical Research (Excluding Clinical*

- Trials*), pages 1392–1392. American Association for Cancer Research, July 2019.
- [19] Madeleine Shaver, Paul Kohanteb, Catherine Chiou, Michelle Bardis, Chanon Chantaduly, Daniela Bota, Christopher Filippi, Brent Weinberg, Jack Grinband, Daniel Chow, and Peter Chang. Optimizing Neuro-Oncology Imaging: A Review of Deep Learning Approaches for Glioma Imaging. *Cancers*, 11(6):829, 2019.
 - [20] P. Korfiatis and B. Erickson. Deep learning can see the unseeable: predicting molecular markers from MRI of brain gliomas. *Clinical Radiology*, 74(5):367–373, 2019.
 - [21] Jeffrey D. Rudie, Andreas M. Rauschecker, R. Nick Bryan, Christos Davatzikos, and Suyash Mohan. Emerging applications of artificial intelligence in neuro-oncology. *Radiology*, 290(3):607–618, March 2019.
 - [22] Javier E. Villanueva-Meyer, Peter Chang, Janine M. Lupo, Christopher P. Hess, Adam E. Flanders, and Marc Kohli. Machine learning in neurooncology imaging: From study request to diagnosis and treatment. *American Journal of Roentgenology*, 212(1):52–56, January 2019.
 - [23] Eyal Lotan, Rajan Jain, Narges Razavian, Girish M. Fatterpekar, and Yvonne W. Lui. State of the art: Machine learning applications in Glioma Imaging. *American Journal of Roentgenology*, 212(1):26–37, 2019.
 - [24] M Zhou, J Scott, B Chaudhury, L Hall, D Goldgof, K W Yeom, M Iv, Y Ou, J Kalpathy-Cramer, S Napel, R Gillies, O Gevaert, and R Gatenby. Radiomics in Brain Tumor: Image Assessment, Quantitative Feature Descriptors, and Machine-Learning Approaches, 2018.
 - [25] Thomas C. Booth, Matthew Williams, Aysha Luis, J. Cardosa, Ashkan Keyoumars, and Haris Shuaib. Machine learning and glioma imaging biomarkers. *Clinical Radiology*, July 2019.
 - [26] Choong Wan Woo, Luke J. Chang, Martin A. Lindquist, and Tor D. Wager. Building better biomarkers: Brain models in translational neuroimaging. *Nature Neuroscience*, 20(3):365–377, 2017.
 - [27] Jianxing He, Sally L. Baxter, Jie Xu, Jiming Xu, Xingtao Zhou, and Kang Zhang. The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, 25(1):30–36, January 2019.
 - [28] Coryandar Gilvary, Neel Madhukar, Jamal Elkhader, and Olivier Elemento. The Missing Pieces of Artificial Intelligence in Medicine. *Trends in Pharmacological Sciences*, 40(8):555–564, 2019.
 - [29] Pratik Shah, Francis Kendall, Sean Khozin, Ryan Goosen, Jianying Hu, Jason Laramie, Michael Ringel, and Nicholas Schork. Artificial intelligence and machine learning in clinical development: a translational perspective. *npj Digital Medicine*, 2(1):69, December 2019.
 - [30] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
 - [31] Alzheimer’s Disease Neuroimaging Initiative (ADNI). <http://adni.loni.usc.edu/> [Accessed: November 25, 2019].
 - [32] A Di Martino, C-G Yan, Q Li, E Denio, F X Castellanos, K Alaerts, J S Anderson, M Assaf, S Y Bookheimer, M Dapretto, B Deen, S Delmonte, I Dinstein, B Ertl-Wagner, D A Fair, L Gallagher, D P Kennedy, C L Keown, C Keysers, J E Lainhart, C Lord, B Luna, V Menon, N J Minshew, C S Monk, S Mueller, R-A Müller, M B Nebel, J T Nigg, K O’Hearn, K A Pelphrey, S J Peltier, J D Rudie, S Sunaert, M Thioux, J M Tyszka, L Q Uddin, J S Verhoeven, N Wenderoth, J L Wiggins, S H Mostofsky, and M P Milham. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular Psychiatry*, 19(6):659–667, June 2013.
 - [33] Adriana Di Martino, David O’Connor, Bosi Chen, Kaat Alaerts, Jeffrey S. Anderson, Michal Assaf, Joshua H. Balsters, Leslie Baxter, Anita Beggiato, Sylvie Bernaerts, Laura M. E. Blanken, Susan Y. Bookheimer, B. Blair Braden, Lisa Byrge, F. Xavier Castellanos, Mirella Dapretto, Richard Delorme, Damien A. Fair, Inna Fishman, Jacqueline Fitzgerald, Louise

- Gallagher, R. Joanne Jao Keehn, Daniel P. Kennedy, Janet E. Lainhart, Beatriz Luna, Stewart H. Mostofsky, Ralph-Axel Müller, Mary Beth Nebel, Joel T. Nigg, Kirsten O’Hearn, Marjorie Solomon, Roberto Toro, Chandan J. Vaidya, Nicole Wenderoth, Tonya White, R. Cameron Craddock, Catherine Lord, Bennett Leventhal, and Michael P. Milham. Enhancing studies of the connectome in autism using the autism brain imaging data exchange II. *Scientific Data*, 4(1), March 2017.
- [34] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Transactions on Medical Imaging*, 34(10):1993–2024, 2014.
- [35] Oskar Maier, Bjoern H. Menze, Janina von der Gablentz, Levin Häni, Mattias P. Heinrich, Matthias Liebrand, Stefan Winzeck, Abdul Basit, Paul Bentley, Liang Chen, Daan Christiaens, Francis Dutil, Karl Egger, Chaolu Feng, Ben Glocker, Michael Götz, Tom Haeck, Hanna-Leena Halme, Mohammad Havaei, Khan M. Iftikharuddin, Pierre-Marc Jodoin, Konstantinos Kamnitsas, Elias Kellner, Antti Korvenoja, Hugo Larochelle, Christian Ledig, Jia-Hong Lee, Frederik Maes, Qaiser Mahmood, Klaus H. Maier-Hein, Richard McKinley, John Muschelli, Chris Pal, Linmin Pei, Janaki Raman Rangarajan, Syed M.S. Reza, David Robben, Daniel Rueckert, Eero Salvi, Paul Suetens, Ching-Wei Wang, Matthias Wilms, Jan S. Kirschke, Ulrike M. Krämer, Thomas F. Münte, Peter Schramm, Roland Wiest, Heinz Handels, and Mauricio Reyes. ISLES 2015 - a public evaluation benchmark for ischemic stroke lesion segmentation from multispectral MRI. *Medical Image Analysis*, 35:250–269, January 2017.
- [36] DK Kermany and M Goldbaum. Labeled optical coherence tomography (oct) and chest x-ray images for classification. *Mendeley Data*, 2, 2018.
- [37] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn Ball, Katie Shpanskaya, Jayne Seekins, David A. Mong, Safwan S. Halabi, Jesse K. Sandberg, Ricky Jones, David B. Larson, Curtis P. Langlotz, Bhavik N. Patel, Matthew P. Lungren, and Andrew Y. Ng. CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:590–597, July 2019.
- [38] International skin imaging collaboration: Melanoma project website. <https://isic-archive.com/>.
- [39] Jure Zbontar, Florian Knoll, Anuroop Sriram, Matthew J. Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdal, Adriana Romero, Michael Rabbat, Pascal Vincent, James Pinkerton, Duo Wang, Nafissa Yakubova, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastmri: An open dataset and benchmarks for accelerated MRI. *CoRR*, abs/1811.08839, 2018.
- [40] Kenneth Clark, Bruce Vendt, Kirk Smith, John Freymann, Justin Kirby, Paul Koppel, Stephen Moore, Stanley Phillips, David Maffitt, Michael Pringle, Lawrence Tarbox, and Fred Prior. The cancer imaging archive (TCIA): Maintaining and operating a public information repository. *Journal of Digital Imaging*, 26(6):1045–1057, 2013.
- [41] Shaina Sta Cruz, Ivo D Dinov, Megan M Herting, Clio González-Zacarias, Hosung Kim, Arthur W Toga, and Farshid Sepehrband. Imputation strategy for reliable regional mri morphological measurements. *Neuroinformatics*, pages 1–12, 2019.
- [42] Amod Jog, Aaron Carass, Snehashis Roy, Dzung L Pham, and Jerry L Prince. Random forest regression for magnetic resonance image synthesis. *Medical image analysis*, 35:475–488, 2017.
- [43] Adrian V. Dalca, Katherine L. Bouman, William T. Freeman, Natalia S. Rost, Mert R. Sabuncu, and Polina Golland. Medical image imputation from image collections. *IEEE Transactions on Medical Imaging*, 38(2):504–514, February 2019.
- [44] Agisilaos Chatsias, Thomas Joyce, Mario Valerio Giuffrida, and Sotirios A Tsaftaris. Multimodal mr synthesis via modality-invariant latent representation. *IEEE Transactions on*

- Medical Imaging*, 37(3):803–814, 2017.
- [45] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
 - [46] Per Welandar, Simon Karlsson, and Anders Eklund. Generative adversarial networks for image-to-image translation on multi-contrast MR images - a comparison of CycleGAN and UNIT. *arXiv preprint arXiv:1806.07777*, 2018.
 - [47] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
 - [48] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in Neural Information Processing Systems*, pages 700–708, 2017.
 - [49] Mauricio Orbes-Arteaga, Manuel Jorge Cardoso, Lauge Sørensen, Marc Modat, Sebastien Ourselin, Mads Nielsen, and Akshay Sadananda Uppinakudru Pai. Simultaneous synthesis of FLAIR and segmentation of white matter hypointensities from T1 MRIs. In *1st Conference on Medical Imaging with Deep Learning (MIDL 2018)*, 2018.
 - [50] Salman UH Dar, Mahmut Yurt, Levent Karacan, Aykut Erdem, Erkut Erdem, and Tolga Çukur. Image synthesis in multi-contrast mri with conditional generative adversarial networks. *IEEE Transactions on Medical Imaging*, 38(10):2375–2388, October 2019.
 - [51] Anmol Sharma and Ghassan Hamarneh. Missing MRI pulse sequence synthesis using multi-modal generative adversarial network. *IEEE Transactions on Medical Imaging*, pages 1–1, 2019.
 - [52] Mahmut Yurt, Salman Ul Hassan Dar, Aykut Erdem, Erkut Erdem, and Tolga Çukur. mustGAN: Multi-stream generative adversarial networks for MR image synthesis. *arXiv preprint arXiv:1909.11504*, 2019.
 - [53] Gihyun Kwon, Chihye Han, and Dae-shik Kim. Generation of 3D brain MRI using auto-encoding generative adversarial networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 118–126. Springer, 2019.
 - [54] Dongwook Lee, Won-Jin Moon, and Jong Chul Ye. Which contrast does matter? towards a deep understanding of MR contrast using collaborative GAN. *arXiv preprint arXiv:1905.04105*, 2019.
 - [55] Changhee Han, Hideaki Hayashi, Leonardo Rundo, Ryosuke Araki, Wataru Shimoda, Shinichi Muramatsu, Yujiro Furukawa, Giancarlo Mauri, and Hideki Nakayama. Gan-based synthetic brain MR image generation. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 734–738. IEEE, 2018.
 - [56] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
 - [57] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
 - [58] Koshino Kazuhiro, Rudolf A Werner, Fujio Toriumi, Mehrbod S Javadi, Martin G Pomper, Lilja B Solnes, Franco Verde, Takahiro Higuchi, and Steven P Rowe. Generative adversarial networks for the creation of realistic artificial brain magnetic resonance images. *Tomography*, 4(4):159, 2018.
 - [59] Hoo-Chang Shin, Neil A Tenenholtz, Jameson K Rogers, Christopher G Schwarz, Matthew L Senjem, Jeffrey L Gunter, Katherine P Andriole, and Mark Michalski. Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In *International Workshop on Simulation and Synthesis in Medical Imaging*, pages 1–11. Springer, 2018.
 - [60] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.

- [61] Cheng-Bin Jin, Hakil Kim, Mingjie Liu, Wonmo Jung, Seongsu Joo, Eunsik Park, Young Saem Ahn, In Ho Han, Jae Il Lee, and Xuenan Cui. Deep ct to mr synthesis using paired and unpaired data. *Sensors*, 19(10):2361, 2019.
- [62] Soheil Damangir. Segmentation of white matter lesions – using multispectral mri and cascade of support vector machines with active learning. Master’s thesis, KTH, School of Computer Science and Communication (CSC), 2011.
- [63] Ksenia Konyushkova, Raphael Sznitman, and Pascal Fua. Introducing geometry in active learning for image segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2974–2982, 2015.
- [64] Ksenia Konyushkova, Raphael Sznitman, and Pascal Fua. Learning active learning from data. In *Advances in Neural Information Processing Systems*, pages 4225–4235, 2017.
- [65] David D. Lewis and William A. Gale. A sequential algorithm for training text classifiers. In *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’94, page 3–12, Berlin, Heidelberg, 1994. Springer-Verlag.
- [66] Wei-Ning Hsu and Hsuan-Tien Lin. Active learning by learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 2659–2665, 2015.
- [67] Jennifer Wortman Vaughan. Making better use of the crowd: How crowdsourcing can advance machine learning research. *Journal of Machine Learning Research*, 18:1–46, 2018.
- [68] Silas Ørting, Andrew Doyle, Matthias Hirth Arno van Hilten, Oana Inel, Christopher R Madan, Panagiotis Mavridis, Helen Spiers, and Veronika Cheplygina. A survey of crowdsourcing in medical image analysis. *arXiv preprint arXiv:1902.09159*, 2019.
- [69] Jane Roskams and Zoran Popović. Power to the people: Addressing big data challenges in neuroscience by creating a new cadre of citizen neuroscientists. *Neuron*, 92(3):658–664, 2016.
- [70] Vivien Marx. Neuroscience waves to the crowd. *Nature Methods*, 10(11):1069–1074, 2013.
- [71] Benjamin Timmermans, Zoltán Szilávik, and Robert-Jan Sips. Crowdsourcing ground truth data for analysing brainstem tumors in children. In *Belgium Netherlands Artificial Intelligence Conference (BNAIC)*, 2016.
- [72] Martin Rajchl, Matthew CH Lee, Franklin Schrans, Alice Davidson, Jonathan Passerat-Palmbach, Giacomo Tarroni, Amir Alansary, Ozan Oktay, Bernhard Kainz, and Daniel Rueckert. Learning under distributed weak supervision. *arXiv preprint arXiv:1606.01100*, 2016.
- [73] Anisha Keshavan, Jason D. Yeatman, and Ariel Rokem. Combining citizen science and deep learning to amplify expertise in neuroimaging. *Frontiers in Neuroinformatics*, 13, 2019.
- [74] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [75] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [76] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015.
- [77] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [78] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4700–4708, 2017.
- [79] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE*

- conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [80] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
 - [81] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
 - [82] Saining Xie, Alexander Kirillov, Ross Girshick, and Kaiming He. Exploring randomly wired neural networks for image recognition. *arXiv preprint arXiv:1904.01569*, 2019.
 - [83] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *Journal of Machine Learning Research*, 20(55):1–21, 2019.
 - [84] Martin Wistuba, Ambrish Rawat, and Tejaswini Pedapati. A survey on neural architecture search. *arXiv preprint arXiv:1905.01392*, 2019.
 - [85] John R. Zech, Marcus A. Badgeley, Manway Liu, Anthony B. Costa, Joseph J. Titano, and Eric Karl Oermann. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLOS Medicine*, 15(11):e1002683, November 2018.
 - [86] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
 - [87] Lutz Prechelt. Early stopping—but when? In *Neural Networks: Tricks of the trade*, pages 55–69. Springer, 1998.
 - [88] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
 - [89] Shahab Aslani, Vittorio Murino, Michael Dayan, Roger Tam, Diego Sona, and Ghassan Hamarneh. Scanner invariant multiple sclerosis lesion segmentation from MRI. *arXiv preprint arXiv:1910.10035*, 2019.
 - [90] Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. In *Advances in Neural Information Processing Systems*, pages 3342–3352, 2019.
 - [91] Mohsen Ghafoorian, Alireza Mehrtash, Tina Kapur, Nico Karssemeijer, Elena Marchiori, Mehran Pesteie, Charles RG Guttman, Frank-Erik de Leeuw, Clare M Tempny, Bram van Ginneken, et al. Transfer learning for domain adaptation in MRI: Application in brain lesion segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 516–524. Springer, 2017.
 - [92] Chaitanya Baweja, Ben Glocker, and Konstantinos Kamnitsas. Towards continual learning in medical imaging. *arXiv preprint arXiv:1811.02496*, 2018.
 - [93] Karin van Garderen, Sebastian van der Voort, Fatih Incekara, Marion Smits, and Stefan Klein. Towards continuous learning for glioma segmentation with elastic weight consolidation. In *International Conference on Medical Imaging with Deep Learning – Extended Abstract Track*, London, United Kingdom, 08–10 Jul 2019.
 - [94] Neerav Karani, Krishna Chaitanya, Christian Baumgartner, and Ender Konukoglu. A lifelong learning approach to brain MR segmentation across scanners and protocols. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 476–484. Springer, 2018.
 - [95] Martín Abadi, H. Brendan McMahan, Andy Chu, Ilya Mironov, Li Zhang, Ian Goodfellow, and Kunal Talwar. Deep learning with differential privacy. *Proceedings of the ACM Conference on Computer and Communications Security*, 24-28-Octo(Ccs):308–318, 2016.
 - [96] Le Trieu Phong, Yoshinori Aono, Takuya Hayashi, Lihua Wang, and Shiho Moriai. Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions*

- on *Information Forensics and Security*, 13(5):1333–1345, 2018.
- [97] Praneeth Vepakomma, Tristan Swedish, Ramesh Raskar, Otkrist Gupta, and Abhimanyu Dubey. No peek: A survey of private distributed deep learning. *arXiv preprint arXiv:1812.03288*, 2018.
 - [98] Praneeth Vepakomma, Otkrist Gupta, Tristan Swedish, and Ramesh Raskar. Split learning for health: Distributed deep learning without sharing raw patient data. In *ICLR 2019 Workshop on AI for social good*, 2018.
 - [99] Micah J Sheller, G Anthony Reina, Brandon Edwards, Jason Martin, and Spyridon Bakas. Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 92–104. Springer, 2018.
 - [100] Veronika Cheplygina, Marleen de Bruijne, and Josien PW Pluim. Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *Medical image analysis*, 54:280–296, 2019.
 - [101] Reza Azmi, Boshra Pishgoo, Narges Norozi, and Samira Yeganeh. Ensemble semi-supervised frame-work for brain magnetic resonance imaging tissue segmentation. *Journal of medical signals and sensors*, 3(2):94, 2013.
 - [102] SA Blessy and C Helen Sulochana. Performance analysis of unsupervised optimal fuzzy clustering algorithm for MRI brain tumor segmentation. *Technology and Health Care*, 23(1):23–35, 2015.
 - [103] Jonas Grande-Barreto and Pilar Gómez-Gil. Unsupervised brain tissue segmentation in MRI images. In *2018 IEEE International Autumn Meeting on Power, Electronics and Computing (ROPEC)*, pages 1–6. IEEE, 2018.
 - [104] Xiaoran Chen and Ender Konukoglu. Unsupervised detection of lesions in brain MRI using constrained adversarial auto-encoders. *arXiv preprint arXiv:1806.04972*, 2018.
 - [105] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
 - [106] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
 - [107] Varghese Alex, Kiran Vaidhya, Subramaniam Thirunavukkarasu, Chandrasekharan Kesavadas, and Ganapathy Krishnamurthi. Semisupervised learning using denoising autoencoders for brain lesion detection and segmentation. *Journal of Medical Imaging*, 4(4):041311, 2017.
 - [108] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research*, 11(December):3371–3408, 2010.
 - [109] Joseph Enguehard, Peter O’Halloran, and Ali Gholipour. Semi-supervised learning with deep embedded clustering for image classification and segmentation. *IEEE Access*, 7:11093–11104, 2019.
 - [110] Adrian V Dalca, Evan Yu, Polina Golland, Bruce Fischl, Mert R Sabuncu, and Juan Eugenio Iglesias. Unsupervised deep learning for bayesian brain MRI segmentation. *arXiv preprint arXiv:1904.11319*, 2019.
 - [111] Qiang Zheng, Yihong Wu, and Yong Fan. Integrating semi-supervised and supervised learning methods for label fusion in multi-atlas based image segmentation. *Frontiers in neuroinformatics*, 12:69, 2018.
 - [112] Christoph Baur, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. Fusing unsupervised and supervised deep learning for white matter lesion segmentation. In *International Conference on Medical Imaging with Deep Learning*, pages 63–72, 2019.
 - [113] Pawel Mlynarski, Hervé Delingette, Antonio Criminisi, and Nicholas Ayache. Deep learning with mixed supervision for brain tumor segmentation. *Journal of Medical Imaging*, 6(3):1 – 13, 2019.

- [114] Adnan Ahmad, Srinjay Sarkar, Apurva Shah, Sonal Gore, Vani Santosh, Jitender Saini, and Madhura Ingalthalikar. Predictive and discriminative localization of IDH genotype in high grade gliomas using deep convolutional neural nets. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 372–375. IEEE, 2019.
- [115] P Chang, J Grinband, B D Weinberg, M Bardis, M Khy, G Cadena, M.-Y. Su, S Cha, C G Filippi, D Bota, P Baldi, L M Poisson, R Jain, and D Chow. Deep-learning convolutional neural networks accurately classify genetic mutations in gliomas. *American Journal of Neuroradiology*, 39(7):1201–1207, 2018.
- [116] Sérgio Pereira, Raphael Meier, Victor Alves, Mauricio Reyes, Carlos A. Silva, and Carlos A Silva B. Automatic Brain Tumor Grading from MRI Data Using Convolutional Neural Networks and Quality Assessment. In *First International Workshop on Interpretability of Machine Intelligence in Medical Image Computing, iMIMIC 2018*, volume 11038, pages 106–114, 2018.
- [117] Fangyan Ye, Jian Pu, Jun Wang, Yuxin Li, and Hongyuan Zha. Glioma grading based on 3d multimodal convolutional neural network and privileged learning. In *Proceedings - 2017 IEEE International Conference on Bioinformatics and Biomedicine, BIBM 2017*, volume 2017-Janua, pages 759–763. Institute of Electrical and Electronics Engineers Inc., 2017.
- [118] Bum Sup Jang, Seung Hyuck Jeon, Il Han Kim, and In Ah Kim. Prediction of pseudoprogression versus progression using machine learning algorithm in glioblastoma. *Scientific Reports*, 8(1), 2018.
- [119] Lele Chen, Yue Wu, Adora M DSouza, Anas Z Abidin, Axel Wismüller, and Chenliang Xu. MRI tumor segmentation with densely connected 3D CNN. In *Medical Imaging 2018: Image Processing*, volume 10574, page 105741F. International Society for Optics and Photonics, 2018.
- [120] Mohammad Havaei, Axel Davy, David Warde-Farley, Antoine Biard, Aaron Courville, Yoshua Bengio, Chris Pal, Pierre Marc Jodoin, and Hugo Larochelle. Brain tumor segmentation with deep neural networks. *Medical Image Analysis*, 35:18–31, 2017.
- [121] Pu Huang, Dengwang Li, Zhicheng Jiao, Dongming Wei, Guoshi Li, Qian Wang, Han Zhang, and Dinggang Shen. CoCa-GAN: Common-Feature-Learning-Based Context-Aware Generative Adversarial Network for Glioma Grading. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, pages 155–163, Cham, 2019. Springer International Publishing.
- [122] Wei Luo, Dinh Phung, Truyen Tran, Sunil Gupta, Santu Rana, Chandan Karmakar, Alistair Shilton, John Yearwood, Nevenka Dimitrova, Tu Bao Ho, Svetha Venkatesh, and Michael Berk. Guidelines for Developing and Reporting Machine Learning Predictive Models in Biomedical Research: A Multidisciplinary View. *Journal of medical Internet research*, 18(12):e323, December 2016.
- [123] Christopher J Kelly, Alan Karthikesalingam, Mustafa Suleyman, Greg Corrado, and Dominic King. Key challenges for delivering clinical impact with artificial intelligence. *BMC medicine*, 17(1):195, 2019.
- [124] Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.
- [125] Lauren Gordon, Teodor Grantcharov, and Frank Rudzicz. Explainable Artificial Intelligence for Safe Intraoperative Decision Support. *JAMA Surgery*, pages 1–2, 2019.
- [126] Andreas Holzinger, Chris Biemann, Constantinos S Pattichis, and Douglas B Kell. What do we need to build explainable ai systems for the medical domain? *arXiv preprint arXiv:1712.09923*, 2017.
- [127] Muhammad Aurangzeb Ahmad, Carly Eckert, and Ankur Teredesai. Interpretable machine learning in healthcare. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 559–560. ACM, 2018.
- [128] Weina Jin, Sheelagh Carpendale, Ghassan Hamarneh, and Diane Gromala. Bridging AI Developers and End Users: an End-User-Centred Explainable AI Taxonomy and Visual

- Vocabularies, 2019.
- [129] Sana Tonekaboni, Shalmali Joshi, Melissa D McCradden, and Anna Goldenberg. What clinicians want: Contextualizing explainable machine learning for clinical end use. *arXiv preprint arXiv:1905.05134*, 2019.
 - [130] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2921–2929, 2016.
 - [131] Kyu Sung Choi, Seung Hong Choi, and Bumseok Jeong. Prediction of IDH genotype in gliomas with dynamic susceptibility contrast perfusion MR imaging using an explainable recurrent neural network. *Neuro-Oncology*, 21(9):1197–1209, September 2019.
 - [132] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 618–626, 2017.
 - [133] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for Simplicity: The All Convolutional Net. In *ICLR workshop*, December 2014.
 - [134] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?". In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD 2016*, pages 1135–1144. ACM Press, 2016.
 - [135] Sérgio Pereira, Raphael Meier, Richard McKinley, Roland Wiest, Victor Alves, Carlos A. Silva, and Mauricio Reyes. Enhancing interpretability of automatically extracted machine learning features: application to a RBM-random forest system on brain lesion segmentation. *Medical Image Analysis*, 44:228–244, February 2018.
 - [136] Danding Wang, Qian Yang, Ashraf Abdul, and Brian Y Lim. Designing theory-driven user-centric explainable AI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, page 601. ACM, 2019.
 - [137] Carrie J Cai, Emily Reif, Narayan Hegde, Jason Hipp, Been Kim, Daniel Smilkov, Martin Wattenberg, Fernanda Viegas, Greg S Corrado, Martin C Stumpe, and Michael Terry. Human-centered tools for coping with imperfect algorithms during medical decision-making. In *Conference on Human Factors in Computing Systems - Proceedings*, page 14. ACM, 2019.
 - [138] Andreas Holzinger. Interactive machine learning for health informatics: when do we need the human-in-the-loop? *Brain Informatics*, 3(2):119–131, 2016.
 - [139] Aiguo Li, Serdar Bozdog, Yuri Kotliarov, and Howard A. Fine. GliomaPredict: A clinically useful tool for assigning glioma patients to specific molecular subtypes. *BMC Medical Informatics and Decision Making*, 10(1), 2010.
 - [140] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115–118, January 2017.
 - [141] Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, et al. CheXnet: Radiologist-level pneumonia detection on chest X-rays with deep learning. *arXiv preprint arXiv:1711.05225*, 2017.
 - [142] Nicholas Bien, Pranav Rajpurkar, Robyn L. Ball, Jeremy Irvin, Allison Park, Erik Jones, Michael Bereket, Bhavik N. Patel, Kristen W. Yeom, Katie Shpanskaya, Safwan Halabi, Evan Zucker, Gary Fanton, Derek F. Amanatullah, Christopher F. Beaulieu, Geoffrey M. Riley, Russell J. Stewart, Francis G. Blankenberg, David B. Larson, Ricky H. Jones, Curtis P. Langlotz, Andrew Y. Ng, and Matthew P. Lungren. Deep-learning-assisted diagnosis for knee magnetic resonance imaging: Development and retrospective validation of MRNet. *PLOS Medicine*, 15(11):e1002699, November 2018.