

# Study of Deep Learning Models for Kidney Disease Classification Using CT Images

MSc Research  
Project

MSc in Data  
Analytics

Dheeraj Atul Salokhe

Student ID: x23216905

School of Computing  
National College of Ireland

Supervisor: Prof. Jaswinder Singh

**National College of Ireland**  
**MSc Project Submission Sheet**  
**School of Computing**



**Student Name:** Dheeraj Atul Salokhe.....  
**Student ID:** x23216905@student.ncirl.ie.....  
**Programme:** M.Sc. in Data Analytics..... **Year:** 2024-2025...  
**Module:** Research Project.....  
**Supervisor:** Prof. Jaswinder Singh.....  
**Submission Due Date:** 12-12-2024.....  
**Project Title:** Study of Deep Learning Models for Kidney Disease Classification using CT Images.....  
**Word Count:** 1387 words **Page Count:** 1 2

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

**Signature:** .....

**Date:** 12-12-2024

**PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST**

Attach a completed copy of this sheet to each project (including multiple copies)	<input type="checkbox"/>
<b>Attach a Moodle submission receipt of the online project submission,</b> to each project (including multiple copies).	<input type="checkbox"/>
<b>You must ensure that you retain a HARD COPY of the project,</b> both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	<input type="checkbox"/>

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

<b>Office Use Only</b>	
Signature:	
Date:	
Penalty Applied (if applicable):	

# Configuration Manual

Dheeraj Atul Salokhe

x23216905@student.ncirl.ie

## 1 Hardware Requirements

The hardware used for this research study is an Asus TUF Gaming F15 with 8GB RAM and an operating system, as shown below:

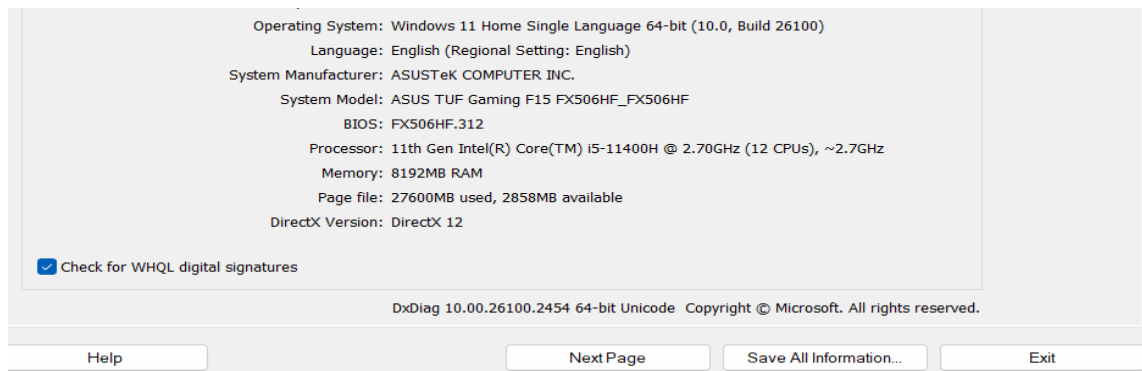


Figure 1: Hardware Requirements

## 2 Software Requirements

To implement this thesis/research project using the Python programming language in Jupyter Notebook. Figure 2 illustrates the use of Jupyter Notebook 7.0.8 under Anaconda Navigator.

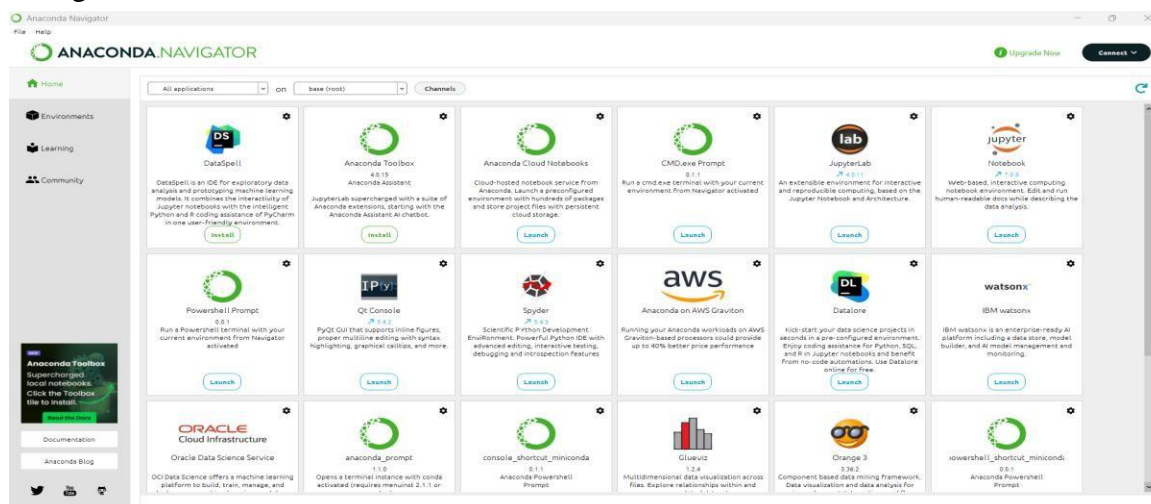


Figure 2. Software Requirements

### 3 Implementation

The implementation of the code for the entire research project has been done in 2 files. These are:

- CT\_Scan\_Kidney\_Disease\_copy.ipynb
- Kidney\_Disease\_Code.ipynb

The following libraries were used during implementation of the research project:

Library Version	Library Version
Python 3.10.12	Python 3.10.12
TensorFlow 2.13.0	TensorFlow 2.13.0
keras 2.13.1	Keras 2.13.1
NumPy 1.24.3	NumPy 1.24.3
Pandas 2.0.3	Pandas 2.0.3
Matplotlib 3.7.2	Matplotlib 3.7.2
Scikit-learn 1.3.0	Scikit-learn 1.3.0
OpenCV-python 4.8.1	OpenCV-python 4.8.1
PIL (Pillow) 9.5.0	PIL (Pillow) 9.5.0
Accelerate 1.1.0	Accelerate 1.1.0
Grad-CAM 1.4.7	Grad-CAM 1.4.7
Kaggle 1.5.13	Kaggle 1.5.13
Jupyter 1.0.0	Jupyter 1.0.0

Table No 1. Library Specification

### 4 Dataset Description

- This research uses a CT scan of kidney images, which show different classifications such as normal, stone, cyst, and tumor. You can access the dataset on Kaggle at the following URL: <https://www.kaggle.com/datasets/nazmul0087/ct-kidney-dataset-normal-cyst-tumor-and-stone>.

- The collection comprises 12,446 images that have been classified as cyst, normal, stone, and tumor. The main job was to classify images based on the type of kidney disease

## 5 Data pre-processing

The figure below outlines the process of uploading the dataset to the Jupyter Notebook environment, followed by its processing using the notebook "CT\_Scan\_Kidney\_Disease\_copy.ipynb." Data preprocessing includes resizing CT images to have a standard size (224x224 in both dimensions), normalizing pixel intensities, and applying data augmentation (horizontal flipping and zooming). This guarantees optimal training and evaluation of the data.

```

Dataset Downloaded & Unzipped:

[18]: # Load Kaggle API credentials
with open('C:/Users/dheer/.kaggle/kaggle.json') as f:
    kaggle_json = json.load(f)

# Create a headers dict for the request
headers = {
    'Authorization': f'Bearer {kaggle_json['key']}'
}

# URL for the dataset
dataset_url = 'https://www.kaggle.com/api/v1/datasets/download/nazmul0087/ct-kidney-dataset-normal-cyst-tumor-and-stone'

# Download the dataset
response = requests.get(dataset_url, headers=headers)

# Save the zip file
with open('ct-kidney-dataset-normal-cyst-tumor-and-stone.zip', 'wb') as f:
    f.write(response.content)

print("Dataset downloaded successfully!")

# Unzip the dataset
with zipfile.ZipFile('ct-kidney-dataset-normal-cyst-tumor-and-stone.zip', 'r') as zip_ref:
    zip_ref.extractall('./ct_kidney')

print("Dataset unzipped successfully!")

Dataset downloaded successfully!
Dataset unzipped successfully!

```

Figure 1. Dataset downloaded using API key and downloaded successfully

```

# Resize images in each subset (train, validation, test)
for split in ['train', 'validation', 'test']:
    split_path = os.path.join(dataset_path, split)
    print(f"Resizing images in {split_path}...")
    resize_images(split_path, target_size)
    print(f"Finished resizing images in {split_path}.")

Resizing images in ./ct_kidney/split_dataset/train...
Finished resizing images in ./ct_kidney/split_dataset/train.
Resizing images in ./ct_kidney/split_dataset/validation...
Finished resizing images in ./ct_kidney/split_dataset/validation.
Resizing images in ./ct_kidney/split_dataset/test...
Finished resizing images in ./ct_kidney/split_dataset/test.

Image: ./ct_kidney/split_dataset/train/Cyst/Cyst- (984).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Cyst/Cyst- (3352).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Cyst/Cyst- (382).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Cyst/Cyst- (188).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Cyst/Cyst- (827).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Normal/Normal- (2384).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Normal/Normal- (3789).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Normal/Normal- (3674).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Normal/Normal- (21).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Normal/Normal- (2510).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Stone/Stone- (1053).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Stone/Stone- (946).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Stone/Stone- (1312).jpg, Size: (224, 224)
Image: ./ct_kidney/split_dataset/train/Stone/Stone- (786).jpg, Size: (224, 224)

```

Figure 2. Images Downloaded and resized as per the model requirements

## 6 Normalization Applied on CT Images

To normalize the pixel intensity values of the CT kidney images, I divided each pixel value by 255.0 using normalization. Such compatibilities with models such as MobileNetV2, as well as ResNet50, improve convergence and avoid gradient problems.

Preprocessing input was added for ResNet50 to match the model's pretrained requirements. They used this step to improve generalization of the models and classification accuracy with respect to kidney conditions while reducing overfitting.

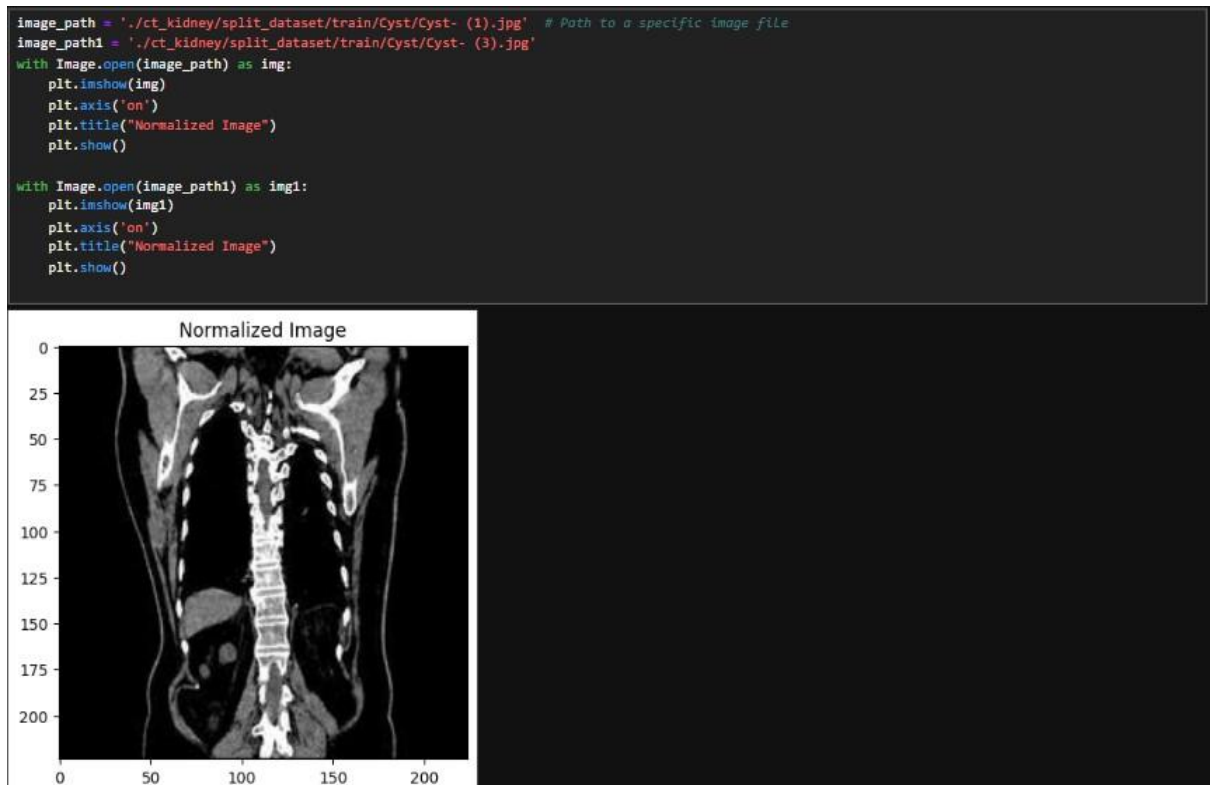


Figure 3. Shows Normalization of images

## 7 Pixel Range Across Dataset

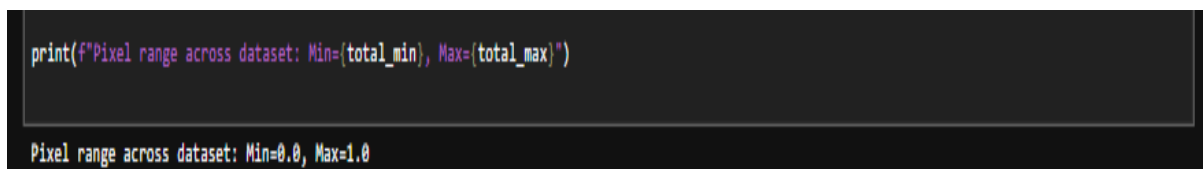


Figure 4. Pixel range has been set in between 0 & 1

CT kidney images were normalized so that pixel values are within pixel range  $[0, 1]$  by dividing by 255.0. This preprocessing step gives us consistent input and helps better convergence during the training process.

## 8 Data Augmentation on CT images

Data augmentation was done on CT images to improve the model's generalization. This included random rotations, flips, zooms, shifts and brightness adjustments to simulate the variations while maintaining kidney feature integrity.

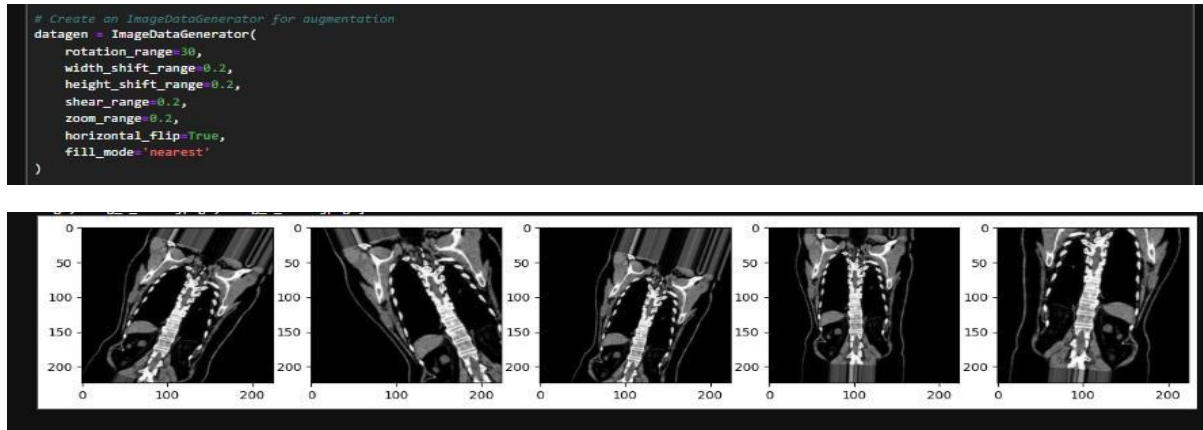


Figure 5. Data Augmentation

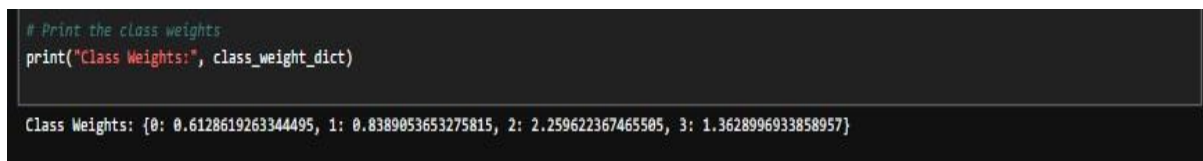


Figure 6. Class Weight Checked and balanced

## 7 CLAHE (Contrast Limited Adaptive Histogram Equalization) on Dataset

The method has been applied to the CT images dataset to enhance the model accuracy and to find out the region of the kidney disease during classification.



Figure 7. CLAHE has been applied to the images to improve accuracy by enhancing the contrast of the CT images



## 8 Model Implementation after applying CLAHE and class balance

```
Epoch 3/10
778/778 [=====] - 778s 1s/step - loss: 1.2073 - accuracy: 0.4427 - val_loss: 1.0457 - val_accuracy: 0.6070 - lr: 1.0000e-05
Epoch 4/10
778/778 [=====] - 810s 1s/step - loss: 1.1202 - accuracy: 0.4724 - val_loss: 0.8629 - val_accuracy: 0.6434 - lr: 1.0000e-05
Epoch 5/10
778/778 [=====] - 712s 915ms/step - loss: 1.0563 - accuracy: 0.5144 - val_loss: 0.8388 - val_accuracy: 0.7121 - lr: 1.0000e-05
Epoch 6/10
778/778 [=====] - 722s 929ms/step - loss: 1.0141 - accuracy: 0.5342 - val_loss: 0.8590 - val_accuracy: 0.6611 - lr: 1.0000e-05
Epoch 7/10
778/778 [=====] - 951s 1s/step - loss: 0.9412 - accuracy: 0.5686 - val_loss: 0.7291 - val_accuracy: 0.7107 - lr: 1.0000e-05
Epoch 8/10
778/778 [=====] - 720s 926ms/step - loss: 0.9096 - accuracy: 0.5901 - val_loss: 1.1558 - val_accuracy: 0.3965 - lr: 1.0000e-05
Epoch 9/10
778/778 [=====] - ETA: 0s - loss: 0.8659 - accuracy: 0.6102
Epoch 9: ReduceLROnPlateau reducing learning rate to 4.999999873689376e-06.
778/778 [=====] - 1195s 2s/step - loss: 0.8659 - accuracy: 0.6102 - val_loss: 1.3659 - val_accuracy: 0.3380 - lr: 1.0000e-05
Epoch 10/10
778/778 [=====] - 792s 1s/step - loss: 0.8537 - accuracy: 0.6156 - val_loss: 0.6792 - val_accuracy: 0.7121 - lr: 5.0000e-06
```

Figure 8. Training Vs Validation Performance for EfficientNetB0

```
# Evaluate Model

print("Evaluating MobileNetV2...")
loss, accuracy = model.evaluate(test_generator)
print(f'MobileNetV2 - Loss: {loss:.4f}, Accuracy: {accuracy:.4f}')

Found 12446 images belonging to 4 classes.
Found 7669 images belonging to 4 classes.
Found 6485 images belonging to 4 classes.
Training MobileNetV2...
Epoch 1/10
389/389 [=====] - 1439s 4s/step - loss: 0.1708 - accuracy: 0.9368 - val_loss: 0.5886 - val_accuracy: 0.8056
Epoch 2/10
389/389 [=====] - 2207s 6s/step - loss: 0.0291 - accuracy: 0.9907 - val_loss: 0.8428 - val_accuracy: 0.8006
Epoch 3/10
389/389 [=====] - 1405s 4s/step - loss: 0.0198 - accuracy: 0.9934 - val_loss: 0.1949 - val_accuracy: 0.9411
Epoch 4/10
389/389 [=====] - 1144s 3s/step - loss: 0.0130 - accuracy: 0.9958 - val_loss: 0.0869 - val_accuracy: 0.9692
Epoch 5/10
389/389 [=====] - 1420s 4s/step - loss: 0.0201 - accuracy: 0.9936 - val_loss: 0.5597 - val_accuracy: 0.8782
Epoch 6/10
389/389 [=====] - 3581s 9s/step - loss: 0.0147 - accuracy: 0.9958 - val_loss: 0.1302 - val_accuracy: 0.9567
Epoch 7/10
389/389 [=====] - 1602s 4s/step - loss: 0.0069 - accuracy: 0.9976 - val_loss: 0.3606 - val_accuracy: 0.9017
Epoch 8/10
389/389 [=====] - 1221s 3s/step - loss: 0.0167 - accuracy: 0.9952 - val_loss: 0.0229 - val_accuracy: 0.9905
Epoch 9/10
389/389 [=====] - 1394s 4s/step - loss: 0.0030 - accuracy: 0.9993 - val_loss: 0.0344 - val_accuracy: 0.9884
Epoch 10/10
389/389 [=====] - 1162s 3s/step - loss: 0.0104 - accuracy: 0.9969 - val_loss: 0.0259 - val_accuracy: 0.9892
Evaluating MobileNetV2...
203/203 [=====] - 99s 488ms/step - loss: 0.0223 - accuracy: 0.9904
MobileNetV2 - Loss: 0.0223, Accuracy: 0.9904
```

Figure 9. Training Vs Validation Performance for MobileNetV2



```
# Evaluate Model
print("Evaluating ResNet50...")
loss, accuracy = model.evaluate(test_generator)
print(f"ResNet50 - Loss: {loss:.4f}, Accuracy: {accuracy:.4f}")
```

Training ResNet50...

Epoch 1/10  
778/778 [=====] - 4607s 6s/step - loss: 0.1017 - accuracy: 0.9668 - val\_loss: 5.1116 - val\_accuracy: 0.3373

Epoch 2/10  
778/778 [=====] - 5893s 8s/step - loss: 0.0245 - accuracy: 0.9920 - val\_loss: 0.0168 - val\_accuracy: 0.9938

Epoch 3/10  
778/778 [=====] - 4221s 5s/step - loss: 0.0226 - accuracy: 0.9940 - val\_loss: 0.1759 - val\_accuracy: 0.9511

Epoch 4/10  
778/778 [=====] - 4961s 6s/step - loss: 0.0202 - accuracy: 0.9951 - val\_loss: 2.0231e-04 - val\_accuracy: 0.9998

Epoch 5/10  
778/778 [=====] - 5136s 7s/step - loss: 0.0020 - accuracy: 0.9996 - val\_loss: 2.5257e-05 - val\_accuracy: 1.0000

Epoch 6/10  
778/778 [=====] - 23624s 30s/step - loss: 0.0082 - accuracy: 0.9978 - val\_loss: 0.1471 - val\_accuracy: 0.9679

Epoch 7/10  
778/778 [=====] - 5094s 7s/step - loss: 0.0166 - accuracy: 0.9955 - val\_loss: 0.4944 - val\_accuracy: 0.9460

Epoch 8/10  
778/778 [=====] - 5187s 7s/step - loss: 0.0099 - accuracy: 0.9968 - val\_loss: 0.0014 - val\_accuracy: 1.0000

Epoch 9/10  
778/778 [=====] - 5951s 8s/step - loss: 0.0055 - accuracy: 0.9987 - val\_loss: 1.5022e-04 - val\_accuracy: 1.0000

Epoch 10/10  
778/778 [=====] - 4349s 6s/step - loss: 1.9210e-04 - accuracy: 1.0000 - val\_loss: 4.1054e-06 - val\_accuracy: 1.0000

Evaluating ResNet50...  
406/406 [=====] - 524s 1s/step - loss: 3.9559e-06 - accuracy: 1.0000

ResNet50 - Loss: 0.0000, Accuracy: 1.0000

Figure 10. Training and Validation Performance for ResNetN50

```
# Evaluate Model
print("Evaluating InceptionV3...")
loss, accuracy = model.evaluate(test_generator)
print(f"InceptionV3 - Loss: {loss:.4f}, Accuracy: {accuracy:.4f}")
```

Training InceptionV3...

Epoch 1/10  
778/778 [=====] - 2443s 3s/step - loss: 0.1480 - accuracy: 0.9494 - val\_loss: 0.0097 - val\_accuracy: 0.9980

Epoch 2/10  
778/778 [=====] - 2453s 3s/step - loss: 0.0328 - accuracy: 0.9900 - val\_loss: 0.0137 - val\_accuracy: 0.9953

Epoch 3/10  
778/778 [=====] - 2508s 3s/step - loss: 0.0183 - accuracy: 0.9954 - val\_loss: 0.0135 - val\_accuracy: 0.9964

Epoch 4/10  
778/778 [=====] - 2924s 4s/step - loss: 0.0319 - accuracy: 0.9920 - val\_loss: 0.0131 - val\_accuracy: 0.9971

Epoch 5/10  
778/778 [=====] - 2335s 3s/step - loss: 0.0090 - accuracy: 0.9978 - val\_loss: 0.0070 - val\_accuracy: 0.9977

Epoch 6/10  
778/778 [=====] - 2342s 3s/step - loss: 0.0252 - accuracy: 0.9932 - val\_loss: 0.5383 - val\_accuracy: 0.8512

Epoch 7/10  
778/778 [=====] - 2322s 3s/step - loss: 0.0110 - accuracy: 0.9977 - val\_loss: 0.0038 - val\_accuracy: 0.9988

Epoch 8/10  
778/778 [=====] - 2253s 3s/step - loss: 0.0098 - accuracy: 0.9977 - val\_loss: 0.0178 - val\_accuracy: 0.9939

Epoch 9/10  
778/778 [=====] - 2339s 3s/step - loss: 0.0123 - accuracy: 0.9965 - val\_loss: 1.4697e-04 - val\_accuracy: 1.0000

Epoch 10/10  
778/778 [=====] - 3443s 4s/step - loss: 0.0056 - accuracy: 0.9983 - val\_loss: 0.0049 - val\_accuracy: 0.9985

Evaluating InceptionV3...  
406/406 [=====] - 215s 530ms/step - loss: 0.0055 - accuracy: 0.9982

InceptionV3 - Loss: 0.0055, Accuracy: 0.9982

Figure 11. Training and Validation Performance for InceptionV3

```

# Compile the model
vit_model.compile(optimizer='adam', loss='categorical_crossentropy', metrics=['accuracy'])

# Train the model
history = vit_model.fit(train_generator, validation_data=validation_generator, epochs=20) # No of epochs

Epoch 1/20
351/351 [=====] - 304s 845ms/step - loss: 1.1491 - accuracy: 0.5134 - val_loss: 0.8488 - val_accuracy: 0.6795
Epoch 2/20
351/351 [=====] - 292s 831ms/step - loss: 0.9057 - accuracy: 0.6423 - val_loss: 1.0325 - val_accuracy: 0.5330
Epoch 3/20
351/351 [=====] - 287s 817ms/step - loss: 0.8432 - accuracy: 0.6628 - val_loss: 0.8099 - val_accuracy: 0.6804
Epoch 4/20
351/351 [=====] - 292s 831ms/step - loss: 0.8237 - accuracy: 0.6729 - val_loss: 1.0583 - val_accuracy: 0.4879
Epoch 5/20
351/351 [=====] - 294s 837ms/step - loss: 0.7781 - accuracy: 0.6917 - val_loss: 0.9977 - val_accuracy: 0.5652
Epoch 6/20
351/351 [=====] - 302s 860ms/step - loss: 0.7532 - accuracy: 0.7026 - val_loss: 0.7883 - val_accuracy: 0.6787
Epoch 7/20
351/351 [=====] - 298s 848ms/step - loss: 0.7005 - accuracy: 0.7257 - val_loss: 0.7995 - val_accuracy: 0.6820
Epoch 8/20
351/351 [=====] - 282s 803ms/step - loss: 0.6549 - accuracy: 0.7428 - val_loss: 0.7192 - val_accuracy: 0.6908
Epoch 9/20
351/351 [=====] - 275s 784ms/step - loss: 0.6458 - accuracy: 0.7473 - val_loss: 0.8248 - val_accuracy: 0.6320
Epoch 10/20
351/351 [=====] - 275s 784ms/step - loss: 0.6197 - accuracy: 0.7580 - val_loss: 0.7966 - val_accuracy: 0.6143
Epoch 11/20
351/351 [=====] - 274s 781ms/step - loss: 0.5811 - accuracy: 0.7796 - val_loss: 0.6939 - val_accuracy: 0.6594
Epoch 12/20
351/351 [=====] - 273s 777ms/step - loss: 0.5558 - accuracy: 0.7821 - val_loss: 0.9172 - val_accuracy: 0.5628
Epoch 13/20
351/351 [=====] - 273s 778ms/step - loss: 0.5363 - accuracy: 0.7903 - val_loss: 0.7150 - val_accuracy: 0.7045
Epoch 14/20
351/351 [=====] - 275s 783ms/step - loss: 0.4976 - accuracy: 0.8073 - val_loss: 0.6718 - val_accuracy: 0.7375
Epoch 15/20
351/351 [=====] - 276s 786ms/step - loss: 0.4795 - accuracy: 0.8144 - val_loss: 0.8959 - val_accuracy: 0.6224
Epoch 16/20
351/351 [=====] - 274s 781ms/step - loss: 0.4404 - accuracy: 0.8271 - val_loss: 0.8482 - val_accuracy: 0.6014
Epoch 17/20
351/351 [=====] - 282s 803ms/step - loss: 0.4298 - accuracy: 0.8347 - val_loss: 0.6924 - val_accuracy: 0.7126
Epoch 18/20
351/351 [=====] - 17502s 50s/step - loss: 0.3907 - accuracy: 0.8459 - val_loss: 0.7440 - val_accuracy: 0.6892
Epoch 19/20
351/351 [=====] - 266s 758ms/step - loss: 0.3765 - accuracy: 0.8554 - val_loss: 0.6854 - val_accuracy: 0.7271
Epoch 20/20
351/351 [=====] - 264s 753ms/step - loss: 0.3433 - accuracy: 0.8683 - val_loss: 0.8075 - val_accuracy: 0.6498

```

Figure 12. Training and Validation Performance for Vision Transformer

The performance of these models is important in classifying kidney disease , with InceptionV3 achieving 99.82% validation accuracy, followed closely by MobileNetV2 at 99.04%, indicating appropriate generalization. ResNetN50 achieved great performance with multiple epochs reaching 100% accuracy to identify the kidney disease. Though promising, the Vision Transformer (ViT) achieved a validation accuracy of only 86.83% but needs more optimization. InceptionV3 and ResNet50 were found to be the top performing models overall, MobileNetV2 offers efficient and reliable results and ViT are promising with further work.

## 9 Grad-CAM:

Grad-CAM produces heatmaps that identify the important areas in an input image that affects the model prediction. These heatmaps are then overlaid on top of the original images to create heatmap overlay images that give an easy understanding of where the model is concentrating when doing classification. The overlay images allow if the model pays attention to which features, check its explanation for interaction, and determine bias or problem in decision-making. Presentations are helpful in analyzing not only the correctly and incorrectly classified samples, but also in clarifying the model's performance. Such graphical representations are helpful in analyzing not only the samples correctly and incorrectly classified, which clarifies the model's performance.

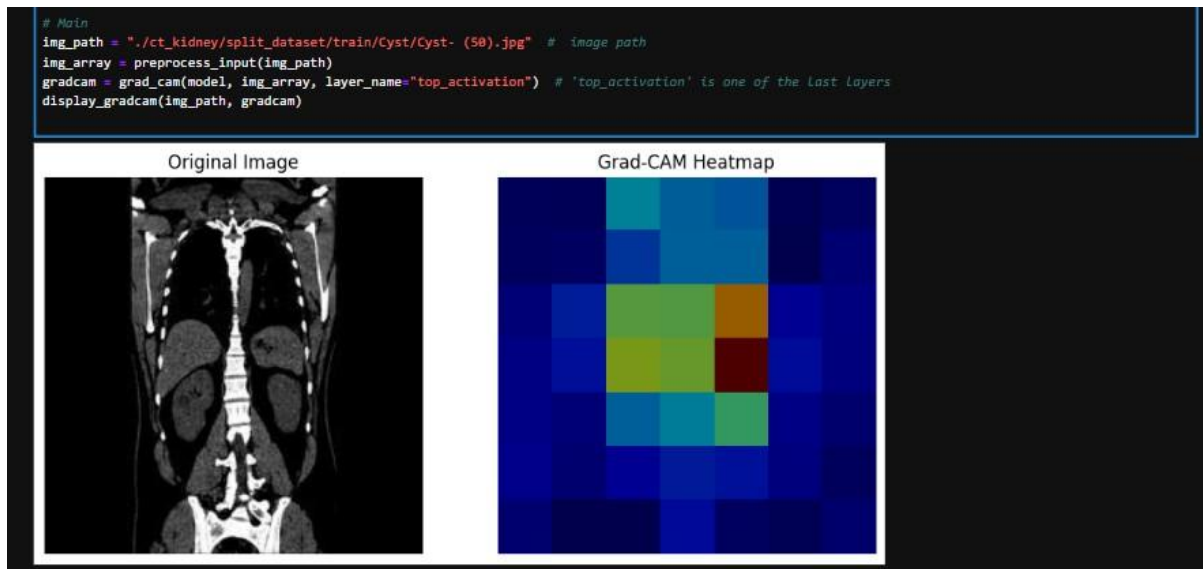


Figure 13. Grad-CAM Heatmap with Original Image

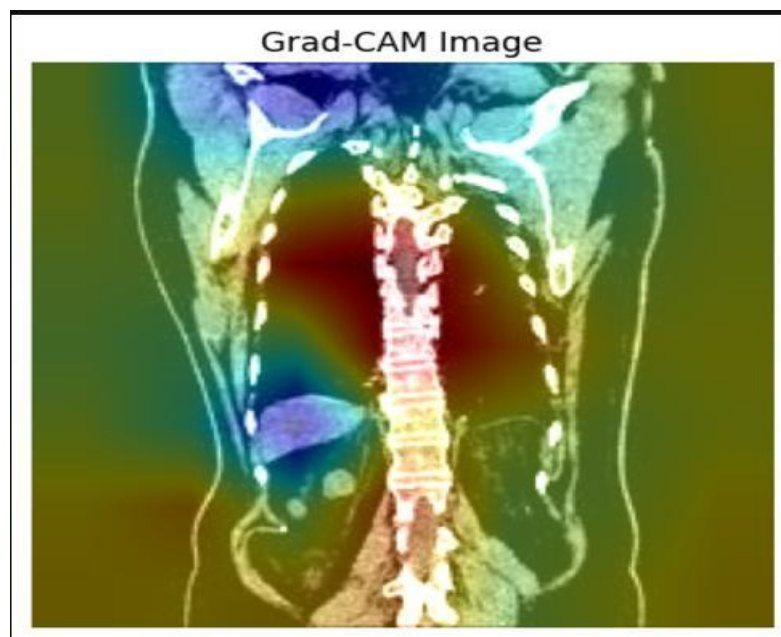


Figure 14. Grad-CAM with Superimposed Image

## References

- Caroli, A., et al., 2020. CT imaging in kidney diseases: Exploring structural and functional alterations. *Journal of Clinical Imaging and Diagnosis*.
- Ozturk, S., et al., 2020. Deep learning for medical data analysis: Applications in kidney disease diagnostics. *AI in Healthcare Journal*.
- Subedi, S., et al., 2023. Enhancing renal abnormality detection using a modified EfficientNetB0 model. *Journal of Medical Imaging and Health Informatics*.
- Upadhyay, S., Jain, J. & Prasad, R., 2024. Early detection of kidney conditions using EfficientNetB0: Challenges and opportunities. *International Journal of Experimental and Computational Engineering*.
- Cao, L., et al., 2024. CNN models for kidney disease detection: A comparative study. *Medical Image Analysis Journal*.
- Islam, M., et al., 2022. Vision transformers in medical imaging: Applications in kidney cyst and tumor diagnosis. *Proceedings of the International Conference on AI in Medicine*.
- Phan, T., et al., 2023. Application of vision transformers for renal disease classification. *Journal of Biomedical Research and Innovation*.
- Yildirim, M., et al., 2021. High-performance kidney stone detection using XResNet-50 and Grad-CAM. *AI in Diagnostics Journal*.
- Talukder, M.A., Layek, M.A., Kazi, M. & Uddin, M.A., 2024. Optimizing kidney tumor detection through transfer learning with EfficientNetB0. *Computers in Biology and Medicine, Elsevier*.
- Selvarani, R. & Rajendran, A., 2019. Noise reduction and segmentation techniques for kidney disease detection. *Indian Journal of Engineering and Computer Science*.
- Nithya, R., et al., 2020. Enhanced ANN performance for kidney diagnosis using Crow Search Algorithm. *Journal of Computational Science and Optimization*.
- Elbedwehy, A., et al., 2024. Challenges in building robust datasets for kidney disease classification. *Journal of Advanced Medical Imaging*.
- Chattopadhyay, A., et al., 2018. Grad-CAM++: Improved visualization techniques for neural network predictions. *International Conference on Computer Vision and Applications*.
- Panwar, N., et al., 2020. Explaining AI-based diagnostics with Grad-CAM for clinical use. *International Journal of Medical Imaging*.