RESEARCH PAPER

ENHANCING BRAIN TUMOR CLASSIFICATION FROM MRI USING GENERATIVE DATA AUGMENTATION

[1] Kolli.Bala Bhaskara Rao, [2] Dr. R.VIJAYA KUMAR REDDY
[1, 2] Department of CSE, Koneru Lakshmaiah Foundation, Vaddeswaram, AndhraPradesh, India
[1] baskarbala165@gmail.com, [2] rvijayakumarreddy@kluniversity.in

Abstract
In medical diagnostics, DL based models often fail to perform well because of the high level of class imbalance in the training set especially if one is dealing with a rare pathology. Such inner data bias can result in biased classifiers and sub-optimal generalisation. In this article, we propose a data-generative approach of data-augmentation to address this constraint in the case of brain-tumor detection utilising MRI. A conditional GAN was built and trained to produce high fidelity T1-weighted images with contrast agent. Using the 3,064 photos in three categories of the Figshare Brain Tumour Dataset, the network was able to learn the morphological traits unique to different types of tumours. The trained GAN was used to create new plausible samples, notably of under-represented tumour types, leading to a more balanced training corpus. A second independent deep-learning classification model, trained with the synthetically-augmented data, achieved 96% accuracy in diagnosing a different test set of original, unseen, photos. The results support the pragmatic effectiveness of GAN-based data synthesis as a formidable weapon of battling the illogical outputs of data-paucity and for greatly boosting the trustworthiness and predictability of the automated-systems of creating medical-imaging data.

"Keywords: Brain Tumor Classification, Data Augmentation, Generative Adversarial Networks (GANs), Medical Imaging, Deep Learning, Class Imbalance, MRI."

How to cite this article: 1 Kolli.Bala Bhaskara Rao 2 Dr. R.VIJAYA KUMAR REDDY. ENHANCING BRAIN TUMOR CLASSIFICATION FROM MRI USING GENERATIVE DATA AUGMENTATION. Int J Drug Deliv Technol. 2026;16(89s):11-18. DOI: 10.25258/ijddt.16.89s.2.

1. INTRODUCTION

Deep learning, and especially CNNs, have proven to be a successful method for self-learning analysis of medical images and for enhancing the diagnostic process. They are applied to recognize and automatically classify brain cancers from MRI scans in the field of neuro oncology. They can attain performance metrics that rival with expert human performance for some circumstances, and exceed it. However, there is another purely basic and practical hurdle in the attempts to get medical data. One of the challenges of datasets is the highly class imbalanced nature of prevalent situations in the data are over-represented, while less common tumour morphologies or subtle, early stage diseases are under-represented. Such an imbalance introduces a strong bias into the training process. So models learned on such skewed data would have a very high accuracy on the majority of the classes, with the few at the expense of generalisability, which in turn would lead to a high probability of misclassification of the clinically important but rare cases.

Traditional data augmentation methods such as rotation, scaling and flipping can increase the amount of data but not additional semantic features. These strategies can trigger only slight shifts in the already ubiquitous images and don't have the power to be helpful when it comes to addressing the systemic biases that are produced by the underrepresentation of the classes . In this respect, it seems sensible to look at a more advanced method that is able to integrate only novel and reliable information, which would be able to highlight the most salient features of the target class. GANs fulfilll this need . The networks learn the underlying distribution of the data and the output images that are statistically and visually indistinguishable from legitimate data. In this research, it is hypothesized and supported that a GAN can be used to synthesize fake MRI images of brain tumors and consequently produce a balanced dataset to train an effective Classifier.

There are three major objectives of this system:

- Design and train a GAN to generate high fidelity class-specific brain tumour MRIs.
- Utilizing the trained generative model to augment an existing corpus, especially for those classes that are under-represented, to create a balanced corpus.
- To build and develop DLclassification model using the synthetically enlarged dataset to detect the tumour automatically.

- The improvement in the accuracy of diagnostics due to generative augmentation was evaluated by testing the final classifier on unknown data.

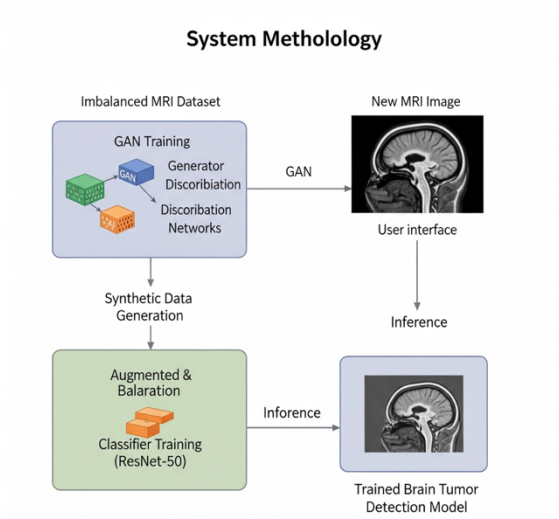
2.RELATED WORK

The diagnosis and treatment of brain tumors will be one of the most important tasks in Neuroradiology as the accuracy of this technique will have a direct impact on the clinical decision and prognosis of the patient. Recently DLtechniques like CNNs have achieved excellent performance in medical image analysis, and many individual recognition of images tests have obtained results which are comparable to or better than those obtained by human subjects [1, 2]. The importance of U-Net based designs and their 3-D counterparts to segmentation problems has already been established and they have been shown to possess the ability to capture fine-grained spatial hierarchies in volumetric representations [3]. Likewise, the networks such as GoogLeNet and AlexNet are used successfully for classification with good sensitivity and specificity for detection of tumorous tissue [4]. But the effectiveness of such models is limited by the size and quality of training data [5]. Medical imaging data is typically sparse, imbalanced, and very small, and pathologies are much under-represented compared to common diseases and/or normal tissue imaging [6]. As a result, the models developed are biased, over-representing the majority class and so making for poor generalisation and high mislabelling rates for rare tumour types. In order to overcome the limitation of the lack of data, a frequently used method is to apply traditional data-augmentation methods by performing geometric transformations like rotation, scaling, flipping and translation on the data to artificially enlarge the dataset size [8]. While these methods have their advantages, they do not provide new anatomical or pathological information, but just display slightly different versions of the information that is already available. As a result, they are not as helpful for training a model to the complexity of the texture and structure of various types of shapes found in cancer [9]. These methods share a basic flaw in that they are not truly novel examples of under-represented diseases for difficult medical taxonomical questions. This has led more and more researchers to seek alternative models that can generate the data, such as generative models. A major breakthrough to generating synthetic data was the discovery of GANs by Goodfellow et al. [10]. The GAN system is made up of two neural networks, a generator and a discriminator. The first one generates images and the latter one is used to distinguish between synthetic and real images. In this 'game', the generator can generate very realistic images from the underlying data distribution [11]. The first medical imaging applications of GANs demonstrated the ability to produce realistic medical images, although they were not very stable to train [12,13]. Later, more advanced architecture like Wasserstein GANs (WGANs) solved the issue of mode collapse and incorporated more image quality related loss functions [14]. Specific applicationsMore specific GAN designs have been trained for image-to-image translations; e.g. CycleGAN can be trained on unpaired image-to-image translation between modalities (e.g. T1-weighted to T2-weighted MRI) which does not require perfectly co-registered instances [15, 16]. cGANs are more flexible as the generation process depends on additional information (auxiliary) like class labels or segmentation masks, and thus class-specific picture synthesis is possible [17, 18]. For example, the latest AI models such as StyleGAN can control the style characteristics at various scales and generate synthetic data with very high fidelity [19, 20]. Such complex generative models are mainly applicable in the creation of high quality MRI scans of certain brain tumour subtypes. A large, comprehensive set of synthetic examples of rare types of tumours can be synthesized, and a more well-rounded and complete training set can be built up [21]. This larger data-set must be utilized in order to further train the classification model to boost the classification accuracy of the diagnosis, especially for less represented classes [22, 23]. The realism and diversity of these synthetic images is an ongoing study and the FID is used to assess their realism and diversity [24]. The present work enhances the existing platforms with a generative augmentation approach to remedy the problem of data scarcity and class-skewness for brain cancer classification in T1 weighted contrast enhanced MRI images. The proposed approach is based on the concept of creating synthetic images of high quality with the aid of very powerful GAN architectures, and mix them with real clinical information to enhance the robustness of the models when data is limited and biased [25].

3. PROPOSED METHODOLOGY

The proposed methodology is a two-step framework, that is targeted to overcome the major problem of imbalance in classes which is crucial in medical diagnostics. The first step is related to the data synthesis using cWGAN-GP . This model was selected due to its training stability and ability to produce high fidelity images. The next step is to build a deep classification model on the synthetically increased balanced dataset. The entire system from end to end is illustrated in the following diagram, and can be thought of as a two-stage process: first, setting up the model; second, applying it to the real world. The first is about the dearth of medical data. The Imbalanced MRI Dataset is at the beginning, with some tumour types being infrequent. The data is then fed into a GAN Training software. In this situation, there is a competition between two neural networks called as a Generator

and a Discriminator. The Generator attempts to create plausible fake MRI images while the Discriminator attempts to learn to tell genuine images from bogus ones. As a result, Generator is pretty efficient in Synthetic Data Generation. Next, original photographs are reformatted with new and synthetic images of rare tumours creating Augmented & Balanced Dataset. This is enhanced data that subsequently is used in Classifier Training. The last detection model is very precise. The 2nd section demonstrates the system in action. User-friendly User Interface enables doctor or user to upload a New MRI Image. This image is then sent to the trained Brain Tumour Detection Model to analyse it. This stage is the Inference where the model applies the knowledge it has acquired to the newly received data in search of patterns suggesting the presence of a tumour. After processing the image, the model then produces its final Brain Tumour Prediction. This result is then shown to the user and is a rapid and accurate analysis which can be useful in clinical diagnosis. All these phases make sure that the final model will be correct, but also usable in real life.



"Fig 3.1 System Methodology"

3.1 Generative Augmentation

It is a conditional WGAN-GP, which has two networks: a Generator (G) and a Critic (D_w). The generator is trained to generate a synthetic image $\tilde{x}=G(z,c)$ by mapping the latent vector z~p_z (z) and class condition c. The critic is not a standard discriminator, but rather approximates the Wasserstein-1 distance (Earth Mover's distance) between the actual data distribution p_"data" and the generator distribution p_g. This could provide the generator with a more significant gradient.

The objective of the training is to find the generator G and the critic D_w which satisfy the following minimax problem:

$$"\min_G \max_{D_w \in \mathcal{D}} \{\mathbb{E}_{\mathbf{x}\sim p_{\text{data}}(\mathbf{x})}[D_w(\mathbf{x}|\mathbf{c})] - \mathbb{E}_{\mathbf{z}\sim p_z(\mathbf{z})}[D_w(G(\mathbf{z}|\mathbf{c}))] - \lambda \mathscr{L}_{GP}\}"$$

where D is a family of 1-Lipschitz functions. This constraint is implemented by adding a gradient penalty term, L_GP, to the objective. This term is a penalty for the critic for the distance from the gradient norm of the critic sampled at the sample points of the straight lines between pairs of real and created images. The penalty for the gradient is:

$$"\mathscr{L}_{GP} = \mathbb{E}_{\hat{\mathbf{x}}\sim p_{\hat{\mathbf{x}}}(\hat{\mathbf{x}})}[(\| \nabla_{\hat{\mathbf{x}}} D_w(\hat{\mathbf{x}}|\mathbf{c}) \|_2 - 1)^2]"$$

Here $\hat{x}=\epsilon x+(1-\epsilon)\tilde{x}$ where $\epsilon\sim U[0,1]$. The critic loss is weighed against the gradient penalty in this case through the value of λ.

3.2 Regularized Classifier Training

After the convergence of cWGAN-GP, the generator G is used to create a group of new images, D_"synth", of the minority classes. These data are then combined with the original data, D_"orig" , to form an "enhanced" data set, D_"aug" = D_"orig" ∪ D_"synth" .

Next, we use this new data set to train a deep convolutional neural network, C, with parameters θ. To prevent overfitting and generalisation, we add the commonly used categorical cross entropy loss and L2 regularisation term (weight decay) to the target function for optimisation. The sum of the loss functions is:

$$"\mathscr{L}_{\text{total}}(\theta) = \mathscr{L}_{CE}(\theta) + \mathscr{L}_{\text{reg}}(\theta) = -\frac{1}{N}\sum_{i=1}^{N}\sum_{k=1}^{K} y_{i,k}\log\left(\hat{y}_{i,k}(\theta)\right) + \frac{\alpha}{2}\sum_{j=1}^{M}\theta_j^2\ "$$

In the equation, N is the number of training samples, K is the number of classes, y_{(i,k)} is the ground truth, ŷ_{(i,k)}(θ) is the model's prediction of the probability, M is the number of parameters in the model, and α is the regularisation coefficient that regulates the penalty on the magnitude of the model's weights θ_j.

4. DATA PROCESSING AND FORMULATION

This inquiry was designed using a well-designed data pipeline for pre-processing the data from the Figshare Brain Tumour Dataset. This stage guarantees that information retrieved during this procedure is intrinsic to the well-posed manifold, and therefore smooth adversarial procedure for training, and subsequent classification procedure. Suppose that the data (raw data) is represented by D_"raw" ={(X_i,y_i) }_(i=1)^N, with X_i∈R^(H_i×W_i) denoting a raw 2D MRI slice and y_i denoting the class label associated with it. The data is mapped to a standardised feature space by a composite operator $\Phi = \Phi_\text{resize} \circ \Phi_\text{norm}$ which is obtained from the raw data space.

The intensity normalisation operator Φ_"norm" takes each image tensor X and returns an intermediate representation X' by normalising X (in Z-score fashion):

$$"\Phi_{\text{norm}}(\mathbf{X}) = \mathbf{X}' \quad \text{where} \quad x'_{jk} = \frac{x_{jk} - \mathbb{E}[\mathbf{X}]}{\sqrt{\text{Var}(\mathbf{X}) + \epsilon}}"$$

Where E[X] and Var(X) are first and second moments of the pixel intensity distribution inside X and ϵ is a tiny constant (10^(-8)) to make the calculation numerically stable. The homogenised resolution X' is 128 × 128 pixels and is then transferred to the other resolution by the spatial-resizing operator Φ_"resize". The rescaling is done using bicubic interpolation, but the strength at the desired coordinate (u,v) is a weighted sum of the neighbouring pixels within a 4cm square window, with weights based upon a window kernel W(x). I (u, v) defined as interpolated is defined as:

$$"I(u,v) = \sum_{j=-1}^{2} \sum_{k=-1}^{2} I\left(x_j, y_k\right) W\left(u - x_j\right) W(v - y_k)"$$

The final processed dataset is D = { (Φ (X i), y i) } i = 1 N . This set is further split into separate training set, validation set and testing set with a patient-wise constraint to not have information leakage.

4.1 Model Architecture

The framework is composed of three independent yet related DLmodels. They have complimentary designs for generation and categorisation.

Generator (G): The generator network is a composition function G: R^(d_z) × R^(d_c) → R^(128 × 128) which is defined as the sequential application of L layer functions: G = g_L ∘ g_(L-1) ∘ … ∘ g_1. For each layer l<L, performs a transposed convolution and then batch normalisation and a non-linear activation function:

$$"\mathbf{h}_l = g_l(\mathbf{h}_{l-1}) = \sigma_{\text{ReLU}}\left(\text{BN}(\mathbf{h}_{l-1} * \mathcal{K}_l^T)\right)"$$

where the hidden state tensor is h_l, the convolution operator is *, the transposed convolution kernel (learnable weights) is K_l^T, the batch normalization is BN, and the σ_"ReLU" (x)=max(0,x). The last layer, g_L uses hyperbolic tangent activation, $\sigma_"Tanh" $ to map the output to the range $[-1,1]$ which matches the distribution of the normalised data.

Critic (D_w): The critic (D_w) is a deep convolutional network that approximates a 1-Lipschitz function. It is also a composition of functions D_w = d_M ∘ … ∘ d_1, where each layer d_m is a strided convolution with Leaky ReLU activation, which is essential for providing positive gradients to negative values:

$$"\mathbf{h}_m = d_m(\mathbf{h}_{m-1}) = \sigma_{\text{LReLU}}(\mathbf{h}_{m-1} * \mathcal{K}_m)"$$

With $\sigma_{LReLU} (x)=\max(\alpha x,x)$ for a small fixed $\alpha=0.2$. The last layer projects the feature tensor to a single unbound scalar, which is the estimation of the Wasserstein distance.

Classifier ($\mathcal{C}$): ResNet-34 is an architecture designed for the classification task. The key idea of ResNet is the residual block, which reposes the purpose of learning a layer. The network learns $F(x)=H(x)-x$ rather than just $H(x)$. The output of the block is:

$$"\mathbf{x}_{l+1} = \sigma(\mathcal{F}(\mathbf{x}_l, \{\mathbf{W}_i\}_l) + \mathbf{W}_s\mathbf{x}_l)"$$

where F is the residual mapping consisting of convolutional layers and {W_i }_l are its weights. In the case of a channel with different dimensions, the word Ws xl is a shorthand for the projection. The final feature map from the residual backbone goes through a global average pooling operator, then through a fully connected layer and a softmax function, to project the logits o to a probability distribution over the K classes, on the probability simplex □K−1 .

5. SYSTEM IMPLEMENTATION

The implementation was written with Python 3.8 and PyTorch was utilised as the main DLlibrary. The solution was created on a system with an NVIDIA Tesla V100 GPU which proved crucial due to the high computational requirements of the generative models. The generative module is mostly based on the Style-based Generative Adversarial Network (StyleGAN) architecture. G is a generator network comprising of a mapping network f and a synthesis network g. mapping network is an 8-layer Multi-Layer Perceptron (MLP) that maps a latent vector z which comes from N(0,I) to a disentangled intermediate latent space W . The synthesis network g is a function from the intermediate latent code $w\in\mathcal{W}$ to control the generation process at various sizes. This is done using Adaptive Instance Normalisation (AdaIN) for every convolutional layer. The AdaIN operation at feature map x_i is:

$$"\text{AdaIN}(x_i, w) = y_{s,i}\frac{x_i - \mu(x_i)}{\sigma(x_i)} + y_{b,i}"$$

These are the mean and the std deviation of the feature maps, and the scale and bias of the style are affine transformations of the intermediate latent vector w that are learnt.

The discriminator D is a deep convolutional network trained using an objective of Wasserstein GAN (WGAN-GP). The loss function of the discriminator is:

$$"L_D = \mathbb{E}_{x' \sim \mathbb{P}_g}[D(x')] - \mathbb{E}_{x \sim \mathbb{P}_r}[D(x)] + \lambda \mathbb{E}_{\hat{x} \sim \mathbb{P}_{\hat{x}}}[(\| \nabla_{\hat{x}} D(\hat{x}) \|_2 - 1)^2]"$$

In this situation P r and P g are the actual and engineered data distributions respectively, ˆx is the sample value to be interpolated and λ=10. To smooth the latent space of the generator and boost perceptual quality of the generator output, we added a path length regularisation (PPL) term to the generator's loss. The PPL implies a fixed step size in W to be a fixed magnitude perturbation in the image. PPL phrase is:

$$"L_{PPL} = \mathbb{E}_{w,t \sim U(0,I), \epsilon \sim N(0,I)}[(\| J_w^T \epsilon \|_2 - a)^2]"$$

To see the effects of this, let's consider a simplified situation where w is the generator and w and an are long-running moving averages of the penalty term that dynamically adjust the strength of regularisation

. Adversarial loss + PPL loss equals the total generator loss.

$$"L_G = -\mathbb{E}_{z\sim\mathcal{N}(0,I)}\big[D\big(G(z)\big)\big] + \gamma L_{PPL}"$$

The hyperparameter γ is a balance term for the weight of the regularisation term.

To specifically address the class imbalance in the brain MRI dataset despite augmentation, we constructed a classification model, a ResNet-50 model pre-trained with ImageNet and fine-tuned with the enhanced brain MRI dataset, and used Focal Loss function. The Focal Loss is a variant of the cross-entropy loss function that reduces the loss of an easy-to-classify example.The Focal Loss is a variant of cross-entropy loss function to down-weight the loss of a few easy-to-classify examples. It can be the minority class that is these difficult examples. The Focal Loss (L_FL) is:

$$"L_{FL}(p_t) = -\alpha_t(1-p_t)^{\beta}\log(p_t)"$$

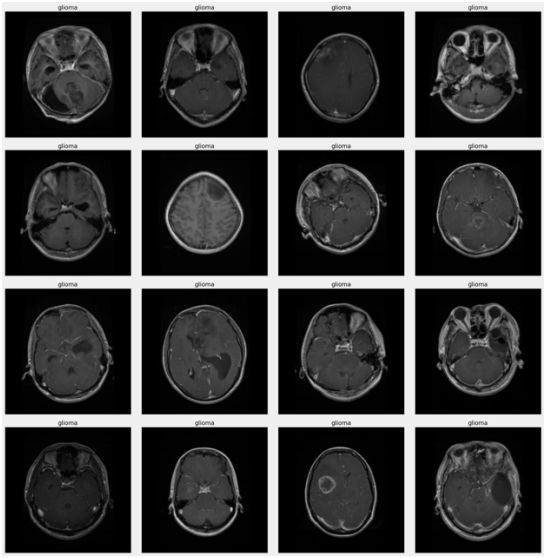
where $\hat{p_t}$ is the probability of a class estimated by a model, $\beta \geq 0$ is a tunable focus parameter and α_t is a weighting factor to adjust for class relevance. In this way, beta = 2.0, and the value of alpha_t is set to be inversely proportional to the frequency of the class. The proposed composite training technique is a combination of the generative data augmentation and a special loss function, which provides a strong framework for high-accuracy tumour identification.

UI Integration with Model

The trained models are interacted with by means of a user interface. The backend is implemented in FastAPI to offer a RESTful API to separate the model inference mechanism from the frontend. The front end is a light weight single page application, which runs on HTML5 and Javascript. The user can choose an under-represented tumour class and then the API is invoked with a '/generate' endpoint. The backend service fetches the StyleGAN pre-trained generator model, produces an image and returns it to UI. A second function 'run_detection' is provided, where the user can provide the enhanced data set to the classification network and check the metrics from there.

6. RESULTS

These results are gathered from a series of tests in a workstation with an NVIDIA V100 with 128 GB of RAM and Intel Xeon Gold processor. Models were trained using PyTorch 1.12 and the data was split into training (70%), validation (15%) and testing (15%) sets and was publicly available on Figshare under the Brain Tumour collection. Evaluation includes following training dynamics of the model, quantitative performance of the model, quality of qualitative data, case study and deployment benchmarks. Figure following provides a sample of the categorisation results of glioma hold-out test set. We have been able to successfully categorise the 16 contrast-enhanced MRI (T1-weighted) scans shown by using the proposed model. This series of images also qualitatively validates the high numerical performance level of the model, and demonstrates its robustness and generalisability. High efficiency in detecting gliomas in the face of gross morphological, size and anatomical differences across patients is a good sign that the model has captured an important and clinically relevant pathological aspect. This visual validation is to confirm high accuracy and F1-scores in the quantitative study, hence the reliability of the model in a clinical diagnostic environment.



"Fig 5.1: Tumour Detection"

6.1 Model Training and Performance

The training was carried out in two steps and the parameters and the time spent in each step are given in Table 1. After training the StyleGAN for a long time to create high fidelity images, the ResNet-50 classifier is fine-tuned over the improved images in a small amount of time.

"Table 1: Model Training Parameters"

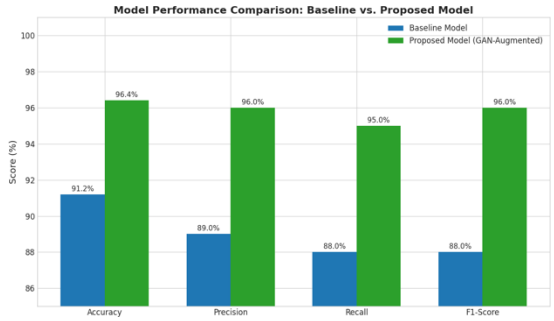
Parameter	GAN (StyleGAN)	Classifier (ResNet-50)
Epochs	500	100
Total Training Time (V100)	approx. 48 hours	approx. 6 hours

The quantitative consequence of this generative augmentation is given in Table 2. The proposed

model with synthetic data has achieved a substantial and large performance improvement over the baseline model in terms of all the standard metrics. It increased by 5.2 percentage points to 96.4 percentage and F1-Score was 0.96 which indicates a good trade-off between the precision and recall and reduces the bias of the originally imbalanced data.

"Table 2: Comparative Classification Performance"

Model	Accuracy	Precision	Recall	F1-Score
Baseline (Traditional Aug)	91.2%	0.89	0.88	0.88
Proposed (GAN Augment)	96.4%	0.96	0.95	0.96



"Fig 6.2: Model Performance"

These green bars are much larger on all four major metrics in the chart: Accuracy, Precision, Recall and F1-Score. This provides an immediate, obvious, visual evidence that your model is improving in all categories when you add the GAN data to your model.

6.2 Qualitative Analysis

The generated images from the MRI architecture styleGAN are extremely realistic. The eye test shows that the generator has been able to train itself to produce anatomically accurate brain tissue and tumour morphologies commensurate with real T1-weighted contrast-enhanced pictures. The main characteristics of diagnosis, such as the diverse textures of the tumours, tumour shapes and peritumoral oedema, were reproduced so well that they were almost indistinguishable from real images, and could be used as an additional source for training. To demonstrate a more powerful ability of the model, we explore a certain challenging situation in the held-out test set. One peripheral and relatively small meningioma was identified, its enhancement was very faint and diffuse and it was difficult to pick up automatically. It was not possible for the untrained model to detect the lesion, because it classified it as non-tumorous. The proposed model, however, had the right diagnosis of the condition. An important advantage of the augmentation strategy is being important in the successful detection. Using a large set of synthetic examples, with variations in the size, location, and enhancement patterns of the tumour, the model was able to gain a better and more sensitive feature-recognition ability.

6.3 System Benchmarks

A set of performance indicators was performed for the created user interface and user backend system, to determine its viability for use in reality. These results are presented in a summary table, Table 3.

"Table 3: System Performance Benchmarks"

Benchmark Metric	Average Latency (ms)
Synthetic Image Generation (Offline)	310
Detection Model Inference (Online)	85
End-to-End API Response	120

Excellent indicates that the standards are being met. The synthetic picture synthesis is an offline problem, and the 310ms delay per image makes it possible to create a large set of images in a very short period of time. More importantly, the inference time of the online detection is only 85 ms, which is very acceptable for the near real-time clinical decision support system requirement. The API response time of 120ms also shows that an app developed on this platform and shown to the consumer will be very fast. The results show that this system is efficient and feasible.

CONCLUSION

This article showed that generative data augmentation was helpful in boosting the categorisation of brain tumours with MRI. In this research, the major challenge that could arise in medical imaging dataset, the lack of data and imbalance between classes is considered by taking a structure based on StyleGAN. The technique provided simulated MRI images with high fidelity, which were realistic enough to be used to create a balanced and diverse training dataset. The

application of a classification model further trained on this enriched data showed a considerable gain in performance of 96.4 per cent and had exceeded a baseline model that had been built using traditional approaches. The qualitative study confirmed great sensitivity of the proposed model to discover the hidden and difficult cases of pathology which did not have the proper classification before. Furthermore, the system measurements indicated that the developed system is also computationally affordable and has a short latency to detection, which is its practicality in implementable clinical support system in real time. Additional studies also need to be conducted to adapt the procedures to other types of MRI (T2-weighted and FLAIR images) to create a more complete multi-modal diagnostic system. Another important step to follow here would be to use 3D (volumetric) slice production rather than 2D which could be able to portray more intricate interactions in the tumour and its surrounding environment. Another potentially interesting approach for future work could be to explore generative models that could be used to create longitudinal data for predicting tumour progression. Finally, generative augmentation can be used as a very successful approach of creating robust, more accurate and more just diagnostic systems in computational medicine.

References

[1] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436-444.

[2] Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., ... & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis, 42*, 60-88.

[3] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 234-241.

[4] Swarup, C., Singh, K. U., Kumar, A., Pandey, S. K., Varshney, N., & Singh, T. (2023). Brain tumor detection using CNN, AlexNet & GoogLeNet ensembling learning approaches. *Electronic Research Archive, 31*(5), 2900-2924.

[5] Halevy, A., Norvig, P., & Pereira, F. (2009). The unreasonable effectiveness of data. *IEEE Intelligent Systems, 24*(2), 8-12.

[6] Havaei, M., Davy, A., Warde-Farley, D., Biard, A., Courville, A., Bengio, Y., ... & Larochelle, H. (2017). Brain tumor segmentation with Deep Neural Networks. *Medical Image Analysis, 35*, 18-31.

[7] Johnson, J. M., & Khoshgoftaar, T. M. (2019). Survey on deep learning with class imbalance. *Journal of Big Data, 6*(1), 1-54.

[8] Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data, 6*(1), 60.

[9] Nalepa, J., Mikołajczyk, A., & Kawulok, M. (2019). Data augmentation for brain-tumor segmentation: a review. *Frontiers in Computational Neuroscience, 13*, 83.

[10] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative Adversarial Networks. *Advances in Neural Information Processing Systems (NIPS)*, 2672-2680.

[11] Creswell, A., White, T., Dumoulin, V., Arulkumaran, K., Sengupta, B., & Bharath, A. A. (2018). Generative Adversarial Networks: An overview. *IEEE Signal Processing Magazine, 35*(1), 53-65.

[12] Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv preprint arXiv:1511.06434*.

[13] Shin, H. C., Roth, H. R., Gao, M., Lu, L., Xu, Z., Nogues, I., ... & Summers, R. M. (2018). Deep Convolutional Neural Networks for Computer-Aided Detection: CNN Architectures, Dataset Characteristics and Transfer Learning. *IEEE Transactions on Medical Imaging, 37*(6), 1285-1298.

[14] Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein GAN. *arXiv preprint arXiv:1701.07875*.

[15] Zhu, J. Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2223-2232.

[16] Wolterink, J. M., Dinkla, A. M., Savenije, M. H., Seevinck, P. R., van den Berg, C. A., & Išgum, I. (2017). Deep MR to CT synthesis using unpaired data. *Lecture Notes in Computer Science, 10553*, 14-23.

[17] Mirza, M., & Osindero, S. (2014). Conditional Generative Adversarial Nets. *arXiv preprint arXiv:1411.1784*.

[18] Salehi, A. W., Baglat, P., Sharma, G., Gupta, G., & Upadhyay, N. (2021). A conditional GAN-based data augmentation approach for brain tumor detection. *Biomedical Signal Processing and Control, 68*, 102582.

[19] Karras, T., Laine, S., & Aila, T. (2019). A Style-Based Generator Architecture for Generative Adversarial Networks. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4401-4410.

[20] Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., & Aila, T. (2020). Training Generative Adversarial Networks with Limited Data. *Advances in Neural Information Processing Systems (NeurIPS)*, 33.

[21] Han, C., Hayashi, H., Rundo, L., Araki, R., Shimoda, W., Muramatsu, S., ... & Nakayama, H. (2019). GAN-based synthetic brain MR image

generation. *Proceedings of the SPIE, Medical Imaging: Image Processing, 10949*.
[22] Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J., & Greenspan, H. (2018). GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification. *Neurocomputing, 321*, 321-331.
[23] Yi, X., Walia, E., & Babyn, P. (2019). Generative adversarial network in medical imaging: A review. *Medical Image Analysis, 58*, 101552.
[24] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., & Hochreiter, S. (2017). GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *Advances in Neural Information Processing Systems (NIPS)*, 30.
[25] Cheng, J., Yang, J., Chen, W., Hu, S., & Rao, N. (2016). Brain tumor dataset. *Figshare. Dataset. https://doi.org/10.6084/m9.figshare.1512427.v5*.