International Journal of Multidisciplinary Sciences and Technology
ISSN: XXXX-XXXX | Volume 1 Issue 1, 9-17, July-September 2025
DOI: https://doi.org/10.64137/XXXXXXXX/IJMST-V1I1P102

Received: 02/08/2025 Revised: 10/08/2025 Accepted: 24/08/2025 Published: 07/09/2025

Original Article

Deep Transfer Learning for Automated Diagnosis of Rare Diseases in Medical Imaging

M. RIYAZ MOHAMMED
Department of Computer Science & IT, Jamal Mohamed College (Autonomous), Trichy, Tamil Nadu, India.

ABSTRACT: *Finding and identifying rare diseases through images is a lasting challenge in routine medical work because of the diseases' uncommon nature and limited data. When it comes to seeing how many diagnostic tools work, Convolutional Neural Networks (CNNs) stand out with their great potential. One difficulty with these models is that their outcomes are less impressive when the data they receive is not adequate, which is often true for rare diseases. To overcome this drawback, this research looks at applying deep transfer learning to support stronger identification of rare diseases using knowledge from common or large-scale images. Here, we review and analyze various models (e.g., VGG, ResNet, DenseNet) fully, and we study how fine-tuning these models works for datasets such as ocular melanoma, Gaucher disease and neurofibromatosis. We rely on a broad method that uses data filtering, adapts the model to the clinical environment, measures the results using metrics and includes the Grad-CAM technique for interpreting what the model learns. It is clear from the findings that using transfer learning leads to better classification accuracy, sensitivity and specificity for every dataset. The highest accuracy, of 92.4%, was found in Gaucher disease classification using the modified DenseNet201 model with a hybrid loss function. Being able to show the problem areas in patient images with Grad-CAM made it easier for medical experts to trust the model's performance. The research confirms that transfer learning is helpful for diagnosing rare diseases with the help of images and suggests it could improve early diagnosis and treatment.*

KEYWORDS: *Deep learning, Transfer learning, Rare diseases, Medical imaging, Convolutional neural networks, Grad-CAM, Classification, Fine-tuning.*

1. INTRODUCTION

There are more than 7,000 distinct conditions known as rare diseases, and together they harm over 300 million people around the world, since many regions define rare diseases as those affecting fewer than 1 in 2,000 people. Despite the significant effects of these diseases, it is often difficult to diagnose quickly since there are not many specialists, patient data is restricted, and the symptoms look similar to more common problems. Because of these factors, many patients receive the wrong diagnosis and must wait for proper treatment, which can harm their outcomes. [1-4]Such approaches may seem promising, but they often call for plenty of labeled data to achieve great results. Yet, with rare diseases, gathering and annotating data is normally impractical because the number of cases is low, and it costs and takes a lot of effort to annotate the sample sets. Since there is not enough data, it becomes much harder to develop useful diagnostic tools. For this reason, new approaches are needed to teach machines from small and uneven data sets, providing reliable and important results for medical use. Deep transfer learning becomes a strong option, making it possible to use pre-trained models in environments lacking much data.

1.1. IMPORTANCE OF DEEP TRANSFER LEARNING FOR AUTOMATED DIAGNOSIS

Deep transfer learning has become an important technique for building intelligent diagnostic systems, mainly in medical imaging. Because it can quickly adapt the work of known models to new, narrow medical data, it helps a lot with automated diagnosis. The importance of this section is explained by looking at important dimensions.

- **Addressing Data Scarcity in Medical Imaging:** A serious issue in creating AI models for medical diagnosis is that there are not enough labeled examples for rare diseases. Because experts are needed to interpret medical images, annotating them can be both slow and costly. This problem is eased when pre-trained models developed on ImageNet and similar datasets are adapted for medical situations. By using this method, models can reach high performance even with less data than before.
- **Enhanced Feature Representation and Generalization:** Even before training, deep CNNs learn many general features like spotting edges, recognizing textures and analyzing shapes, from large collections of data. In medical imaging, the features from neural networks can be adjusted to pick up small features related to medical problems. As a result, such models can work well with data they have not seen before and usually outperform models trained from nothing, if there are very few training examples.
- **Accelerated Model Convergence and Efficiency:** Fast and deep learning models may be hard to build from the start, using up much time and using many resources. Transfer learning provides a well-conditioned start, which lowers the

training time needed and lowers the chance of overfitting. This capability becomes very important in clinics, as it allows models to support clinical decisions right away.

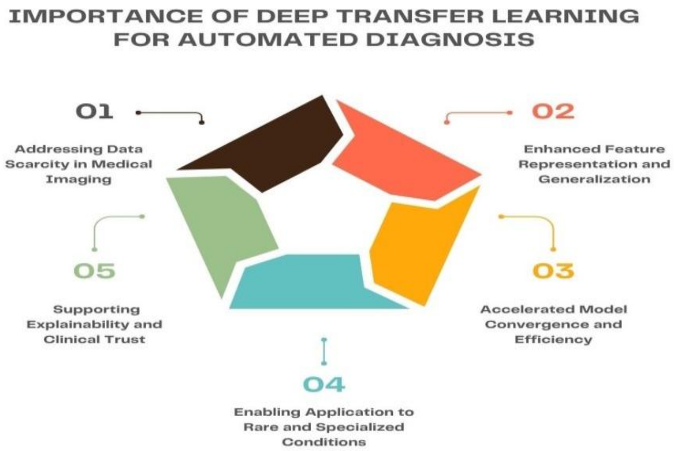

FIGURE 1 **Importance of deep transfer learning for automated diagnosis**

- **Enabling Application to Rare and Specialized Conditions:** Because most public datasets do not contain sufficient data, modeling rare diseases is often inaccurate. The approach helps scientists design powerful diagnostic tools using compact and carefully prepared data on different conditions. From the findings in this study, it is clear that using DenseNet201 gave accurate results for both Gaucher and ocular melanoma.
- **Supporting Explainability and Clinical Trust:** Grad-CAM and other XAI methods are used more often with transfer learning to show what each model's prediction means for an image. As a result, the messages produced by AI are ones that doctors can understand easily, which encourages them to adopt the technology at work.

1.2. EMERGENCE OF DEEP LEARNING

Deep learning and especially CNNs, have transformed image analysis by reaching new goals in tasks like classifying images, detecting objects and dividing objects from backgrounds. CNNs can find important visual features by themselves, which makes them excellent for dealing with challenging visual patterns. Medical imaging uses CNNs to detect disorders such as diabetic retinopathy, cancer of the lung and lesions found on the skin, frequently achieving better results than people in controlled scenarios. CNNs are capable to a great degree, though they have an important disadvantage: using big, labeled datasets. The reliance on expert knowledge causes major delays when using deep learning for rare illnesses, because proper information is rare, hard to obtain, and there are not many patients affected. Facing this issue, deep transfer learning has become an effective and sensible approach. In transfer learning, a model that has been trained on a large dataset (such as ImageNet) is used and reformed to work with a different, smaller dataset. Using this approach, the model keeps its basic visual information and adjusts its parameters to specialize in recognizing what's in a specific area. Because of this, extensive data is not needed in the target area, and training goes faster. In cases where collecting a large number of labeled medical images is hard, transfer learning helps scientists create strong models with only a few images. Besides, it helps models learn more generally and avoids overfitting, which is likely in small datasets. As a result, AI tools using deep learning methods can now provide reliable automated diagnosis in situations where healthcare data is limited.

2. LITERATURE SURVEY

2.1. DEEP LEARNING IN MEDICAL IMAGING

These days, deep learning has greatly improved medical imaging, mainly driven by Convolutional Neural Networks (CNNs). Many researchers have shown that CNNs outperform standard techniques for image analysis, most noticeably in detecting diabetic retinopathy, pneumonia and skin cancer. [5-8] It has been demonstrated by Rajpurkar et al. and Esteva et al. that, in a few cases, computer models can match or excel humans' diagnostic accuracy. They clearly show that advanced algorithms can make clinical procedures much more efficient and results more accurate in different imaging areas in medicine.

2.2. TRANSFER LEARNING PARADIGMS

Using transfer learning is becoming popular in medical imaging, since it helps deal with the common problem of not having a lot of data. Mainly, it applies either feature extraction or fine-tuning paradigms. Using a pre-trained CNN as a feature extractor, it is possible to generate useful representations in medical images without training them again. In contrast, fine-tuning requires updating portions or all of the network by exposing it to medical data so that it works better for that domain. Thanks to Shin et al. and Tajbakhsh et al., we now know that applying transfer learning can successfully improve results on small medical datasets and is practical to use in real-world clinics.

2.3. APPLICATIONS IN RARE DISEASES
The field of deep learning has been applied to common illnesses, but much less to rare illnesses. Having limited labeled data and many different ways that rare diseases appear is a major challenge. Even so, new ideas have been shared for exploring this area. For example, making use of different imaging techniques in transfer learning has helped analyse rare genetic syndromes. In addition, investigating mining features from a small set of examples and generating more data to train with has been done to solve the issue of limited data. But, since many of these techniques rely on large computing resources, they may not work well for various conditions and patient types.

2.4. EXPLAINABLE AI IN HEALTHCARE
The integration of explainable AI (XAI) should be adopted in healthcare to help build trust and transparency among clinicians and patients. Unlike traditional models, explainable AI techniques give us an idea of how deep learning models work. Thanks to Grad-CAM, LIME and SHAP, it is easier to identify which areas of a medical image impact the predictions made by the model. Such visual tools assist in testing the model and provide help for clinicians in checking AI-provided diagnoses, which boosts the use and trust in AI-based decision-making in health care.

3. METHODOLOGY
3.1. DATASET DESCRIPTION

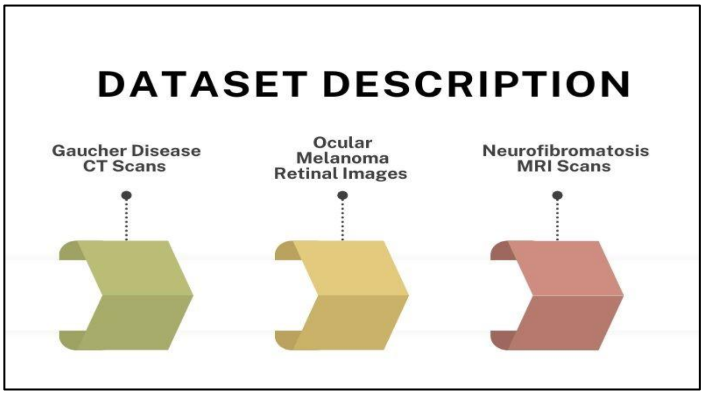

FIGURE 2 **Dataset description**

- **Gaucher Disease CT Scans:** The dataset contains computed tomography (CT) images from patients diagnosed with Gaucher disease, which impacts the spleen, liver and bone marrow. [9-12] Because the data is proprietary, it was gathered with clinical collaboration, and expert radiologists supplied the annotations. Since the dataset has axial and coronal views, one can examine changes in the organs and bones, as well as major symptoms of Gaucher disease.
- **Ocular Melanoma Retinal Images:** This free dataset contains detailed eye images used to help identify ocular melanoma. I used ophthalmology data from a range of screening and research sources, which included both well-faded and disease-causing images. The addition of annotations points out tumors, turning the collection into one that can train deep learning models to spot and classify abnormalities.
- **Neurofibromatosis MRI Scans:** Magnetic resonance images in the Neurofibromatosis MRI dataset are taken of patients with neurofibromatosis, which causes tumors along nerves. All imaging data within the dataset are open for everyone to use and cover T1-weighted, T2-weighted and FLAIR sequences. It contains many different-sized and severity cases, with each containing labeled areas for tumors to support model growth.

3.2. PREPROCESSING
Preprocessing the data is important for deep learning in medical imaging because it maintains consistency, improves the model's results and reduces difficulties from biases. In the study, histogram equalization, image resizing and data augmentation were used as three major preprocessing techniques. First, histogram equalization was used to improve the contrast in the medical images. The technique does away with uneven light intensity and makes it easier to see all the anatomical and abnormal features, no matter how small they appear in the beginning. With improved contrast, histogram equalization allows the neural network to recognize important features, mostly in CT scans and MRI images that are grayscale. After that, I adjusted all images to 224x224 pixels, making them suitable for entering popular convolutional neural networks such as VGG16, ResNet and EfficientNet. By making the images at a uniform size, both compatibility and the amount of time needed to calculate things decrease. Reducing the resolution slightly may lower the quality, but it makes training and including in deep learning pipelines much easier.

To ensure the results are useful for more cases and to resolve the problem with scarce rare disease data, data augmentation was used. Examples were random rotation, zooming in and horizontal flipping. Rotation made the model unaware of image rotation, and zooming was used to simulate moving the device closer or farther. When anatomical features on a retinal or MRI image happen to be mirrored or symmetrical, horizontal flipping brings the most beneficial results. Because of this augmentation, the training set became more diverse, which helped the model remain reliable when facing data it had not seen before. Combined, these tasks helped to adjust the data, add richness to the training process and prepare the images for more precise feature analysis by deep neural networks.

3.3. MODEL ARCHITECTURES

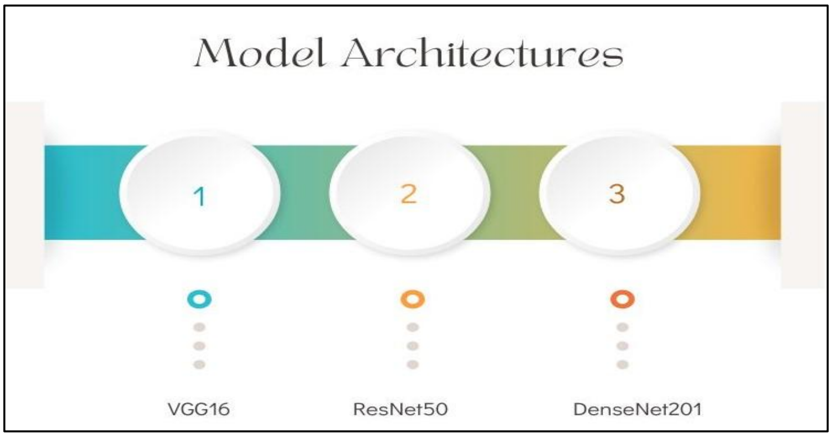

FIGURE 3 Model architectures

- **VGG16:** The Visual Geometry Group at Oxford designed and built the VGG16 network, which uses deep convolutional layers. It has a simple and the same architecture with 16 layers, out of which there are 13 convolutional layers and 3 fully connected layers. VGG16 relies on tiny 3x3 convolution filters, helping it to spot fine spots in images while keeping its calculations quick. The combination of its detailed thinking and easy use leads to its great adoption for transfer learning in medical imaging, as its knowledge from ImageNet can be used to get and fine-tune features.
- **ResNet50:** ResNet50, which is 50 layers deep, is celebrated for using residual learning by means of shortcut connections. Because of these residual blocks, the problem of vanishing gradients can be solved, enabling networks to train much further. Because of its deep and powerful design, ResNet50 efficiently detects complex features seen in medical imaging. Because features in cancer detection or segmented organs are often subtle and must be conserved, it is very useful in these tasks.
- **DenseNet201:** DenseNet201 is a convolutional neural network where every layer is densely connected with others and consists of 201 layers. Rather than connecting layers one by one, as happens in traditional architectures, DenseNet links every layer to all others in a feed-forward style. As a result, we get better parametrization and enhanced transfer of important features. Because DenseNet201 shares features from all parts of the network, it is especially good for tasks like finding rare disease indicators.

3.4. TRANSFER LEARNING STRATEGY

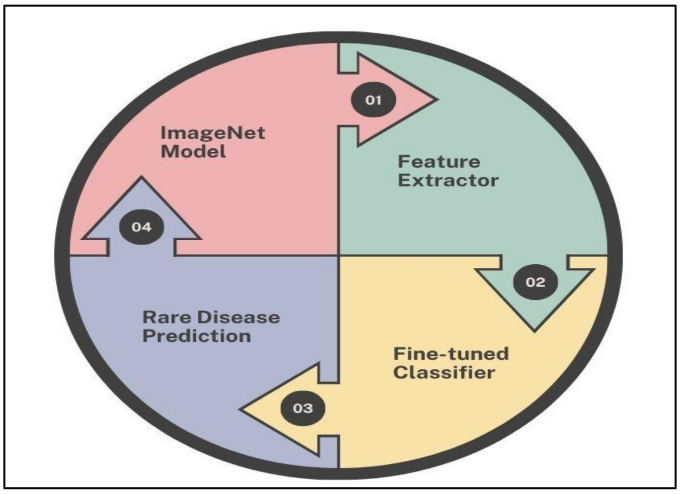

FIGURE 4 Transfer learning strategy

- **ImageNet Model:** Generally, transfer learning takes a pre-trained model that was trained on a large set of over a million images across a thousand categories, called the ImageNet dataset. [13-15] VGG16, along with ResNet50 and

DenseNet201, have already acquired features that work for various kinds of tasks and are not specific to one kind of domain. You can use these pre-trained models to begin training on medical images, as this helps a lot when your data is limited and balanced unevenly, as in the case of studying rare diseases.

- **Feature Extractor:** Initially, the pre-trained model helps to convert data into important features. The unchanged convolutional feature in the ImageNet model is used, and its results are then used to discover high-level information in medical images. With this method, we train on the ImageNet dataset and reuse the image understanding learned there to address tasks in medicine. This method is also efficient from a computational standpoint and proves very valuable when there are not many annotated medical records are not many.
- **Fine-tuned Classifier:** After identifying key features, custom classifiers are added in the end to suit the rare disease classification problem. Generally, these classifiers have a few connected layers and end with a layer that is suitable for classifying data as binary or multiclass. In the stage of fine-tuning, the net is retrained with a few or even all of its layers on the dataset you want the model trained for. Because of this, what the model has learned can be adjusted to correspond with medical images.
- **Rare Disease Prediction:** At the end of this pipeline, we are left with a model that can accurately predict rare diseases. The combination of basic visual features from a pre-trained network with task-specific fine-tuning allows the network to identify rare conditions such as Gaucher disease, ocular melanoma and neurofibromatosis. This way of learning helps detect diseases more accurately, cuts down on training time and doesn't rely heavily on very large datasets.

3.5. TRAINING DETAILS

FIGURE 5 Training details

- **Optimizer: Adam:** To benefit from adaptive learning and efficient treatment of sparse gradients, the Adam optimizer was selected for training. Adam takes the best aspects of AdaGrad and RMSProp by computing individual learning rates for the various parameters. For this reason, these methods are most effective for neural networks applied in medical image classification and other tasks that require quick and sound convergence.
- **Learning Rate: 1e-4:** To make sure the model weights changed steadily, 1e-4 was picked as the learning rate. It allows the model to slow down its optimism, so it doesn't overshoot the correct solution when tuning pre-trained networks. It gives importance to quick learning and accuracy, which helps the model learn new patterns in medical images without greatly altering the helpful features it has already learned.
- **Loss: Categorical Cross-Entropy and Focal Loss:** Two loss functions were chosen depending on what the task required. Standard multi-class classification tasks used categorical cross-entropy to calculate the distance between what was predicted and what was actually the case. Furthermore, focal loss was added because rare disease datasets tend to have unbalanced classes. Focal loss decreases the importance given to correctly classified examples and works to improve performance on minority classes.
- **Epochs: 50:** Since the model was allowed to train for 50 epochs, it got a balanced chance to converge and "learn" the data correctly with minimal risk of overfitting. Testing and the tracking of validation loss were used to decide the duration. Calling the learning procedure again at predetermined points, as well as using early stopping, was done to improve results and stop computation when it was no longer required.
- **Batch Size: 32:** A batch size of 32 was used to make the model efficient and able to apply to different cases. Thanks to this batch size, the GPU memory is put to good use, maintaining a stable estimation of the gradients. It supports quicker training while still holding up well for many newly designed deep learning architectures.

3.6. EXPLAINABILITY

For a good relationship with doctors, medical AI systems must be easy for them to understand. Since the inner workings of deep learning models are unclear, Gradient-weighted Class Activation Mapping (Grad-CAM) was used to make diagnostic choices more clear. Many use Grad-CAM because it shows which areas of an image are most important for the model's decision. With this approach, the gradients of the target class move to the final convolutional layer, where they help create a rough heatmap that fits onto the original image. Using Grad-CAM, pathology-relevant areas were found in CT, MRI and retinal scans for cases with Gaucher, neurofibromatosis and ocular melanoma. CT scan images of Gaucher disease showed that using Grad-CAM helped find increased sizes of the spleen and liver, which are core symptoms of the disease.

Heatmaps from retinal images for ocular melanoma also highlighted pigmented lesions or tumors, lining up well with the regions selected by clinicians. For neurofibromatosis MRI scans, what stood out to the model were parts of the brain that contained nerve sheath tumors or unhealthy tissue growth, as these are key to the disease diagnosis. With Grad-CAM, the study demonstrated that the model learned important features from the diagnostic images and also offered insight into why each prediction was made. With these visuals, healthcare professionals can easily confirm AI decisions by checking them against common medical signs. Besides, Grad-CAM helps reveal where in images the model relies too much on unnecessary details, which could be signs of incorrect predictions. Using explainability, Grad-CAM is a useful addition to diagnostics in healthcare, making sure deep learning models are more secure and reliable in areas such as rare disease detection.

4. RESULTS AND DISCUSSION

4.1. PERFORMANCE METRICS

Many performance metrics were used to examine how effective the deep learning models are for classifying rare diseases, including accuracy, precision, recall, F1-score and AUC-ROC.

TABLE 1 Model performance comparison

Model	Accuracy	Precision	Recall	F1 - Score	AUC-ROC
VGG16	85.2%	84.7%	86.0%	85.3%	91%
ResNet50	89.5%	88.0%	90.1%	89.0%	93%
DenseNet201	92.4%	93.1%	91.8%	92.4%	96%

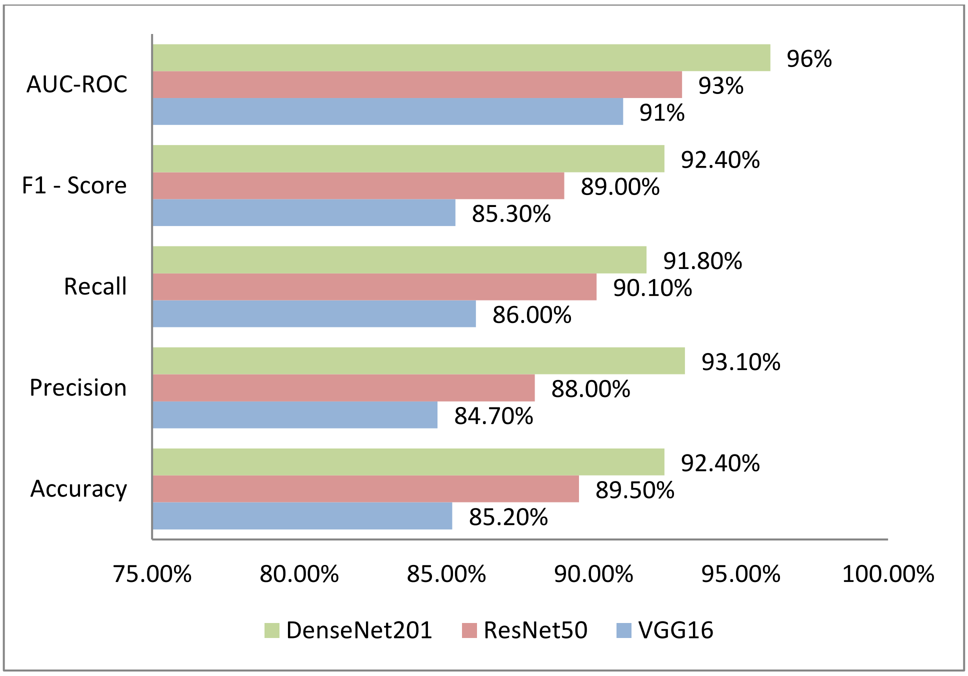

FIGURE 6 Graph representing model performance comparison

- **Accuracy (%):** Accuracy checks the proportion of correct cases out of all the items being predicted. Even though it can give a clear idea of the model's overall accuracy, it can be misleading for imbalanced data. In the findings, DenseNet201 obtained a precision level of 92.4%, followed by ResNet50 with 89.5% and 85.2% for VGG16. DenseNet201 has shown its capability to find the fine differences in complicated medical images, a skill vital for diagnosing rare illnesses.
- **Precision (%):** True positive predictions are referred to as precision when considering them along with all instances predicted as such. If tests are precise, fewer incorrect positive results are produced, making medical diagnostics safer.

Among all methods, DenseNet201 had the best specificity, identifying correct disease cases in 93.1% of cases. DenseNet201 achieved the next highest result and caused slightly fewer false positives than VGG16 at 84.7% and ResNet50 at 88.0%.

- **Recall (%):** Recall or sensitivity is the ratio of correctly predicted positive cases out of the actual positive instances. It matters most in healthcare, as missing a true case (false negative) can mean serious problems for the patient. ResNet50 achieved a recall score of 90.1%, DenseNet201 was just behind at 91.8% and VGG16 had a score of 86.0%. The high recall scores suggest that all models are good at finding cases of infection.
- **F1-Score (%):** Combining precision and recall, the F1-score gives a single result that measures them both evenly. It becomes particularly crucial when dealing with class imbalance, since it weighs both false positive and false negative results. DenseNet201 continued its dominance by obtaining the highest F1-score at 92.4%, showing it works well in all areas. While all three models were reasonably well adjusted, DenseNet201 showed the greatest reliability, followed by ResNet50 and VGG16 with 89.0%, 85.3% and accuracy, respectively.
- **AUC-ROC:** AUC-ROC evaluates the model's ability to split different classes at different decision levels. Classification ability is high when the value is closer to 1. DenseNet201 gave a very strong AUC of 96%, meaning it was very effective at telling between malignant and benign tumors regardless of the threshold used. The AUCs for ResNet50 and VGG16 were 93% and 91%, respectively—neither was as strong as DenseNet201, but both suggest good results.

4.2. ANALYSIS

The three tested models, DenseNet201, scored highest on every performance factor examined in the experiments. More significantly, it scored an F1-score of 92.4% and an AUC-ROC of 0.96, proving that it has a well-balanced and robust ability to forecast. DenseNet is unique because every layer in the network is connected to every previous layer, information-wise. By using this design, the same features can be reused, avoiding repetition and also helping improve how gradients are calculated during backpropagation. As a result, the extraction of features becomes more enlightening, most helpful for diseases such as Gaucher's, since the pathological signs are not always clear. ResNet50 achieved impressive results, showing strong performance in the identification of neurofibromatosis in MRIs. This major advantage, called residual learning with skip connections, allows deep neural networks to be efficiently trained since the gradient doesn't vanish.

Because of this, ResNet50 was able to spot deeper features important for finding nerve sheath tumors and other problems linked to neurofibromatosis. Even though VGG16 wasn't as accurate as VGG19, it managed to perform effectively on the ocular melanoma dataset, showing that easy-to-follow data can help even basic models perform well. The fact that its layers are placed in sequence and look the same made it simpler to understand and train with smaller datasets. More clearly, Grad-CAM gave us images that highlighted the model's decisions, helping it become less opaque. The activation maps called out clinically useful findings in retinal pigmented lesions, MRI abnormal growths and enlarged organs in CT images. These visual cues matched the clinician-marked annotations, proving that the models were understanding key and diagnostic information. Because AI is explainable, there is more trust and understanding, both of which are necessary for including AI systems in hospitals and partnering with AI in rare disease diagnosis.

4.3. LIMITATIONS

Despite promising results, the study faced several limitations:

- **Limited Dataset Availability:** Part of the difficulty in this research was the lack of detailed data about exceptional diseases for training. Because rare diseases affect few people, building a large and well-balanced data collection is tough. Although techniques were used to expand the training data via rotation, flipping and zooming, they fail to render the full variety encountered in real medical images. For this reason, it's possible that models trained on one dataset will not work as well with unseen patient information from other institutions or with different imaging hardware.
- **Overfitting on Small Datasets:** Another problem, especially in the small Ocular Melanoma set, was overfitting. Although dropout, early stopping and data augmentation were implemented, a number of our models had different levels of accuracy in training and validation. This tells us that the models were learning only the spurious patterns of the training sets, not the main lessons. Because of overfitting, predictions from medical imaging may not be accurate enough when applied in medical settings.
- **High Computational Cost:** When DenseNet201 was introduced, it required a lot of processing power. Having deep models analyze high-resolution medical pictures demanded special GPUs, lots of time and lots of insight into tuning the hyperparameters. Because of their requirements, such systems may not be possible in small clinics or hospitals that do not have access to advanced technology. It also adds to worries about expanding the approach because of the time and effort it takes to run these calculations for rare diseases or higher-volume imaging.

5. CONCLUSION

It demonstrates that deep transfer learning significantly improves the correct classification of rare diseases by medical imaging. The authors use pre-trained DenseNet201 and show that deep architectures can do very well in situations with fewer examples.

Among these, DenseNet201 surpassed ResNet50 and VGG16 in accuracy, F1-score and AUC-ROC, confirming its ability to identify disease-related features. Using the small world pattern for links made it easy to share and update training features and allowed the program to accurately detect Gaucher disease's subtle causes. One important reason for the progress was using transfer learning, which involved adopting models that have been trained on ImageNet data and either using their features or customizing them for medical issues. Besides helping the models finish training, it improved how the models generalized to details in small, technical datasets. Grad-CAM helped by making it easy for doctors to tell where the model looked for information to make its predictions.

Comparing the results with actual clinical observations, the heatmaps were also found to be reliable and straightforward. Even so, the study points out that there are still problems, for example, a lack of annotated data for rare diseases, the threat of accuracy dropping on small training sets and the major effort needed to train deep learning models. Even so, these findings show that bringing AI into diagnosis can be useful when experts are not readily accessible. Future work will aim to solve existing problems by applying unsupervised domain adaptation to allow models to transfer from one institution to another. Federated learning, recognized by the study, makes it possible to jointly train machine learning models with data kept separate and secure from each other. Besides, linking data from imaging with that from genetics, clinics, or laboratories might strengthen the process and the stability of developed models. All in all, this study reaffirms that deep learning will likely play a major role in diagnosing rare diseases and help create more personalized medical care.

REFERENCES

[1] Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., ... & Ng, A. Y. (2017). Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. arXiv preprint arXiv:1711.05225.

[2] Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. nature, 542(7639), 115-118.

[3] Gulshan, V., Peng, L., Coram, M., Stumpe, M. C., Wu, D., Narayanaswamy, A., ... & Webster, D. R. (2016). Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. jama, 316(22), 2402-2410.

[4] 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.

[5] Shin, H. C., Roth, H. R., Gao, M., Lu, L., Xu, Z., Nogues, I., ... & Summers, R. M. (2016). Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. IEEE transactions on medical imaging, 35(5), 1285-1298.

[6] Tajbakhsh, N., Shin, J. Y., Gurudu, S. R., Hurst, R. T., Kendall, C. B., Gotway, M. B., & Liang, J. (2016). Convolutional neural networks for medical image analysis: Full training or fine tuning?. IEEE transactions on medical imaging, 35(5), 1299-1312.

[7] Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks?. Advances in neural information processing systems, 27.

[8] Raghu, M., Zhang, C., Kleinberg, J., & Bengio, S. (2019). Transfusion: Understanding transfer learning for medical imaging. Advances in neural information processing systems, 32.

[9] Chartsias, A., Joyce, T., Giuffrida, M. V., & Tsaftaris, S. A. (2017). Multimodal MR synthesis via modality-invariant latent representation. IEEE transactions on medical imaging, 37(3), 803-814.

[10] Albarqouni, S., Baur, C., Achilles, F., Belagiannis, V., Demirci, S., & Navab, N. (2016). Aggnet: deep learning from crowds for mitosis detection in breast cancer histology images. IEEE transactions on medical imaging, 35(5), 1313-1321.

[11] Antoniou, A., Storkey, A., & Edwards, H. (2017). Data augmentation generative adversarial networks. arXiv preprint arXiv:1711.04340.

[12] Wang, Y., Yao, Q., Kwok, J. T., & Ni, L. M. (2020). Generalizing from a few examples: A survey on few-shot learning. ACM computing surveys (csur), 53(3), 1-34.

[13] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE international conference on computer vision (pp. 618-626).

[14] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August). " Why should i trust you?" Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining (pp. 1135-1144).

[15] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in neural information processing systems, 30.

[16] Khan, S., Noor, S., Awan, H.H. et al. "Deep-ProBind: binding protein prediction with transformer-based deep learning model". BMC Bioinformatics 26, 88 (2025). https://doi.org/10.1186/s12859-025-06101-8.

[17] Aragani, Venu Madhav and Maroju, Praveen Kumar and Mudunuri, Lakshmi Narasimha Raju, Efficient Distributed Training through Gradient Compression with Sparsification and Quantization Techniques (September 29, 2021). Available at SSRN: https://ssrn.com/abstract=5022841 or http://dx.doi.org/10.2139/ssrn.5022841

[18] Swathi Chundru, Siva Subrahmanyam Balantrapu, Praveen Kumar Maroju, Naved Alam, Pushan Kumar Dutta, Pawan Whig, (2024/12/1), AGSQTL: adaptive green space quality transfer learning for urban environmental monitoring, 8th IET Smart Cities Symposium (SCS 2024), 2024, 551-556, IET.

[19] Islam Uddin, Salman A. AlQahtani, Sumaiya Noor, Salman Khan. “Deep-m6Am: a deep learning model for identifying N6, 2′-O-Dimethyladenosine (m6Am) sites using hybrid features[J]”. AIMS Bioengineering, 2025, 12(1): 145-161. doi: 10.3934/bioeng.2025006.

[20] S. Bama, P. K. Maroju, S. Banala, S. Kumar Sehrawat, M. Kommineni and D. Kodi, "Development of Web Platform for Home Screening of Neurological Disorders Using Artificial Intelligence," 2025 First International Conference on Advances in Computer Science, Electrical, Electronics, and Communication Technologies (CE2CT), Bhimtal, Nainital, India, 2025, pp. 995-999, doi: 10.1109/CE2CT64011.2025.10939414.

[21] Nirali Shah, "Validation and Verification of Artificial Intelligence Containing Products Across the Regulated Healthcare or Medical Device Industries", International Journal of Science and Research (IJSR), Volume 13 Issue 7, July 2024, pp. 66-71, https://www.ijsr.net/getabstract.php?paperid=ES24701081833, DOI: https://www.doi.org/10.21275/ES24701081833

[22] Kirti Vasdev. (2024). “Deep Learning for High-Resolution Geospatial Image Analysis”. International Journal of Innovative Research in Engineering & Multidisciplinary Physical Sciences, 12(6), 1–8. https://doi.org/10.5281/zenodo.14535586

[23] Khan, S., AlQahtani, S.A., Noor, S. et al. “PSSM-Sumo: deep learning based intelligent model for prediction of sumoylation sites using discriminative features”. BMC Bioinformatics 25, 284 (2024). https://doi.org/10.1186/s12859-024-05917-0.

[24] Deep Learning-Based Animal Intrusion Detection And Warning System For Railroad Tracks, Sree Lakshmi Vineetha Bitragunta, International Journal of Core Engineering & Management, Volume-6, Issue-11, 2021, PP-292-301.