

SAGE: Sequential Attribute Generator for Analyzing Glioblastomas using Limited Dataset

Padmaja Jonnalagedda*[§], Brent Weinberg, MD, PhD[†], Jason Allen, MD, PhD[†], Taejin L. Min, MD[†],
Shiv Bhanu, MD[‡] and Bir Bhanu, PhD*

*Dept. of Electrical and Computer Engineering, University of California at Riverside, Riverside, California

[†]Dept. of Radiology and Imaging Sciences, Emory University, Atlanta, Georgia

[‡]Dept. of Radiology, Riverside Community Hospital, Riverside, California

[§]Email: sjonn002@ucr.edu

Abstract—While deep learning approaches have shown remarkable performance in many imaging tasks, most of these methods rely on availability of large quantities of data. Medical image data, however, is scarce and fragmented. Generative Adversarial Networks (GANs) have recently been very effective in handling such datasets by generating more data. If the datasets are very small, however, GANs cannot learn the data distribution properly, resulting in less diverse or low-quality results. One such limited dataset is that for the concurrent gain of 19/20 chromosomes (19/20 co-gain), a mutation with positive prognostic value in Glioblastomas (GBM). In this paper, we detect imaging biomarkers for the mutation to streamline the extensive and invasive prognosis pipeline. Since this mutation is relatively rare, i.e. small dataset, we propose a novel generative framework the Sequential Attribute Generator (SAGE), that generates detailed tumor imaging features while learning from a limited dataset. Experiments show that not only does SAGE generate high quality tumors when compared to standard Deep Convolutional-GAN (DC-GAN) and Wasserstein GAN with Gradient Penalty (WGAN-GP), it also captures the imaging biomarkers accurately.

I. INTRODUCTION

With recent attempts towards Assistive AI and Computer-Aided Diagnosis in the medical world, a common problem encountered is the lack of curated data to train the networks on. Generative Adversarial Networks (GANs) have been shown to learn the data distribution and generate variety of data from it, which can serve as additional data for training. However, GANs also need data to train on and the need for data to create more data poses a classic causality dilemma. In the proposed research, we demonstrate a technique to generate tumor images using limited dataset, by expanding the latent representation of the tumor features. Such an approach would prove significant in studying rare mutations such as 19/20 co-gain an indicator of prognosis in brain tumors.

Clinical Background: Glioblastoma multiforme (GBM) is the most common and aggressive form of malignant tumor, comprising of 54% of all primary brain tumors [1], reporting a 5-year survival rate of 5% [2]. A significant statistic is that the survival rate has not improved in the last three decades [2], which therefore, brings to attention the need for accurate evaluation of prognosis and efficacy of chemotherapy. Assessment of overall clinical outcomes typically requires a combination of clinical, molecular and multi-modal imaging

data. This process is time consuming, invasive, cumbersome and overloads the clinical workforce. Some of the many factors that contribute to this are increasing incidence of GBM, high resolution imaging, paucity of resources for molecular testing, lack of follow-up, inconsistent data recording across modalities, etc. As an answer to this problem, researchers have started looking into optimizing the radiology pipeline by estimating prognostic markers in imaging, in an attempt to bypass the overhead caused by molecular data collection and analysis. Among the many prognostic markers [3], [4], [5], the 19/20 co-gain is yet to be explored. In this paper, we show that there are discriminatory imaging biomarkers indicating mutation and that we can recreate them using SAGE.

Motivation for SAGE: One of the major challenges in the medical analysis is lack of data. 19/20 co-gain is a rare mutation, and is therefore, a limited dataset. Furthermore, it also has a high inter-class similarity and intra-class diversity (Fig 1), which makes the visual assessment of biomarkers very unapparent to the naked eye. Any analysis using GANs will be successful only if the GAN is capable of generating synthetic images that capture the nuance and diversity of features in the mutation. This implies the existence of a data distribution that GANs can learn from. Additionally, since 19/20 co-gain has never before been analyzed for visual manifestation in MR Imaging (MRI), we begin by demonstrating that there is, indeed, a presence of consistent imaging biomarkers that have a positive correlation with presence or absence of 19/20 co-gain. Following this, we will proceed to learn and recreate these biomarkers using SAGE. The motivation behind developing SAGE as opposed to using DC-GAN or WGAN is that SAGE is designed to increase the apparent size of the dataset. In other words, to mitigate the issue of lack of data samples to use in GAN, SAGE increases the size of dataset in latent space by deconstructing features and offers a wider search space of feature combinations (i.e., data distributions). We show that SAGE learns useful features from the dataset to generates good quality and diverse synthetic images.

The rest of the paper is organized as follows: Section 2 contains related work and contributions. Section 3 details the technical approach. Section 4 contains the description of dataset and evaluation setup. Section 5 presents the results and Section 6 is conclusions.

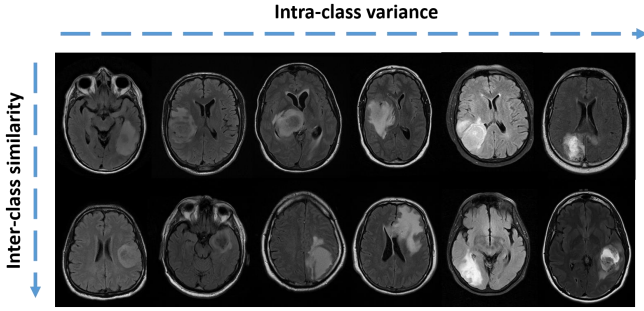


Fig. 1: Inter-class variability and inter-class similarity in 19/20 co-gain mutation. Images are all shown in FLAIR (fluid-attenuated inversion recovery) MR modality. Top: No mutation, Bottom: Mutation present

II. RELATED WORK AND CONTRIBUTIONS

A. Related Work

Deep learning has recently achieved remarkable success in various aspects of automated workflow involving MR Image Analysis such as diagnosis [6], grading [7], segmentation [8] and other clinical outcomes [9]. Generative Adversarial Networks (GANs) have recently gained attention in the medical field due to their efficacy in modeling and recreating data distributions to tackle common issues with medical datasets [10], [11], [12], the primary goal being data augmentation. For MRI, GANs have been used for data augmentation [10], segmentation [13], data anonymization [14], etc. Most of these GANs generate the whole image at once. In our case, this approach does not necessarily generate the most detailed images. Since we are analyzing unapparent visual manifestations, it is critical that the tumor features be generated with sufficient detail. For High- and Low-Grade Gliomas, there has been some research in generating images for data augmentation [15], [10]. Some of the research has attempted to use manifold learning [16], attention-based learning [17], VAEs [18] etc. for improving learning quality. Real-world medical data, however, resides in a higher dimension with many complex attributes operating together. Adding to this complexity, medical data is often fragmented and scarce. With SAGE, we establish the impact of separating features and latent space re-sampling to increase the apparent sample space of the dataset to match the complexity and dimensionality of the data. We propose that this can result in significantly better-quality images.

B. Contributions

- A novel approach that can generate diverse synthetic images from very limited datasets using feature recasting,
- Feature disentanglement and sequential generation for higher-resolution images and added control over generated tumor properties,
- Quantitative analysis of efficacy of proposed method in learning and recreating visually unapparent data distribution compared to naive GANs

III. TECHNICAL APPROACH

A. Overview of Approach

The Sequential Attribute Generator (SAGE) framework has three modules: (a) Shape Generation module, (b) Texture Generation module and (c) Merge module. Essentially, SAGE generates tumor crops using (a) and (b) and then merges the tumors with tumor-less brain slices in (c). The detailed explanation of each of these modules is given in Section III-B. From the overall dataset, we separate the slices of MR Images that contain tumor vs the ones that don't. Tumors are segmented manually from the slices that contain them to create Tumor Crops (TC). Developing deep learning approaches for brain tumor segmentation is an actively growing field. Some of the best performing networks [19] will also have some errors. If a tumor is wrongly segmented, the error will be propagated through the other modules of SAGE, generating incorrect tumors. We, therefore, opted for manual segmentation to obtain TCs. The TCs are then subjected to simple binary thresholding to obtain the shape masks. These will be referred to as Binary Tumor Crops (BTCs) henceforth. The slices that do not contain tumors are used in the merge block input (Tumor-less Images in Figure 2) and are called Pseudo-healthy Images (PHI). When a tumor grows in the brain, it pushes through the healthy brain tissue. Since there is yet no mathematical model that emulates this impact, we have used PHIs. PHIs were a part of the brain that contained a tumor and contain residual impact: thereby, the closest model we can use for merging with the tumors realistically. The properties of tumors and PHIs correspond to where the tumor can be located in the brain. BTCs are a representation of shape of tumors whereas TCs have the texture properties. So, we now separate and re-sample from three macro-features: shape, texture and location. Recasting and re-sampling of these features gives us an apparent increase in sample space since we now have three data distributions instead of one. Figure 2 outlines the overall SAGE approach. The following sections provide details on the SAGE workflow.

B. Description of Framework

In this section, we will detail each of the blocks of SAGE as described in Section III-A. The shape generation module of SAGE consists of a Latent-space Multiplexed Sampling GAN (LMS-GAN). The texture generation module consists of a Texture Alignment Network (TAN). The overall approach is shown in Figure 2.

LMS-GAN: The Latent-space Multiplexed Sampling GAN (LMS-GAN) is the shape generator module of SAGE. It is a densely connected GAN framework that multiplexes between latent space vectors of data and noise to generate a binary mask of the tumor. The input of LMS-GAN are the BTCs. LMS-GAN performs a non-linear mapping from the input image or noise to the latent space. From the latent space, the network randomly samples data points and gives it as an input to the generator, periodically switching the latent vector between noise (Z_{noise}) and real data (Z_{real}). By randomly

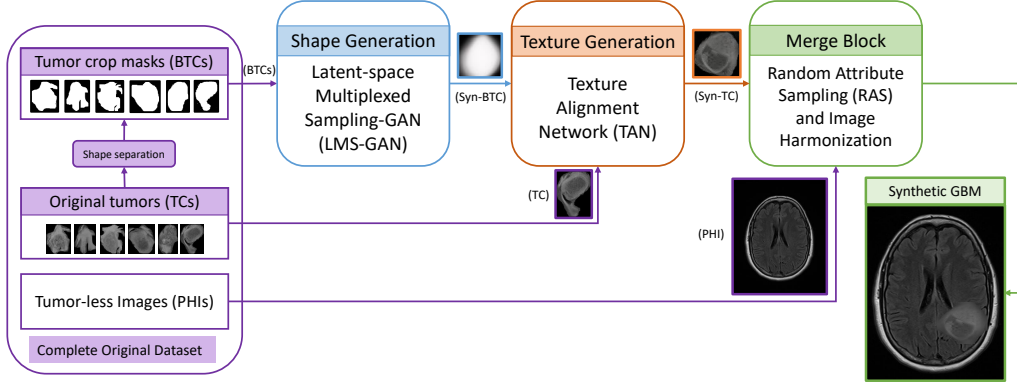


Fig. 2: Overall framework of the Sequential Attribute Generator (SAGE). BTC, TC and PHIs are the isolated features of real data which are re-sampled in the Shape Generation, Tumor Generation and Merge (Location generation) blocks respectively.

sampling points from the latent space, it can generate as many variations of an input vector to the generator as needed. However, since the latent vector input to the generator is no longer a feature vector from a coherent image, the generator block is designed as a densely connected network and not a convolutional network. To train the generator, the loss function switches with the input data type. If the sampling is done from Z_{real} , the network is trained on Mean Squared Error (MSE) loss (Eq 1a) and if the samples are Z_{noise} , loss is Binary Cross-Entropy (BCE) loss (Eq 1b). The overall loss function is shown in (Eq 2).

$$\mathcal{L}_{X \sim p(real)} : \mathbb{E}_{x \sim p_{real}(x)} \|x_{real} - y\|_2 \quad (1a)$$

$$\mathcal{L}_{X \sim p(noise)} : \mathbb{E}_{x \sim p_{noise}(x)} \left[\mathfrak{b} \left(1 - \mathbf{D}(\mathbf{G}(z_{noise})) \right) \right] \quad (1b)$$

$$\mathcal{L}_{LMS} = \mathbb{I}_x \times \mathcal{L}_{X \sim p(noise)} + (1 - \mathbb{I}_x) \times \mathcal{L}_{X \sim p(real)} \quad (2)$$

where \mathbb{I}_x is the indicator function which is 0 when X is data and 1 when X is noise. The noise is sampled from random normal distribution. \mathbf{G} and \mathbf{D} are Generator and Discriminator, respectively. y is the generated image and x_{real} is the real image. $Z_{real}/noise$ are the latent space representations of real and noisy inputs. This is depicted in Fig 3.

Texture Alignment Network (TAN): This module generates texture. TAN takes the Synthetic-BTCs (Syn-BTCs) generated by LMS-GAN and assigns a texture. TAN randomly samples a texture from the pool of TCs, learns its feature

representation and aligns the distribution of Syn-TC with that of the sampled TC distribution. The output of this block is a synthetic tumor crop (Syn-TC). Some researchers have performed feature reassignment [20], [21] in the past. In our approach, we use features from Conv3, 4, 5 blocks of VGG19 and train on perceptual loss (Eq 3). We only include the deeper layers because texture of an image is learnt in the deeper layers of a CNN.

$$\mathcal{L}_{TAN} : \sum_i \omega_i \times \|g_i(T) - g_i(I)\|_2 + \|\mathbb{F}_i(B) - \mathbb{F}_i(I)\|_2 \quad (3)$$

where: \mathcal{L}_{TAN} : overall loss function, ω_i : weight of i^{th} layer, T: sampled TC, B: input Syn-BTC, I: output Syn-TC, \mathbb{F}_i output of i^{th} layer and g_i : Gram matrix of i^{th} layer output.

Merge Block: When we merge the Syn-TCs with PHIs, we must ensure maximum diversity and uniformity to generate realistic images. For diversity, we randomly select location, relative size of tumor (with respect to whole image) and rotation angle for the Syn-TC. The values of these attributes are chosen from a pre-defined allowed range of values. This range is obtained from the mean and standard deviation of the original dataset attributes so we can choose attribute properties with a certain confidence. Once the tumor is merged, we apply edge-tapering Gaussian filters for smoothness. The overall framework of SAGE allows us to address the issues of image quality because of data scarcity commonly experienced by GAN approaches. In the following section, the results of using SAGE framework to generate GBM images are detailed.




Fig. 3: Framework of Latent-space Multiplexed Sampling (LMS) GAN

IV. EXPERIMENTAL SETUP

A. Description of Dataset

We use a dataset containing FLAIR MR images for a cohort of 25 patients with known 19/20 co-gain status. The data is divided into two classes: control class that has 19/20 co-gain absent and the mutated class that has co-gain present. This dataset was provided by our collaborators at Emory University. The cohort is divided into 14 control and 11 mutated patients. Each patient has an average of only 9 FLAIR images with tumor, making it an average of 110 images per class (without

train-test split), this makes for an even smaller training set. Thus, this is a very limited dataset to train GANs.

B. Evaluation Protocol

1. *Qualitative Evaluation:* For qualitative evaluation, we present three types of results, each of which demonstrates a specific strength of SAGE. One result evaluates how well SAGE can learn the inherent data distribution of the data and translate it into generating new tumors. This is shown by comparing SAGE generated images with real images as well as DC-GAN and WGAN-GP results. We chose four real tumor crops with very different visual properties for each class and picked corresponding generated images from all three GANs that most closely resembled these tumors. This provides evidence of how well SAGE can learn the properties of highly diverse and limited data. For the second type of results, we demonstrated how SAGE can generate a diverse range of synthetic tumors compared to other GANs. This is an evaluation of the robustness of SAGE towards mode-collapse and loss of detail in generated images. For the third type of results, we show how well SAGE can capture the quality and detail of tumors. Thus, we zoom in into an example set of images for each class and compare resolution and detail between real, SAGE, DC-GAN and WGAN-GP.

2. *Quantitative Evaluation:* Inception Score and Structural Similarity: We report Inception Score (IS) and Structural Similarity Index (SSIM) to demonstrate the quality of generated images. The Inception Score evaluates two properties: (a) Image quality and (b) Image diversity. To assess image quality, meaningful generated objects should have a conditional label distribution with low entropy [22]. For ensuring diversity in generated images, the marginal probability should have high entropy [22]. Inception score is a combination of these two constraints. It achieves this by calculating the KL Divergence between the conditional and marginal distributions. The average KL divergence value of all generated images is the score. The lowest value of IS is 1 and the highest value is the number of classes. The SSIM is a comprehensive measure of similarity between images. The value of SSIM lies between 0 and 1, where 0 means the images aren't similar at all. Ideally, we want our images to follow the same distribution as real data but also be different enough to ensure diversity. Therefore, we can identify the range of 0.35-0.75 as the preferred range of SSIM to evaluate our synthetic image quality.

3. *Blind Test by Radiologists:* We conducted tests for assessing the quality of generated tumor crops via a blind prediction test with radiologists. We asked radiologists to manually distinguish between real and synthetic tumors. The evaluation was done with 3 radiologists. They were provided with control and mutated folders, each containing random number of real and synthetic images. These images were randomly chosen and resized to 100x100 for consistency. Furthermore, the names of these images were replaced by random numbers to hide any indicators that may point towards the nature of the image. For a completely blind assessment, the statistics of data split were not revealed to the radiologists.

The control group had a total of 92 images: 45 real and 47 synthetic while the mutated group had 82 images: 35 real and 47 synthetic images. The radiologists were then asked to mark the images R/F (Real or Fake). These results are summarized in Table II. Here, accuracy (ACC) is the proportion of correct predictions made, False Positive Rate (FPR) is the proportion of synthetic images that were mistaken as real, False Negative Rate (FNR) is the proportion of real images mistaken as synthetic and Precision (PR) is the proportion of true positives over all positive calls. During our design of experiment for this test, there were two key questions which we answered. These questions and answers are given below:

Question 1: What is your hypothesis that you are trying to show by having us classify them? Do you predict that a human will or will not be able to tell the difference?

Answer 1: We want to have an expert evaluation of the quality of generated tumor images. If the generated images are of high quality, one will not be able to distinguish them from real images. Based on this we can have a metric of the quality of generated images. This will also help us to improve the quality of generated images in the future.


Question 2: What is the reasoning for the tumor being so tightly cropped? For us, it is very unusual to look at an image that is only the segmented tumor. In this way, all the images seem very strange to us. It is difficult to reliably tell the difference because there is not much context.

Answer 2: The images are tightly cropped since only the tumor is (synthetically) generated. Cropped images are used for data augmentation for the training of deep learning algorithms (to distinguish mutated/control groups) assuming the detection/segmentation has been done perfectly. We aim to evaluate how well the tumor features are learnt when isolated. Yes, the context is very important for image interpretation and currently we take that into account in the classification step where we learn the discriminatory features.


4. *Learning Discriminatory Features:* We focus our evaluation on our primary goal: to evaluate whether SAGE can consistently and precisely learn the desired feature distribution. To evaluate this, we use the generated images to classify real images. The rationale is: if SAGE-generated images can detect discriminatory features in a real test set, then SAGE has captured the desired data distribution unique to each class. We compare these results using synthetic images generated using DC-GAN and WGAN-GP to demonstrate the effectiveness of SAGE in learning from limited data. We report accuracy, sensitivity, specificity and dice score for a complete evaluation of performance. Here accuracy is the proportion of total correct predictions, sensitivity is the true positive rate, specificity is the true negative rate and dice score is the harmonic mean of precision (proportion of true positives over all positive calls) and recall (sensitivity).

V. RESULTS AND DISCUSSION

We perform evaluations for qualitative and quantitative analysis of generated tumor crops and the results of blind test by radiologists. The results are summarized below:



(a) Control class



(b) Mutated class

Fig. 4: Tumor crops and Whole Images generated by SAGE vs Real data and compared against DC-GAN and WGAN-GP. The green (solid) box contains real samples, the red (dotted) box contains samples generated by SAGE, blue (dashes-medium) contain samples generated by DC-GAN and yellow (dashes-large) contain WGAN-GP samples.

1. *Qualitative Evaluation of Generated Images:* We evaluate the ability of SAGE to learn from the data distribution accurately. To show that SAGE can learn the features compared to other GANs, the tumor crops and whole slide images generated by SAGE, DC-GAN [23] and WGAN-GP [24] are shown in Figure 4. We use WGAN-GP over standard WGAN because the gradient penalty has been shown to provide superior results [24]. SAGE generates attributes sequentially, giving us complete control over how and what kind of images we generate. It can be seen from Figure 4 that we are able to recreate specific tumor attributes. For comparison, we cherry-picked the DC-GAN and WGAN-GP images that looked

visually similar to the corresponding real and SAGE image. It can be seen that SAGE is able to recreate the detailed attributes more faithfully.

Our goal, however, is not limited to recreating features exactly. We also want to generate good quality images and a diverse set of synthetic images that learn from the real distribution. Figure 6 shows the diversity of images generated using each of the GANs. It can be seen that DC-GAN suffers from mode-collapse due to lack of data. WGAN-GP is relatively better at handling it, however, given the limited data available, it cannot combat mode-collapse and generate detailed attributes at the same time. Therefore, we can see




Fig. 5: Comparing quality of similar-looking generated images. The purple dashed line divides control and mutated samples. The green (solid) box contains real samples, the red (dotted) box contains samples generated by SAGE, blue (dashes-medium) contain samples generated by DC-GAN and yellow (dashes-large) contain WGAN-GP samples.

that while WGAN-GP generates diverse images compared to DC-GAN, it is unable to achieve the same quality of tumors as SAGE. SAGE not only generates better quality images, but also generates more diverse set of synthetic images. Additionally, Figure 5, zooms in on some examples of generated images where DC-GAN and WGAN-GP fail to generate an acceptable detail and quality of tumor image.

2. *Quantitative Evaluation*: As a quantitative measure, we calculated the Inception Score of generated images. For two classes, the score can be between 1 and 2. The closer the score is to 2, the better our GAN performs. For SAGE, we calculated Inception score for a total of 1000 images (500 each of control and mutated). We calculated SSIM over the same set. We used similar settings to evaluate Inception score and SSIM for DC-GAN and WGAN-GP. The results are shown in Table I. We note that SAGE gives an inception score value of 1.71 which shows that it is effective in learning both quality and diversity of data. The SSIM value above 0.5 indicates that visual properties of generated images are similar to real images but since the SSIM is not too close to 1, we can infer that there is a certain level of diversity associated with data.

3. *Evaluation by Radiologists*: We also report results from a blind test conducted with our radiologists (Table II) to demonstrate that SAGE can generate realistic looking good

TABLE I: Inception score (IS) and Structural Similarity (SSIM) for generated images using SAGE, DC-GAN and WGAN-GP. The number of generated images per class is 500 for each GAN.

Generative Model	IS	SSIM
SAGE (ours)	1.71	0.68
WGAN-GP	1.35	0.57
DC-GAN	1.12	0.32

TABLE II: Blind test for radiologists to distinguish between real and synthetic tumor crops. The metrics shown are: ACC - Accuracy, FPR - False Positive Rate, TNR - True Negative Rate, PR - Precision. Row (non-shaded) are values for each radiologist and row (shaded) is the mean value across all radiologists.

Data class	Radiologist	ACC	FPR	FNR	PR
Control	1	0.59	0.42	0.38	0.55
	2	0.67	0.34	0.31	0.66
	3	0.70	0.30	0.36	0.64
	Mean	0.66	0.35	0.33	0.64
Mutated	1	0.74	0.33	0.19	0.72
	2	0.82	0.19	0.18	0.87
	3	0.76	0.30	0.20	0.77
	Mean	0.77	0.27	0.19	0.79

quality tumors. Referring to Table II, we note a FPR of 35% and 27% for control and mutated classes, respectively. We also observe a FNR of 33% and 19%, respectively. This indicates that SAGE is also able to generate realistic new tumors, i.e., learning how a new tumor would realistically grow. For the tumors that the radiologists could detect as being synthetically generated, we obtained feedback from the radiologists about what features look artificial. There was an agreement among radiologists that an image was computer generated because one of two things: (a) a specific artifact which seemed to be present on some of the images. Many of them had sharp lines or a broken appearance, discontinuity. Those images were suspected as computer generated. They also classified "smudgy" appearance as computer generated or looking fake, and (b) the brain cortex seemed discontinuous or abnormal on some of the images, which was labeled as computer generated. Further, they did not find anything about the tumors themselves that was different between real and generated images.




Fig. 6: Diversity of images generated by SAGE vs WGAN-GP vs DC-GAN for both classes of tumors. The purple dashed line divides control and mutated samples. The green (solid) box contains real samples, the red (dotted) box contains samples generated by SAGE, blue (dashes-medium) contain samples generated by DC-GAN and yellow (dashes-large) contain WGAN-GP samples.

TABLE III: Mutation detection using only synthetic data generated using SAGE, DC-GAN and WGAN-GP. Data is generated using the training set of shaded row and tested on the corresponding real test set for all 10-folds. The results reported are mean (standard deviation) over 10 folds. Row: cyan is baseline trained on real data. Metrics are - ACC: Accuracy, SEN: Sensitivity, SPEC: Specificity and DIC: Dice Score for IL: Image Level and PL: Patient Level analyses.

Input Type	Input source	ACC (PL)	ACC (IL)	SEN (IL)	SPEC (IL)	DIC (IL)
Tumor Crops	Real Data	0.85 (0.12)	0.80 (0.09)	0.69 (0.09)	0.84 (0.06)	0.70 (0.07)
	SAGE (ours)	0.85 (0.10)	0.81 (0.07)	0.72 (0.07)	0.87 (0.06)	0.72 (0.08)
	WGAN-GP	0.78 (0.08)	0.76 (0.09)	0.70 (0.12)	0.84 (0.06)	0.72 (0.09)
	DC-GAN	0.74 (0.09)	0.71 (0.10)	0.68 (0.12)	0.82 (0.08)	0.68 (0.10)
Whole Images	Real Data	0.92 (0.08)	0.90 (0.06)	0.85 (0.07)	0.95 (0.05)	0.86 (0.08)
	SAGE (ours)	0.92 (0.09)	0.89 (0.07)	0.84 (0.08)	0.94 (0.04)	0.86 (0.08)
	WGAN-GP	0.85 (0.06)	0.80 (0.07)	0.76 (0.09)	0.88 (0.05)	0.75 (0.08)
	DC-GAN	0.80 (0.08)	0.78 (0.10)	0.74 (0.12)	0.84 (0.06)	0.77 (0.12)

4. *Learning Discriminatory Features*: This test is to demonstrate that SAGE can consistently learn unapparent yet discriminative visual properties of tumors between presence and absence of mutation. Table III (cyan) reports the evidence of existence of discriminative features between the two classes (consistent high accuracy in mutation detection). Once we establish that mutation can be detected using the classifier, our objective is to evaluate how well SAGE is able to learn from a limited sample space. If synthetic images learn features accurately, then they should be equally effective in detecting mutation. Thus, we replace the training set with synthetic images generated using SAGE. We get similar performance in mutation detection using the synthetic training set as we get using real set. For the experiments in Table III, the dataset was divided patient-wise via a 80-20 split for 10 folds. We train ResNet18 for classification and report mean and standard deviation. We describe two types of results: patient-level (PL) classification and image-level (IL) classification. IL results are the results by considering each slice independently.

PL results are obtained by computing the weighted mean of all images of a patient. The rest of Table III reports the results obtained using a similar training protocol, but trained on various training sets. The three training sets are: generated images using DC-GAN, WGAN-GP and SAGE, respectively. The generated images for 10-folds are learnt from the training sets of the corresponding 10-fold real data and evaluated on corresponding test set. We evaluate whether the generative model can consistently and accurately capture the required distribution to distinguish between control and mutated groups (via a 10-fold cross-validation). The test set for every corresponding fold was kept constant across all three experiments (SAGE, DC-GAN and WGAN-GP) for fair comparison with baseline (cyan, real data). For training we used pre-trained ResNet18 using Adam optimizer, learning rate of 0.0001 with step learning rate scheduler and binary cross-entropy loss. The visual discriminatory features between classes are unapparent and yet classification using SAGE images consistently performs at par with real data. Due to

this, we observe that SAGE is able to learn and recreate subtle indicators of tumors that discern the status of mutation.

VI. CONCLUSIONS

Medical image generation is a challenging task due to the restrictive nature of data, diversity in individual cases and subsequently complex feature space. In this work, We have shown that synthetic images generated by SAGE match very closely with real data, thereby indicating that SAGE is able to accurately capture data distribution from limited dataset. Comparing with DC-GAN and WGAN-GP, we can see that SAGE has superior performance in terms of capturing diversity and detail, with an Inception score of 1.71. Upon closer inspection of Figure 5, we can see that where other GANs become blurry or lose features, SAGE generates detailed tumors. Furthermore, due to the sequential generation of features, we have complete control over the tumors we generate. Therefore, the situations where other GANs overfit and generate the same tumors repeatedly can be avoided in SAGE. The classification tasks indicates that the classifier trained on SAGE images performs just as well as real images, indicating the discriminatory features correlating to presence or absence of mutation are learnt accurately. The quantitative and qualitative evaluation showed that SAGE is an effective approach for generating synthetic data using limited datasets. In addition to these analyses, the visual quality and realism of the images was tested via the blind test. From the blind test it can be concluded that SAGE generates realistic images and not just recreate images from random re-sampling of features. We notice that SAGE has an improvement of 7-12% as compared to other methods in learning features and a consistently high performance. Summarizing these results, we note that SAGE learns the features and statistical model accurately. We can, therefore, conclude that SAGE can potentially be applied to tackle many disease and anomaly detection problems with limited and/or fragmented data.

REFERENCES

- [1] Kaja Urbańska, Justyna Sokółowska, Maciej Szmidt, and Paweł Sysa, "Glioblastoma multiforme—an overview," *Contemporary oncology*, vol. 18, no. 5, pp. 307, 2014.
- [2] Ahmad Faleh Tamimi and Malik Juweid, "Epidemiology and outcome of glioblastoma," in *Glioblastoma [Internet]*. Codon Publications, 2017.
- [3] 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, et al., "Mgmt gene silencing and benefit from temozolomide in glioblastoma," *New England Journal of Medicine*, vol. 352, no. 10, pp. 997–1003, 2005.
- [4] Christian Hartmann, Bettina Hentschel, Matthias Simon, Manfred Westphal, Gabriele Schackert, Jörg C Tonn, Markus Loeffler, Guido Reifenberger, Torsten Pietsch, Andreas von Deimling, et al., "Long-term survival in primary glioblastoma with versus without isocitrate dehydrogenase mutations," *Clinical Cancer Research*, vol. 19, no. 18, pp. 5146–5157, 2013.
- [5] Christoph Geisenberger, Andreas Mock, Rolf Warta, Carmen Rapp, Christian Schwager, Andrey Korshunov, Ann-Katrin Nied, David Capper, Benedikt Brors, Christine Jungk, et al., "Molecular profiling of long-term survivors identifies a subgroup of glioblastoma characterized by chromosome 19/20 co-gain," *Acta neuropathologica*, vol. 130, no. 3, pp. 419–434, 2015.
- [6] Geert Litjens, Clara I Sánchez, Nadya Timofeeva, Meyke Hermesen, Iris Nagtegaal, Iringo Kovacs, Christina Hulsbergen-Van De Kaa, Peter Bult, Bram Van Ginneken, and Jeroen Van Der Laak, "Deep learning as a tool for increased accuracy and efficiency of histopathological diagnosis," *Scientific reports*, vol. 6, pp. 26286, 2016.
- [7] Fusun Citak-Er, Zeynep Firat, İlhami Kovanlıkaya, Ugur Ture, and Esin Ozturk-Isik, "Machine-learning in grading of gliomas based on multi-parametric magnetic resonance imaging at 3t," *Computers in biology and medicine*, vol. 99, pp. 154–160, 2018.
- [8] Sérgio Pereira, Adriano Pinto, Victor Alves, and Carlos A Silva, "Brain tumor segmentation using convolutional neural networks in mri images," *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1240–1251, 2016.
- [9] P Chang, J Grinband, BD Weinberg, M Bardis, M Khy, G Cadena, M-Y Su, S Cha, CG Filippi, D Bota, et al., "Deep-learning convolutional neural networks accurately classify genetic mutations in gliomas," *American Journal of Neuroradiology*, vol. 39, no. 7, pp. 1201–1207, 2018.
- [10] 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)*. IEEE, 2018, pp. 734–738.
- [11] Maayan Frid-Adar, Idit Diamant, Eyal Klang, Michal Amitai, Jacob Goldberger, and Hayit Greenspan, "Gan-based synthetic medical image augmentation for increased cnn performance in liver lesion classification," *Neurocomputing*, vol. 321, pp. 321–331, 2018.
- [12] Maayan Frid-Adar, Eyal Klang, Michal Amitai, Jacob Goldberger, and Hayit Greenspan, "Synthetic data augmentation using gan for improved liver lesion classification," in *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*. IEEE, 2018, pp. 289–293.
- [13] Chaoyi Zhang, Yang Song, Sidong Liu, Scott Lill, Chenyu Wang, Zihao Tang, Yuyi You, Yang Gao, Alexander Klistorner, Michael Barnett, et al., "Ms-gan: Gan-based semantic segmentation of multiple sclerosis lesions in brain magnetic resonance imaging," in *2018 Digital Image Computing: Techniques and Applications (DICTA)*. IEEE, 2018, pp. 1–8.
- [14] 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*. Springer, 2018, pp. 1–11.
- [15] Changhee Han, Kohei Muroo, Tomoyuki Noguchi, Yusuke Kawata, Fumiya Uchiyama, Leonardo Rundo, Hideki Nakayama, and Shin'ichi Satoh, "Learning more with less: Conditional pggan-based data augmentation for brain metastases detection using highly-rough annotation on mr images," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2019, pp. 119–127.
- [16] Mahyar Khayatkhoei, Maneesh K Singh, and Ahmed Elgammal, "Disconnected manifold learning for generative adversarial networks," in *Advances in Neural Information Processing Systems*, 2018, pp. 7343–7353.
- [17] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena, "Self-attention generative adversarial networks," *arXiv preprint arXiv:1805.08318*, 2018.
- [18] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner, "beta-vae: Learning basic visual concepts with a constrained variational framework.," *Iclr*, vol. 2, no. 5, pp. 6, 2017.
- [19] Andriy Myronenko, "3d mri brain tumor segmentation using auto-encoder regularization," in *International MICCAI Brainlesion Workshop*. Springer, 2018, pp. 311–320.
- [20] Leon A Gatys, Alexander S Ecker, and Matthias Bethge, "A neural algorithm of artistic style," *arXiv preprint arXiv:1508.06576*, 2015.
- [21] Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, Yizhou Yu, and Mingli Song, "Neural style transfer: A review," *IEEE transactions on visualization and computer graphics*, 2019.
- [22] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen, "Improved techniques for training gans," in *Advances in neural information processing systems*, 2016, pp. 2234–2242.
- [23] Alec Radford, Luke Metz, and Soumith Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.

- [24] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville, “Improved training of wasserstein gans,” in *Advances in neural information processing systems*, 2017, pp. 5767–5777.