International Journal of Social Science and Economic Research

ISSN: 2455-8834

Volume:09, Issue:09 "September 2024"

# Virtual Doctor: Enhancing Global Healthcare Accessibility through Image Classification

Vansh Porwal, Sai Sarvanth Vedula, Taruna Agarwal, John Leddo, Dev Agarwal, Ishwarya Ramineni, Dhanvi Kodali, Aryan Sajjad, Arsal Sajjad, Bhargav Subash, Samir Bahl, Deepika Ravi, Dharana Cherukuwada, Collin O'Brien, Sambhav Jain, Michelle Kim, Krish Tripathy, Avyan Illapakurthy, Arnav Illapakurthy, Ryan Park, Yousef Malik, Prardhana Joy Thumma, Sreeja Sambhuni, Eshwar Sai Dokku, Yudhir Vassam, Krisha Dotiwalla, Ryan Jin, Veda Rangareddyvari, Sanvi Vandanapu, Helena Gabrial, Dheenav Rallabandi, Varsha Mupparaju, Aadiv Aitharaju, Srikaran Yelimati, Zhiqing Zhang, Jishnu Bachala, Dhruv Shrivastava, David Nam, Shreyas Verma, Vihaan Udupa, Aarmaan Pahuja, Sarayu Kadasani, Shourya Jha, Aditi Chalasani, Amala Elangovan, Nicole Camacho-Pinto, Kevin Camacho-Pinto, Gabriel Chavarria, Dhruvi Patel, Nivrita Cherukuwada

MyEdMaster, LLC
Leesburg, Virginia, USA

DOI: 10.46609/IJSSER.2024.v09i09.050 URL: https://doi.org/10.46609/IJSSER.2024.v09i09.050

Received: 15 September 2024 / Accepted: 28 September 2024 / Published: 5 October 2024

## ABSTRACT

*Virtual Doctor aims to make healthcare more available worldwide by using image classification to detect diseases from pictures. Virtual Doctor represents a ground-breaking initiative aimed at transforming health assessments worldwide, particularly in areas where healthcare resources are scarce. Utilizing a sophisticated machine learning model, this platform allows users to perform real-time health evaluations by simply capturing images using a smartphone or other camera-equipped mobile devices. Virtual Doctor empowers individuals to conduct self-assessments and receive immediate feedback on their health, facilitating early detection and preventive care. This innovative approach not only educates users about proper hygiene practices but also significantly enhances access to healthcare services. By simplifying the process of health monitoring, Virtual Doctor seeks to improve overall health outcomes in underserved populations. This report will focus on the Dental Health module of the platform, illustrating how the underlying technology can be adapted to other health domains such as facial health, each tailored to meet specific diagnostic needs.*

  

**1. Problem:**

In many regions worldwide, access to oral healthcare resources is severely limited, with the World Health Organization (WHO) estimating that approximately 50% of the human population suffer from untreated oral diseases. The scarcity of dental care facilities means that large segments of the population lack regular access to dental check-ups and screenings. Consequently, many individuals remain unaware of developing oral health issues until they reach an advanced stage, leading to more severe complications. This lack of timely intervention exacerbates the burden of oral diseases and contributes to overall public health challenges, particularly in underserved communities where healthcare resources are scarce.

**2. Challenge:**

Globally, access to oral healthcare remains critically limited, with estimates from the World Health Organization (WHO) indicating that about 50% of the world's population suffers from untreated oral diseases. The shortage of dental care facilities results in many individuals, especially in underserved areas, not having regular access to dental check-ups and screenings. Consequently, numerous people remain oblivious to developing oral health issues until they become severe, leading to more complex health complications. This delay in detection and treatment not only worsens the oral disease burden but also poses significant public health challenges in regions already struggling with limited healthcare resources.

**3. Proposed method**

This section outlines the systematic approach adopted to develop the Virtual Doctor platform, specifically focused on oral health assessments. We detail the advanced machine learning model, starting with the data collection of diverse oral health images, followed by rigorous preprocessing and data augmentation techniques to ensure consistency and quality. We also discuss the selection and architectural design of the ResNet18 model, tailored for high accuracy in image recognition tasks. This section encapsulates the entire workflow from model training to iterative refinement, highlighting the methods used to overcome challenges related to data scarcity and computational limitations.

  

ISSN: 2455-8834

Volume:09, Issue:09 "September 2024"

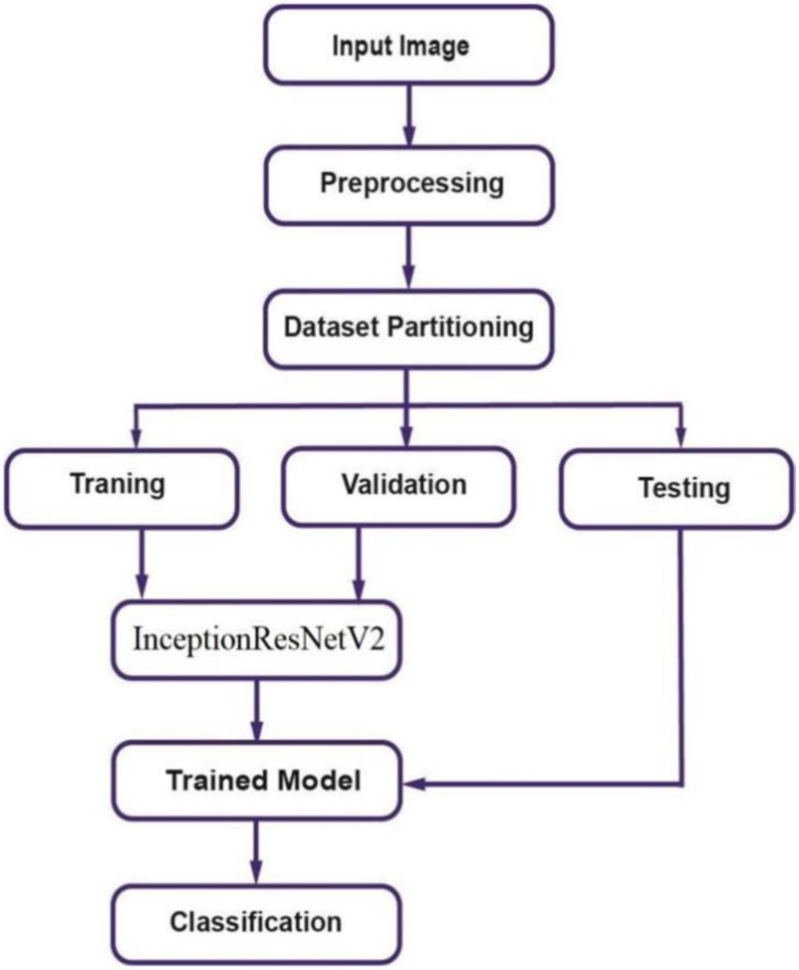


### 4.1. Dataset

The foundation of the Virtual Doctor platform's machine learning model is a meticulously compiled dataset of oral health images. This dataset comprises images sourced from multiple online databases, offering a diverse representation of various dental conditions essential for training our model effectively.

#### Image Distribution among Classes

Our dataset includes a total of 11,252 images, categorized into six primary dental conditions and a healthy class, as detailed below:

- **Caries:** 2,382 images

  

International Journal of Social Science and Economic Research

ISSN: 2455-8834

Volume:09, Issue:09 "September 2024"

- **Gingivitis:** 2,349 images
- **Tooth Discoloration:** 2,541 images
- **Ulcers:** 1,834 images
- **Hypodontia:** 1,251 images
- **Calculus:** 1,296 images
- **Healthy Teeth:** 200 images

The chart below illustrates the distribution of images across these classes, highlighting the imbalance particularly noticeable in the "Healthy" category, which is significantly underrepresented with only 200 images. This skewness poses challenges for training but also reflects common real-world scenarios where healthy dental images are less frequently documented.

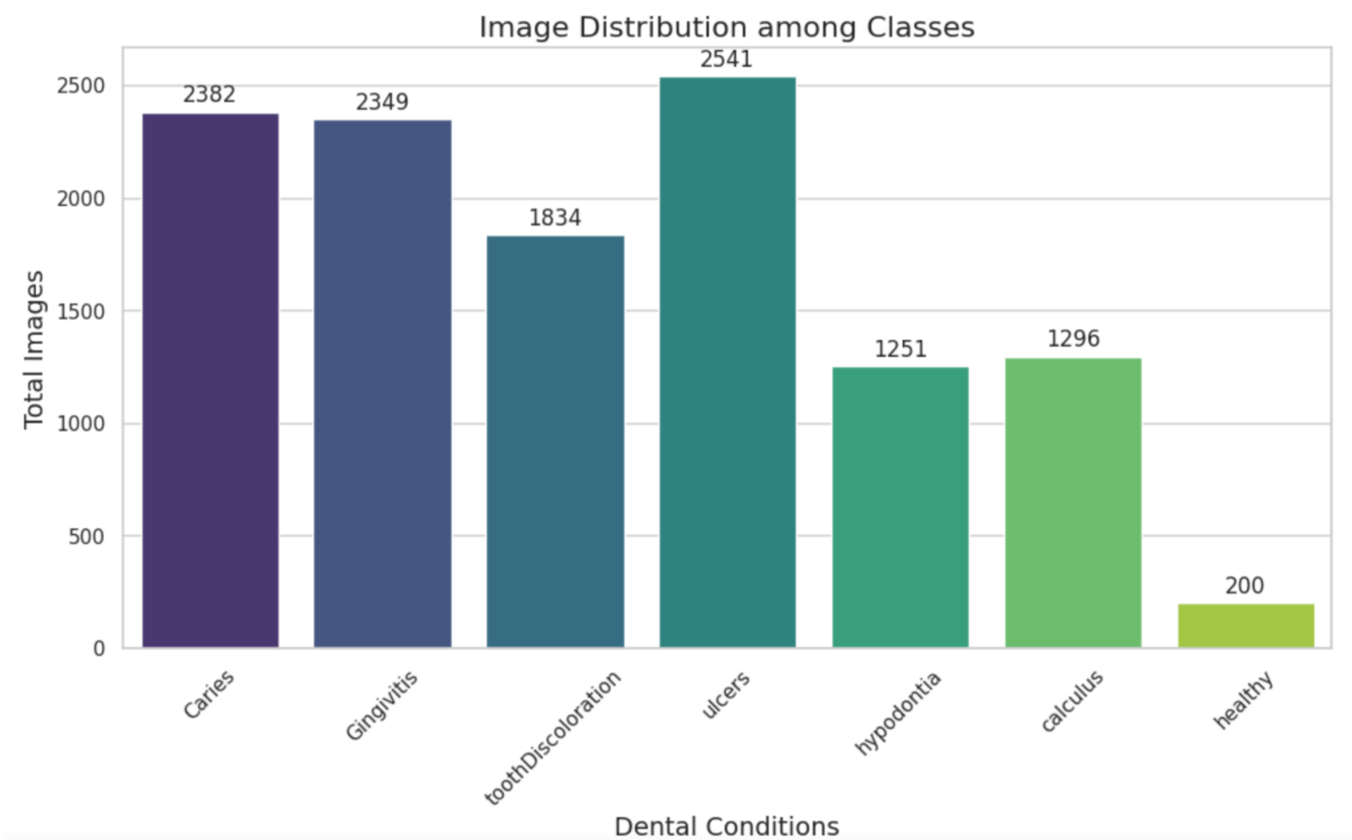


  

ISSN: 2455-8834

Volume:09, Issue:09 "September 2024"

**Data Augmentation**

To address the variability in image quality and the imbalance in class distribution, we employed data augmentation techniques. These include random resizing, horizontal flipping, rotation, and normalization. The transformation parameters for training, validation, and testing sets are as follows:

- **Training Transforms:** Random resizing (80-100% scale), horizontal flipping, 10-degree rotation, conversion to tensor, and normalization (mean: [0.485, 0.456, 0.406], std: [0.229, 0.224, 0.225]).
- **Validation and Testing Transforms:** Resizing to maintain scale, center cropping, conversion to tensor, and normalization using the same mean and standard deviation as training.

These pre-processing steps ensure our model is robust, reducing overfitting and improving its ability to generalize across unseen data.


**4.2. Model Selection and Training**

For the task of classifying dental conditions from images, we selected the ResNet18 architecture, a decision influenced by its proven track record in image recognition tasks. ResNet18 is particularly suited for our needs due to its balance between complexity and performance, capable of achieving high accuracy without the computational heaviness of deeper networks. This

  

ISSN: 2455-8834

Volume:09, Issue:09 "September 2024"

architecture incorporates several CNN layers for feature extraction, followed by pooling layers that perform spatial downsampling, and fully connected layers that culminate in classification.

**Training Process**

The training of our model was conducted in a structured manner:

1. **Dataset Division:** The comprehensive dataset was divided into training, validation, and test sets. This separation is crucial for training robust models and preventing overfitting.

2. **Backpropagation and Gradient Descent:** The ResNet18 model was trained on the training data using backpropagation and gradient descent techniques. These methods are essential for minimizing the loss function and iteratively updating the model parameters.

3. **Performance Monitoring:** Throughout the training process, the model's performance was continuously monitored on the validation set using the learning curve. This monitoring helps in making necessary adjustments to prevent overfitting and ensures that the model generalizes well to new, unseen data.

**Hyperparameter Tuning**

We fine-tuned several hyperparameters to optimize model training and performance. Key hyperparameters adjusted include:

- **Learning Rate:** Critical for determining the step size at each iteration while moving toward a minimum of a loss function.

- **Batch Size:** Influences model accuracy and training speed.

- **Dropout Rate:** Utilized to enhance the generalization of the model.

These parameters were optimized through a series of experiments and validation steps to ensure the best possible performance of our model on the dataset.

**Evaluation and Validation**

Post-training, the model underwent rigorous evaluation and validation:

- **Performance Metrics:** The trained model was assessed on the test set, evaluating its accuracy, precision, recall, and other relevant metrics. These metrics provide insight into the model's effectiveness in accurately classifying dental conditions.

  

ISSN: 2455-8834

Volume:09, Issue:09 "September 2024"

- **Cross-validation:** To test the robustness of the model against different subsets of data, cross-validation was performed. This method helps in understanding the variability and reliability of the model predictions.

- **Iterative Refinement:** Based on the outcomes of these evaluations, the model was iteratively refined. This refinement involved adjusting the hyperparameters, augmenting the dataset, or modifying the architectural elements as needed to boost performance and address identified shortcomings.

Through this meticulous process of selection, training, hyperparameter tuning, and validation, we developed a robust machine learning image classification model that met our objectives. This approach not only ensured high accuracy and generalization but also provided a clear understanding of the model's capabilities and limitations when applied to real-world dental imaging data.

## 5. Results

### Evaluation Metrics

The performance of the Virtual Doctor model on the test set demonstrates strong predictive accuracy, achieving an overall accuracy of 92%. The classification report details the precision, recall, and F1-score for each dental condition, summarized as follows:

- **Calculus:** Achieved a precision of 0.73 and recall of 0.72, resulting in an F1-score of 0.72.

- **Caries:** Exhibited excellent precision and recall, both at or near 1.00, yielding a near-perfect F1-score of 0.99.

- **Gingivitis:** Recorded a precision of 0.82 and a recall of 0.85, with an F1-score of 0.84.

- **Healthy Teeth:** Showed high precision of 0.97 and a recall of 0.85, leading to an F1-score of 0.91.

- **Hypodontia:** Both precision and recall were 0.89, with a corresponding F1-score of 0.89.

- **Tooth Discoloration:** Demonstrated very high precision and recall, both at 0.98, achieving an F1-score of 0.98.

- **Ulcers:** Reached perfect precision and nearly perfect recall, culminating in an F1-score of 0.99.

  

ISSN: 2455-8834

Volume:09, Issue:09 "September 2024"

These results suggest that the model performs exceptionally well across most categories, particularly in identifying ulcers, caries, and tooth discoloration with high accuracy and precision.

Classification Report:

| | precision | recall | f1-score | support |
|---|---|---|---|---|
| Calculus | 0.73 | 0.72 | 0.72 | 130 |
| DataCaries | 0.98 | 1.00 | 0.99 | 239 |
| Gingivitis | 0.82 | 0.85 | 0.84 | 234 |
| Healthytooth | 0.97 | 0.85 | 0.91 | 41 |
| Hypodontia | 0.89 | 0.89 | 0.89 | 35 |
| ToothDiscoloration | 0.98 | 0.98 | 0.98 | 184 |
| Ulcers | 1.00 | 0.98 | 0.99 | 255 |
| accuracy | | | 0.92 | 1118 |
| macro avg | 0.91 | 0.89 | 0.90 | 1118 |
| weighted avg | 0.92 | 0.92 | 0.92 | 1118 |

**Confusion Matrix Analysis**

The confusion matrix provides deeper insight into the model's performance across the different classes. Key observations include:

- **Calculus:** Some confusion with caries and gingivitis, indicating similar imaging characteristics that may overlap.
- **Gingivitis:** A few instances were misclassified as calculus or healthy teeth, which might be due to the subtle signs of inflammation captured in images.
- **Healthy Teeth:** Mostly accurate but some confusion with hypodontia and tooth discoloration, highlighting challenges in distinguishing very subtle anomalies.
- **Ulcers:** Very high accuracy with minimal misclassification, demonstrating the model's effectiveness in identifying more distinct pathological features.

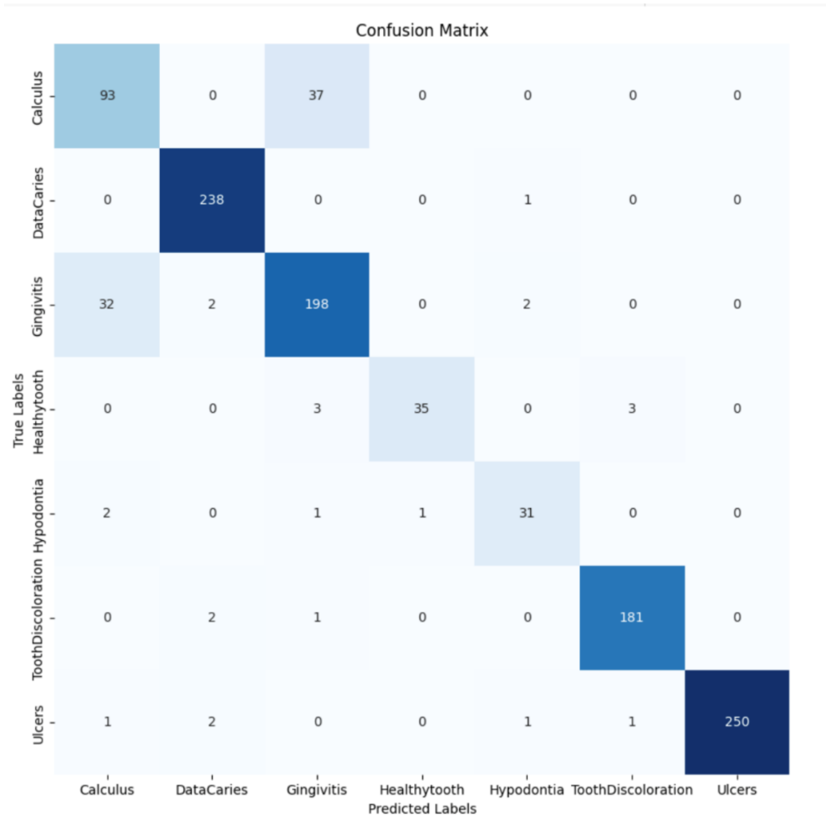


**Visual Grad-CAM Analysis**

The Gradient-weighted Class Activation Mapping (Grad-CAM) provides a visual explanation for decisions from the convolutional part of the model. This analysis helps understand which regions of the image influenced the model's predictions. Below is a summary of the visual results for each class:

- **Calculus, Caries, Gingivitis, Hypodontia, Tooth Discoloration, Ulcers:** The model generally highlights areas relevant to each condition accurately, demonstrating the model's ability to focus on significant regions for diagnosis.

- **Healthy Teeth:** Some images showed misclassifications where the model identified areas as problematic when they were not, indicating potential over-sensitivity or confusion with similar textures or colors in other classes.

The corresponding images and CAM visualizations reveal the model's performance in real scenarios:

  

ISSN: 2455-8834

Volume:09, Issue:09 "September 2024"

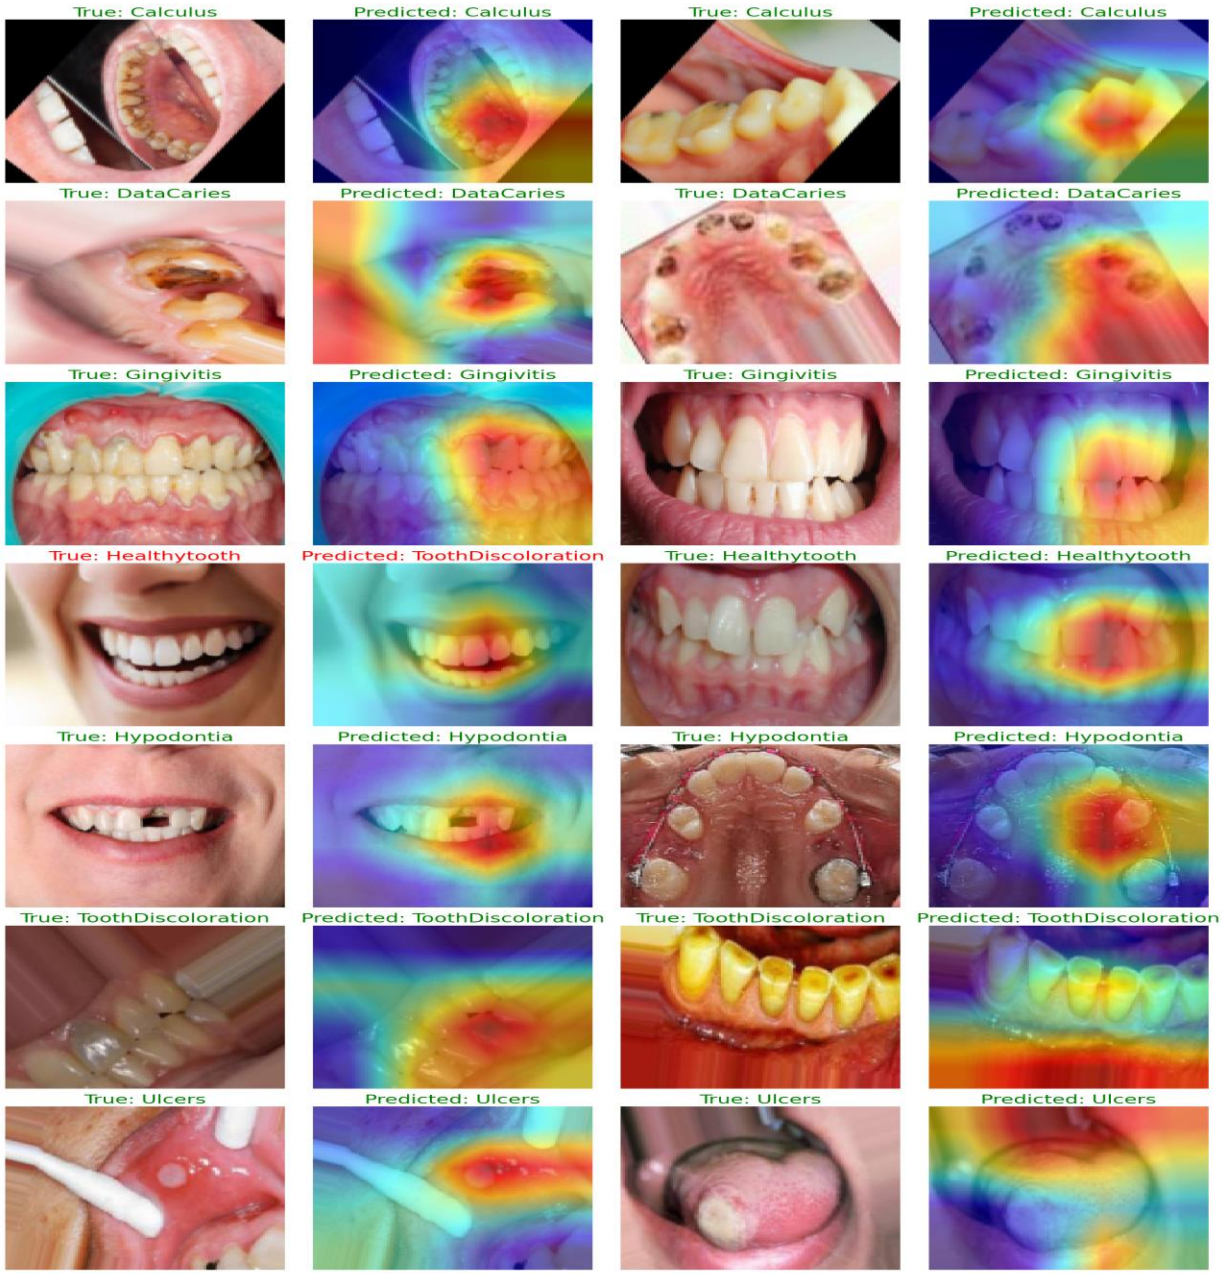


## Discussion

- **Accuracy in Prediction:** The Grad-CAM images clearly illustrate that the model accurately identifies and focuses on the relevant areas in most cases. For diseases like ulcers and caries, the model's focus is highly precise, aligning with the areas of actual disease manifestations.

  

- **Misclassifications:** For healthy teeth, some misclassifications were observed. These could be due to the underrepresentation of this class in the training set, as previously noted in the dataset distribution. Enhancements in the model's training could involve more balanced datasets or improved regularization techniques to mitigate this.

- **Predictive Performance:** The Grad-CAM visualizations also support the quantitative findings from the confusion matrix and classification report, providing a comprehensive understanding of the model's strengths and areas for improvement. The high precision and recall in most categories are visually confirmed through the focus areas in the CAM outputs.

This visual analysis not only reaffirms the quantitative results but also offers insightful details into the model's decision-making process, which is crucial for refining the model and improving its diagnostic capabilities. Further iterations could focus on enhancing the model's sensitivity to less represented conditions and refining its ability to generalize across more diverse oral health scenarios.

## 5. Conclusion

The Virtual Doctor project demonstrates significant advancements in the field of oral health diagnostics through the use of deep learning technologies. The deployment of the ResNet18 architecture within our model has successfully enabled high-precision identification of various dental conditions, as evidenced by the model's overall accuracy rate of 92%. The precision, recall, and F1-scores across different categories further validate the efficacy of the model in clinical-like settings, despite the challenges presented by the skewed distribution of the dataset.

## 6. Future Scope

The Virtual Doctor platform has been initially focused on revolutionizing dental health assessments, but the scope of this project extends far beyond just oral health. As we continue to develop and refine our technology, the platform is poised to expand into several other critical areas of healthcare diagnostics and education.

### Expansion into Facial Disease Detection

A natural extension of our current capabilities would be into the detection and analysis of facial skin conditions. Leveraging similar machine learning technologies, we plan to develop models that can diagnose common facial skin diseases such as:

- **Acne:** Automated recognition and severity classification to guide treatment options.

  

ISSN: 2455-8834

Volume:09, Issue:09 "September 2024"

- **Eczema:** Detection and monitoring over time to evaluate treatment effectiveness.
- **Psoriasis:** Identification and differentiation from other skin conditions.
- **Rosacea:** Early detection and lifestyle management to prevent flare-ups.

These models will use facial imaging to provide quick assessments and help manage chronic conditions through regular monitoring, significantly aiding in patient care.

**Additional Health Modules**

Beyond dental and facial health, the Virtual Doctor platform is envisaged to include several other modules to provide comprehensive health assessments. Some of these future additions could include:

- **Eye Health:** Using image recognition to detect signs of conditions like conjunctivitis, dry eyes, or more serious issues such as glaucoma.
- **Nutritional Deficiencies:** Analysis of physical symptoms visible in the eyes, nails, and skin that may indicate nutritional problems.
- **Mental Health Monitoring:** Using behavioral cues picked up through video interactions to assess signs of mental health issues such as depression or anxiety.

**Educational Components**

To further enhance the platform's utility, we intend to incorporate educational modules that provide users with information about diseases, treatment options, and preventive healthcare measures. These modules will be tailored to the specific conditions identified during the assessments, offering users a personalized health education experience.

**Integration with Health Systems**

Looking forward, we also plan to integrate Virtual Doctor with existing health systems to facilitate seamless information flow between patients and healthcare providers. This integration will support telemedicine initiatives and enable remote monitoring, particularly valuable for patients in remote or underserved regions.

In summary, while the current focus on dental health marks a significant first step, the future of the Virtual Doctor platform promises a holistic approach to health diagnostics and education. By continuously expanding its capabilities and integrating with broader health systems, Virtual Doctor aims to play a pivotal role in improving global health outcomes.

  

**References**

https://link.springer.com/article/10.1007/s11042-023-16776-x

https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9572157/

https://arxiv.org/abs/1512.03385

https://paperswithcode.com/paper/resnet18-model-with-sequential-layer-for

https://ieeexplore.ieee.org/abstract/document/8053547

https://arc.aiaa.org/doi/abs/10.2514/6.2024-2618

  