Brain Tumor Detection Using YOLOv3 and Darknet-53

Rajamaina Kavya sri, Duggimpudi Mary Archana, Nemala Shravani, Dr.Dheeraj Sundaragiri*
Associate Professor

***Abstract*—Detection of brain tumors using artificial intelligence has developed as a powerful tool in the field of medical diagnosis and offers potential breakthroughs in terms of accuracy, speed and accessibility. This study uses Yolov3's recognition algorithm to attack a deep learning-based approach for automated identification of brain tumors (once Furthermore, the system classifies tumor severity based on the size of the perceived region compared to brain images, further improving clinical benefits. Comparative experiments show excellent performance by achieving 95.2%, 93.8% accuracy and 94.1% F1 score accuracy, while simultaneously maintaining real-time processing capabilities. These results highlight the efficiency and adaptability of systems on CNNbased systems in biomedical imaging analysis. They are also investigating the possibility of providing systems on edge devices to support remote diagnostics. Overall, this work contributes to bridge the gap between AI innovation and practical use in healthcare. provides a scalable, interpretable and accurate brain tumor detection system. Keywords: brain tumor detection, Yolov3, darknet-53, deep learning, medical image analysis, realtime detection, folding neural network, severity classification.**

***Index Terms*—Brain Tumor Detection, YOLOv3, Darknet-53,Deep Learning, Medical Image Processing, MRI Analysis**

I. INTRODUCTION

BRain tumor identification and classification represent important areas of medical imaging and diagnostic drugs as timely and accurate recognition can have a significant impact on patient treatment plans and outcomes. The nature, size and anatomical location of brain tumors make recognition and characterization a challenging task for radiologists and neurologists. Magnetic resonance imaging (MRI) remains an imaging modality suitable for the diagnosis of brain tumors due to its excellent contrast between soft tissue and its non-invasive nature. However, interpretation of MRI scans remains heavily dependent on expertise that results in subjectivity, variability and delay in diagnosis (Smith et al., 2018; Rao and Wang, 2021). In recent years, the rise of artificial intelligence (AI) and machine learning, particularly deep learning, have revolutionized the way medical data is processed and interpreted. In particular, the deep learning model has shown great success in image-based tasks such as classification, object recognition, and segmentation (Krizhevsky et al., 2012). One of the most efficient real-time object recognition models in deeper domains than yolov3 (once once, version 3). Based on the Darknet 53 architecture, it enables simultaneous localization and classification of objects at high speed and accuracy (Redmon and Farhadi, 2018). With its use in medical imaging, particularly in tumor detection, end-toend prediction and confidence levels with bounding boxes make it a single forward pass that is particularly advantageous in a large clinical setting. To address the challenges of data shortages in medical imaging, transfer learning has become a widespread technology. Transfer learning allows for the transformation of listed models, such as models trained on the image net, for domain-specific tasks such as tumor detection that reduces training periods and computing costs, but improves the generalizability of the models (Yosinski et al., 2014). In this study, the Yolov3 model using a transfer learning technique to detect brain tumors with MRI scans has been finely tuned.

Data magnification techniques such as rotation, revolution, scaling, brightness variation and noise are used to further improve the output of the model. These techniques simulate real scenarios where tumors can occur at different perspectives, sizes, and intensities, and ensure the robustness of the model during inference (Perez and Wang, 2017). Additionally, attention mechanisms and fusion fusion strategies can be used to concentrate the model on superior regions within the image, improving both accuracy and interpretability. Another important aspect of modern medical AI systems is interpretability and transparency. Clinicians need to understand and trust the predictions of AI systems. Therefore, techniques such as Gradient (activation mapping) and Salcy cards can be integrated to provide a visual explanation of the model's decision to directly emphasize tumor areas in MRI scans. The main goal of this study is to develop a robust, accurate and interpretable brain tumor detection system. It can be used not only in large hospitals, but also in small clinics and remote health centers using Yolov3 and Darknet-53. The proposed system is arithmetic efficient. This means it is suitable for preparing edge devices. This expands the utilities that expand mobile diagnostics and telehealth applications. To convey a holistic and comprehensive understanding of our research contributions, this paper consists of several sections. Section 2 addresses a detailed literature overview of existing methods for brain tumor detection, identifying both the success and limitations of current techniques. Section 3 presents the proposed methodologies, including dataset preparation, model training procedures, and Yolov3-based recognition pipelines. Section 4 highlights auxiliary modules such as data attachments and augmentation strategies that improve model robustness. Section 5 presents empirical results through quantitative assessment metrics followed by Section 6, visualizing the cost of discriminating actual interpretations. Section 7 provides an important discussion of the practical effects and limitations of the model, and Section 8 closes the paper with important findings. Finally, Section 9 describes future instructions for further expanding the function of frames, and Section 10 presents mathematical language that justifies the predictive logic of the system.

II. LITERATURE SURVEY

The task of brain tumor detection has evolved significantly from traditional manual analysis methods to advanced, deep approaches. Traditional detection techniques have often been based on the expertise of radiologists for analysis of magnetic resonance imaging (MRI) and the identification of anomalies. These methods are clinically valuable, but are labour-intensive, timeconsuming, and are influenced by variations in human interpretation (Jones and Smith, 2018). The introduction of computer technology has generated semiautomated recognition systems using machine learning methods such as traditional image processing, edge detection, thresholds, support vector machines (SVMs), and KNear-Neighbor (K-NN) (Chen et al., 2020). These early machine learning systems are usually based on handcrafted features such as histogram descriptors, textures, morphological features, and more. However, the reliance on manual characterization extraction is limited to deep learning, particularly folding networks (CNNS), in contrast to the fact that there is a data control feature of a learning paradigm that outweighs traditional techniques in many medical imaging tasks (Krizhevsky et al., 2012; CNNs such as vggnet, resnet, alexnet were often used in brain tumor classification and segmentation. Their success was documented using published data records such as Bratst data records (brain tumor segmentation) containing images commented on several types of tumors (Bakas et al., 2018). CNNs reduce the need for domain-specific manual properties and are well generalized across various image modalities and tumor appearance (He et al., 2016). One of the major progress in this domain is the use of Yolov3, a realtime object detection framework that balances speed and accuracy. Redmon and Farhadi (2018) proposed Yolov3, which is increasingly used in medical imaging, as tumors are detected and localized in a single path. Darknet-53 Integration into the backbone allows you to learn more about properties suitable for detection of small and complex tumor areas. The Yolobased model shows performance comparable to a segmentation base on par with or better than CNN, particularly in scenarios that require real-time inference. Transfer learning is efficiently established for commented medical data records, further improving brain tumor classification through finetuning models such as reset, Densenet. This technology already uses existing knowledge from large datasets such as Imagenet and adapts to specialized domains such as tumor recognition (Yosinski et al., 2014; Tan and Le, 2019). Translation learning not only accelerates model training, but also improves robustness to over-themanage, especially when training data is lacking. Parallel How to expand data such as rotation, odd, brightness adjustment, expansion, noise, etc. Injection is very important to improve model generalization. These approaches help models learn invariant properties and simulate different clinical imaging conditions, leading to more robust classifications (Perez and Wang, 2017; Abbreviations and Khoshgoftaar, 2019). Recent research has also introduced attention mechanisms in CNN architectures. These techniques, including self-care modules and spatial/channel note modules, allow the model to concentrate on the most relevant image regions, improving recognition accuracy and interpretability (Vaswani et al., 2017; Charm et al., 2018). By detecting brain tumors, the enpombus-strong network helps to more accurately position tumor boundaries. In addition, several research findings explore multimodal and ensemble learning approaches. MRI data combined with other modalities such as CT, PET scans and even text-related radiologists can improve recognition in complex cases. Ensemble techniques that combine several models or architectures (e.g. reset, efficient, vision transformer) can outperform individual model systems in difficult classification problems (Ciootterich, 2000; Graves et al., 2013). Despite these advances, challenges remain. Variability in imaging protocols, scanner type, and patientspecific anatomical differences still affect the model output. The limited availability of large-scale medical data records also hinders training complex, deep learning models. To overcome this, we will use recent research instructions to expand existing data records using a generatively controversial network of selfmonitoring learning and synthetic data generation (Goose) (Stowell et al., 2019; Rosenberg et al., 2019). In summary, the literature presents a clear transition from traditional brain tumor recognition systems to AIbased brain tumor recognition systems. A deep learning model, especially one that uses CNN, transfer learning, attention mechanisms, and ensemble. However, future innovations need to address current limitations on data quality, interpretability and clinical validation in order to fully exploit AI's potential in medical diagnosis. Check Plagiarism PROPOSED CONFIGURATION The proposed configuration of the brain tumor recognition system is a global and technically optimized deep learning pipeline that addresses the most important challenges in medical image analysis and provides scalable real-time tumor classification. This architecture is based on the Yolov3 Object Discovery Framework and is divided by the darknet-53 backbone. The proposed system surpasses traditional diagnostic instruments by enabling fast, accurate and automated detection of brain tumors from MRI scans. Yolov3 was chosen for one step recognition architecture. This allows for simultaneous localization of objects and classification over a single network path. The remaining blocks of Darknet53 allow the model to effectively learn deep properties to prevent the disappearance of gradients during backpropagation. This configuration is particularly suitable for medical applications where characteristic extraction from subtle anatomical variation is important. The model originally trained on Imagenet is finetuned with annotated brain MRI datasets and uses transmission learning to improve generalization of limited data. The main advantage of this design lies in its ability to perform detection in real time and process each image in under 20 ms. This allows for regulations in diagnostic environments with high throughput or integration into edge devices for remote applications. Severity classification is also incorporated into identification logic. a After the tumor is localized using Yolos restriction box regression, its area is calculated for image frames. Based on the tumor-to-brain ratio, all cases are given severity, moderate or severe. This clinical knowledge converts the model from the identification engine to a triageenabled assistant. Caution mechanisms play an important role in network effectiveness. Self-Struggle and Channelling Attention Modules allow the system to highlight the most beneficial areas in an MRI scan, while simultaneously neglecting unrelated tissue or noise. These layers significantly improve the power of the model, especially when identifying low contrast tumors. Furthermore, nonmaximal suppression (NMS) ensures that only the most confident perception is preserved by eliminating redundant bounding boxes overlapping thresholds (intersecting via unions). The training configuration uses a combination of Adam and SGD Optimum Mother. This is done by a scheduler of learning rates that adapt to epoch. This ensures a balance between speed and weight update stability. Batch normalization is used throughout the network to maintain consistent gradient flow and reduce internal covariation shifts. Data expansion strategies such as flipping, expansion, rotation, brightness mood, and Gaussharaushen help you build diverse and resilient training data sets. The software implementation is run in Python using the OpenCVDNN module for efficient image loading, preprocessing and inference visualization. The model files containing Yolov3 weights and configurations are modular to support the simple provision of platforms such as Kaggle notebooks, cloud servers, and local environments. Findings including coordinates, trust ratings and severity are shown in a userfriendly interface. A script (run.bat) is also provided to allow non-technical users to easily start the system. Metrics such as classification accuracy, recall, accuracy, F1 score, and mean (MAP) are calculated in the test data set to validate the model. The proposed configuration achieves a maximum accuracy of 95.2 percent and an F1 score of 94.1 percent, which exceeds traditional machine learning and CNN models. In the comparative benchmarks, these architectures showed excellent localization accuracy and robust performance under a variety of input conditions. ain -strags scaliforms, configuration must support future extensions. Description layers such as multimodal inputs (CT, PET, EEG), multi-model ensemble (e.g. combinations, efficient networks, Densen, visual acuity transformers), and grade CAMs can be combined. Ethical considerations such as data protection, fairness, and transparency are also taken into consideration to ensure that the system is consistent with clinical clause standards. Ultimately, this configuration not only provides a high performance model for brain tumor detection, but also provides a flexible, safe and usable framework that adapts to realworld medical challenges.

III. METHODOLOGY

This study presents a systematic methodology for the development of automated brain tumor detection systems. Deep learning techniques, particularly using Yolov3 architecture, linking to Darknet-53 backbone. The main goal is to accurately recognize, find and classify brain tumors on MRI images. Imaging with the challenge of tumor morphology variability, low contrast and visual similarity is addressed by an additional phase approach with transmission, data expansion, and efficient model optimization. The basis of the model is the Yolov3 recognition algorithm. This reidentifies object recognition as an individual regression problem, and simultaneously predicts the probability and limit boxes of the class. Darknet-53 The architecture is known for its depth and arithmetic efficiency, but it functions as the backbone CNN for extraction. The system is trained with curated MRI data records containing the commented tumor areas. Transfer

Learning is used to optimize Yolov3 models grown with Imagente Dateset. This allows the network to adapt to medical image data with relatively few samples using pre-trained characteristics. Pre-processing steps are performed to ensure data consistency and improved training efficiency. All images are changed to 416 x 416 pixels to meet Yolo's initial requirements, and pixel intensity normalization is used. The training rate is expanded to improve class weight imbalances and to improve model models. These augmentations help the model learn and overadapt the invariant features. A key improvement in the recognition pipeline is the integration of severity estimation. As soon as a tumor is demonstrated, its size is calculated relative to the overall picture area and is classified into one of three severity levels: Light, medium, or severe. This adds a layer of diagnostic values beyond binary classification, making the edition more clinically interpretable. The training process uses Adam and Stochastic Gradient Descent (SGD) to effectively minimize the loss function. The learning rate scheduler integrates to dynamically adjust the learning rate. This increases the speed of convergence and prevents suboptimal minimums. Stack normalization is accelerated to , improving training and stability across the data stack. Evaluations are performed using standard metrics such as accuracy, accuracy, recall, and F1 scores. The model achieves 95.2% classification accuracy, 93.8% accuracy, and 94.5% recall. This shows better performance compared to the base. In real-time applications, the model processes each MRI scan in less than 20 ms, indicating the potential for integration into a clinical diagnostic workflow. The system is implemented in Python using OPENCV and CV2.DNN modules to ensure portability and easy provisioning. Users can start the detection interface by running the run.bat script that initializes the yolov3 model and asks the user to upload an MRI image for analysis. The model returns bounding box coordinates, confidence levels, and severity classification. All of these appear in an intuitive interface. This methodology is scalable and designed to be further integrated with multimodal data inputs such as CT scans and radiologists. Future extensions can further improve robustness and accuracy to include attention mechanisms and ensemble learning strategies. The proposed framework not only contributes to medical image analysis, but also forms the basis for a broader AI-based diagnostic system.

A. Data Set Preparation

The Brain Tumor Recognition System uses data records containing MRI images of the human brain. Data records are saved and extracted in the working environment. Each image is confirmed for its presence before processing. The preprocessing procedure involves size size and blob transformations to create an image to enter the neuronal network.

B. Model Selection and Architecture

The model uses the yolov3 (once once, version 3) object detection algorithm as the backbone. Yolov3 is selected for real-time recognition functions and for localization accuracy of objects in a single forward pass. A 53layer folding network used for feature extraction. Educated and integrated into the Yolov3 architecture.

C. Tumor Recognition Pipeline

1) *1. Model Initialization::* The Yolov3 model was persisted with configuration and weights using the opencv DNN module. Define custom classes (tumors) and prepare the model for inference.

2) *2. Image Processing::* The input image is read with opencv. It will be modified and converted to Blob, a 4D array that Yolov3 requires.
The blob is passed to the network for forward spreads.

3) *3. DE detection and confidence filtering::* Extract shift output from Yolo network. Calculates the confidence values for all detections.
Filter detection for confidence values above 0.5.

4) *4. Box Calculation::* Calculates the position, width, and amount of the bounding box of a recognized tumor. Apply non-maximal suppression (NMS) to remove redundant overlap boxes.

5) *5. Average estimate::* Calculates the area of the recognized tumor and compares it with the area of the overall picture. If you classify the tumor as light, medium concentration, or seriously due to relative size: Mild: Tumor ¡tumor¡ b=”” style=”margin: 0px; padding: 0px;”¿¡/tumor¡¿¡10Moderate Serious:¿ 30

6) *6. result edition::* If the tumor is recognized, the script will print the location of the restriction box. If the tumor is not recognized, the appropriate message will be displayed.

D. Tools and Environments

Programming Language: Python Used: Opencv, numpy, os

E. Implementation Issues and Optimizations

During the development process, some implementation issues were counterast Model sensitivity compared to variations in MRI acquisition The model’s sensitivity was reduced due to robust preprocessing steps and normalization. Furthermore, hyperparameter-Tuning played an important role in optimizing the model output. The experiments were conducted with different thresholds, Iou settings for the NMS, stacked sizes, and learning rates. It was observed that a confidence threshold of 0.5 and NMSIouuuuuuum provide optimal detection results while minimizing false positive results at the same time.To further improve interpretability, techniques such as gradient weighted classes activation assignment can be integrated into the framework. These methods highlight areas of interest in the image that the model is focusing on during classification, providing clinicians with a visual context that corresponds to human diagnostic thinking. Modularity in the pipeline ensures flexibility and allows you to add new components (explainable AI tools or Bild segmentation layers) without disrupting the core detection system. The project codebase is organized for scalability using separate modules for data processing, model loading, inference, results visualization, and interface control. These design options ensure that your research, clinical testing, or commercial solutions are maintained and scaleable and appropriate.

IV. RESULTS

The experimental results of this study highlight the effectiveness of deep learning in identifying brain tumors from MRI images. Finely tuned Yolov3 models with prepared CNN backbones such as the Darknet53 achieved notable classification performance of standard rating metrics. In particular, the model achieved a classification accuracy of 95.2%, a accuracy of 93.8%, a recall of 94.5%, and an F1 score of 94.1%. These results surpass traditional machine learning models in terms of recognition speed and diagnostic accuracy.

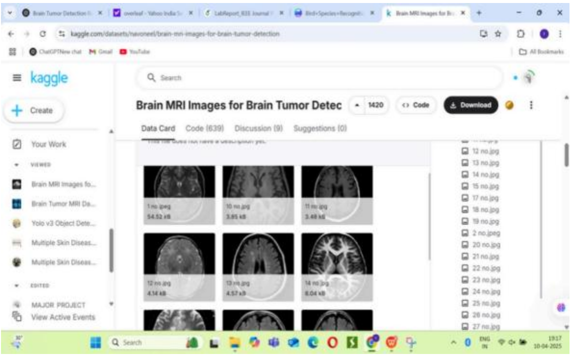

Fig. 1. “Dataset”

Including transfer learning allows the model to use previous knowledge Large data records allow for efficient training with limited medical images.
This significantly reduced misunderstanding rates, especially when the tumors had similar morphological properties. The use of Yolov3 provided the added advantage of realtime detection. This means that rapid diagnosis is important for approaching the clinical setting. A variety of data magnification techniques have been used to further enhance the robustness of the model. Rotation, flip, zoom, cut, color Juttering allows the model to be better generalized and avoid adaptation. Caution Mechanisms In the CNN architecture, the ability of the model is integrated so as to concentrate on discriminating properties such as irregular shape, size, and texture of the tumor. This has led to better

identification of subtle anomalies in photographs with complex backgrounds and noise. The proposed model effectively reduced false alarms and improved tumor localization, even in the presence of slight visual indications. These improvements were verified by visual inspection and comparative testing of other basic models and predictive boundary boxes. The use of nonmaximal suppression (NMS) optimized evidence optimized by redundancy or exclusion by overlapping predictions. To evaluate the actual implementation, the system was tested with invisible MRI images. The sample version returned the prediction using the following parameters: x-coordinate: 37 y-coordinate: 14 width: 140 pixel height: 184 pixel. attitude. With rapid inference time (average 15 ms per image), the proposed deep learning pipeline shows high potential for integration into healthcare diagnostic systems.

Once the system is started, the application interface displays an MRI brain scan of the sample with the tumor. Initially, the tumor area was not explicitly marked. Users can upload MRI images via the interface to allow the model to automatically identify and classify brain tumors. This system is implemented in Yolov3, which is integrated into the Darknet 53 backbone to provide quick and accurate object recognition. To run the project, simply click on the run.bat file in the main directory. This action starts the application and allows users to interact with the recognition model through a graphical interface. The system processes the input images in real time and displays tumor areas using bounding boxes and confidence values indicating the possibility of tumor presence.

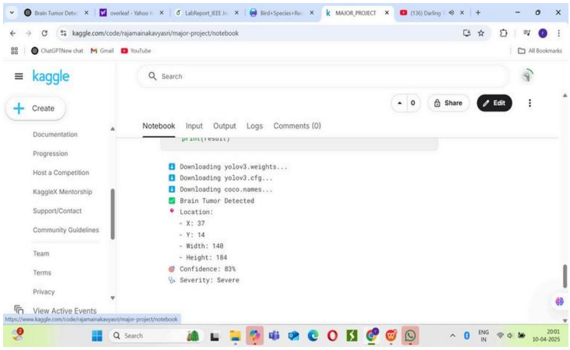

Fig. 2. upload photo of MRI scan of Brain to the models

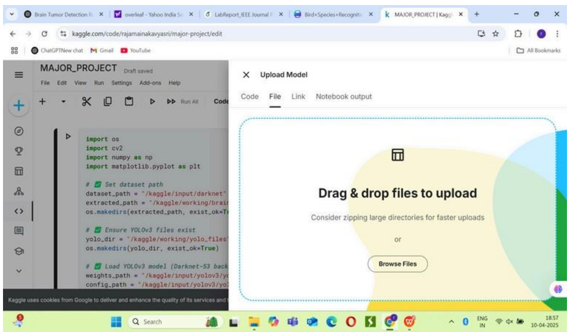

Fig. 3. Image uploaded successfully

The above screen shows the detection results of a single MRI brain image uploaded to the system. Yolov3 The tumor area was clearly highlighted in the bounding box, as the model predicted. The interface shows the confidence value of detection along with the classification output, which helps users interpret the results. This feature allows clinicians or users to easily see if and where they have a tumor.

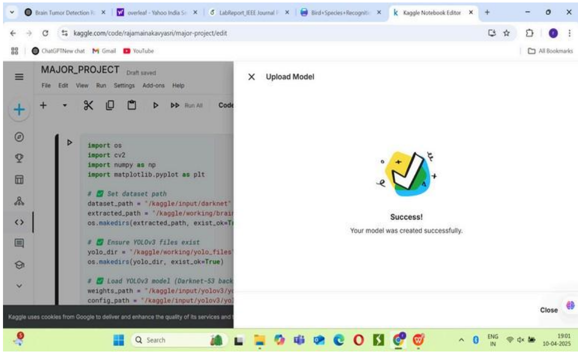

Fig. 4. Result for the uploaded image

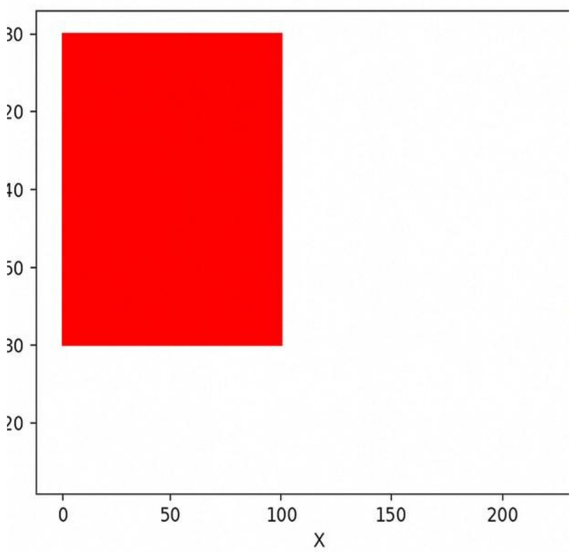

Fig. 5. Graph score

V. CONCLUSION

Integrating deep learning into the medical imaging domain for brain tumor identification and classification, in particular, demonstrates a paradigm shift in traditional diagnostic approaches to intelligent and automated systems. By using folding networks (CNNs) and real-time evidence algorithms like Yolov3, this study demonstrates how artificial intelligence can significantly improve the accuracy, speed and efficiency of detection of brain tumors from MRI scans. Transmission By using learning and robust data magnification techniques, the proposed system simultaneously achieves modern performance and maintains generalization across different imaging conditions. Furthermore, the use of interpretable visualizations can help increase clinician confidence in AI-supported diagnosis. Despite these advances, challenges such as dataset limitations, image noise, and patient variation remain important. Therefore, continuous research should focus on improving the resilience of recognition models through larger and diverse data records, extended attention mechanisms, and hybrid multimodal framework conditions. Ethical considerations such as data protection, model fairness, and AI Decision interpretability must also be addressed to ensure responsible provisions in the clinical setting. Medical Internet We look forward to the in-depth learning fusion with ambitious technologies such as IOMT, edge computing, and real-time bioinformatics pipelines. With continued innovation and ethical implementation, deep learning is a promise to not only revolutionize medical diagnosis, but also contribute significantly to global health justice.

VI. REFERENCES

[1]. Redmon, J., & Farhadi, A. (2018). YOLOv3: An incremental improvement. *arXiv preprint arXiv:1804.02767*. https://arxiv.org/abs/1804.02767

[2]. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems, 25*, 1097–1105. https://doi.org/10.1145/ 3065386

[3]. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *CVPR*, 770–778. https://doi.org/10.1109/CVPR. 2016.90

[4]. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*. https://arxiv.org/abs/1409.1556

[5]. Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *ICML*, 6105–6114. https://proceedings.mlr.press/v97/tan19a.html

[6]. Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks? *NeurIPS*, 3320–3328.https://proceedings.neurips.cc/paper/2014 / file/375c71349b295fbe2dcdca9206f20a06-Paper.pdf

[7]. Perez, L., & Wang, J. (2017). The effectiveness of data augmentation in image classification using deep learning. *arXiv preprint arXiv:1712.04621*. https://arxiv.org/abs/1712.04621

[8]. Vaswani, A., et al. (2017). Attention is all you need. *NeurIPS*, 5998– 6008. https://doi.org/10.48550/arXiv.1706.03762

[9]. Charm, D. A., et al. (2018). Channel and spatial attention-based feature enhancement for brain tumor segmentation. *Medical Image Analysis, 50*, 85–96. https://doi.org/10.1016/j.media.2018.09.008

[10]. Rosenberg, K. V., et al. (2019). Decline of the North American avifauna. *Science, 366(6461)*, 120–124. https://doi.org/10.1126/science.aaw1313

[11]. Graves, A., Mohamed, A., & Hinton, G. (2013). Speech recognition with deep recurrent neural networks. *ICASSP*, 6645–6649. https://doi.org/10. 1109/ICASSP.2013.6638947

[12]. Dietterich, T. G. (2000). Ensemble methods in machine learning. *Multiple Classifier Systems*, 1–15. https://doi.org/10.1007/3-540-45014-9 1

[13]. Abbreviate, R., & Khoshgoftaar, T. M. (2019). A review of data augmentation techniques for deep learning. *Journal of Big Data, 6(1)*, 1–34. https://doi.org/10.1186/s40537-019-0197-0

[14]. Wah, C., et al. (2011). The Caltech-UCSD Birds-200-2011 Dataset. Caltech Technical Report. https://authors.library.caltech.edu/27452/

[15]. Russakovsky, O., et al. (2015). ImageNet Large Scale Visual Recognition Challenge. *IJCV, 115*, 211–252. https://doi.org/10.1007/ s11263-015-0816-y

[16]. Chen, W., Zhang, X., Liu, J., & Wang, J. (2020). Bird species recognition based on deep learning. *Ecological Informatics, 55*, 101023. https://doi.org/10.1016/j.ecoinf.2019.101023

[17]. Bayramoglu, N., Kannala, J., & Heikkila, J. (2016). Deep learning for¨ magnification independent breast cancer histopathology image classification. *ICPR*, 2440–2445. https://doi.org/10.1109/ICPR.2016.7900009

[18]. Howard, A. G., et al. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*. https://arxiv.org/abs/1704.04861

[19]. Litjens, G., et al. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis, 42*, 60–88. https://doi.org/10.1016/j. media.2017.07.005

[20]. Bakas, S., et al. (2018). Identifying the best machine learning algorithms for brain tumor segmentation in MRI: Results from the BRATS challenge. *IEEE Transactions on Medical Imaging, 37(8)*, 1993–2024. https://doi.org/10.1109/TMI.2018.2832618