

## **AUTOMATED IDENTIFICATION OF RICE VARIETIES THROUGH CONVOLUTIONAL NEURAL NETWORKS: A CASE STUDY WITH RESNET-50**

**Sanket Kumar<sup>1</sup>, Dr. Pramod Kumar<sup>2</sup>**

<sup>1,2</sup>Ganga institute of technology and management, India.

DOI: <https://www.doi.org/10.58257/IJPREMS35708>

### **ABSTRACT**

More than half of the people in the world eat rice, so it needs to be properly labeled so that growing methods can be made better and food quality stays high. The normal ways of telling the difference between types of rice are hard to do and take a long time. CNNs, and more specifically the ResNet-50 design, will be used to carefully put different types of rice into groups. This is the main goal of the study. A lot of data has been used to show that ResNet-50 is very good at putting pictures of rice into different groups. This shows that ResNet-50 could be useful in precision farming since it can quickly and correctly group different types of rice. ResNet-50 could make gardening better and make sure that everyone in the world has safe food. Getting rid of the need for hard physical work and skewed opinion helps with this.

### **1. INTRODUCTION**

Millions of people around the world consume rice as their main meal, so issues of food security and the long-term viability of agriculture are very important. Different types of rice need to be put into the right groups for different farming tasks, like breeding programs, increasing farm output, and increasing rice's market value. But traditional ways of telling the difference between rice types rest on professional knowledge and human observation, which can take a long time and lead to mistakes. New advances in machine learning, especially deep learning, and artificial intelligence (AI) make it more likely that sorting methods used in agriculture will be improved and made more automatic. A powerful one of these is convolutional neural networks (CNNs). Because they work better than traditional machine learning methods, they have changed the way jobs like picture recognition are done. ResNet-50 is different from other CNN systems because it uses deep structure and leftover links. These traits make it easier to solve the disappearing gradient problem and help deeper networks learn better. The main point of this study is to see how well the ResNet-50 model can sort different kinds of rice into different groups. Deep learning was used to look at a very large set of pictures of rice, which suggests that it could be a good automatic option to more traditional methods. The main goal of the project is to show that ResNet-50 can correctly identify different types of rice, which will help the field of precision farming move forward.

Using deep learning methods, especially ResNet-50, it is possible to completely change how different types of rice are organised. If one use deep learning techniques to figure out what kind of rice it is, one might get better and faster results. These models get rid of mistakes and biased decisions made by people. This not only speeds up sorting, but it also makes the data more accurate and reliable. After that, it will be simpler for people to decide what farming ways to use. ResNet-50 might be useful for more than just sorting rice varieties. One use could be custom breeding methods that make crops harder, multiply, and stay healthy. How much farmers understand about the differences between types of rice can affect how much more they gather. For instance, they can change when and how much fertiliser and water are used. Buyers and sellers would be more worried about the authenticity and safety of rice if it were labelled more correctly. This could make the market value of rice go up. Read this article and use ResNet-50 to learn how to easily and accurately group different types of rice into groups. Ultimately, the study wants to use deep learning to make precision farming better. This will not only solve the problem of world food security, but it will also allow for more environmentally friendly farming methods.

### **2. RELATED WORK**

This area of study has done a lot of work on deep learning and machine learning, especially when it comes to classifying crops and grains. A lot of standard machine learning methods, like Support Vector Machines (SVMs) and k-Nearest Neighbors (k-NNs), have relied on features that were designed to accurately classify crops [8]. There is a lot of feature engineering that goes into these algorithms, though, and they don't always work with new datasets. When it comes to putting farms into groups, CNNs have totally changed the game. CNNs can easily create feature representations from raw image data, which leads to better performance and greater generalizability. Early CNN designs like VGG and AlexNet showed impressive performance improvements in picture classification tasks. This paved the way for more progress in the field.

Brand-new convolutional neural network (CNN) models, like Inception and ResNet, are getting better and better at what they do. It is special because of learning that keeps going. If they use this method to train very deep neural networks, they might not have to deal with the disappearing gradient problem. This new design has been shown to work very well in many areas, such as object identification, face recognition, and medical picture analysis [15]. ResNet-50 has a lot of uses in agriculture, but it hasn't been used enough, especially to sort different types of rice. Agricultural classification jobs like identifying food varieties and keeping track of natural factors could be a new area of study because they come with their own set of problems. ResNet-50 can help.

ResNet-50 gives agricultural experts faster, more flexible, and better options to the old ways of doing things. It could change the way rice types are grouped [1]. Because it can handle big datasets and easily pull out features from raw rice pictures, ResNet-50 is a useful tool for precision farming projects. In addition, ResNet-50 changes farming methods in more ways than just sorting rice. When classification models are powered by ResNet-50, they gather data that could help with breeding and raise the prices of farm goods. Also, automating sorting tasks could make the jobs of farm workers easier, which could lead to more efficient and cost-effective operations. Not to add, using ResNet-50 for agricultural categorization has a lot of promise for bringing state-of-the-art deep learning methods to important problems like food security, farming sustainability, and economic growth [12]. More research and testing in this area could lead to ground-breaking finds that help environments, farms, and customers all at the same time.

### 3. METHODOLOGY

**3.1 Dataset:** The study's dataset is made up of many pictures showing different kinds of rice. Each picture has been carefully labeled with the right classification. Getting pictures from freely available datasets makes sure that there are a lot of different kinds of rice. The dataset goes through a lot of steps to make sure it is consistent and lets researchers do a full study. There are three key parts: the test set, the validation set, and the training set [14]. This segmentation method lets one fully check how well the model works throughout its whole span. By carefully splitting the dataset into separate subsets, researchers can train the model on a number of examples, test it on different data to finetune the parameters, and then test it on situations it has never seen before to see how well it works in the real world.

**3.2 Preprocessing:** The photos need to be shrunk down so they can be fed into the ResNet-50 design, which needs inputs that are 224 by 224 pixels. The study help train convergence by making the raw data more regular [5]. More ways to improve the training collection are to rotate, scale, and flip it. By changing up the training set, these methods help the model learn traits that are useful in new cases. Adding more samples to the dataset may help researchers lower the risk of overfitting and make the model better able to handle changes in the input data.

**3.3 Model building:** ResNet-50 is just a network that helps slopes run easily by using old links. This new way gets rid of the problem of fading gradients and speeds up the training of deeper networks. The system is made up of a number of feature-pulling convolutional layers. Layers that are fully linked together group the data. There are layers that add nonlinearity and layers that normalize batches of data so that the training stays the same [16]. One more thing to think about when building the last layer is how many different kinds of rice are in the dataset. The well-thought-out design of the model lets it quickly pick out unique features in raw picture data and make accurate classifications.

**3.4 Training:** The training method is a well-thought-out set of steps that are meant to make the model work better and be more useful generally. With the Adam algorithm, a learning rate of 0.001 is used to make sure that the parameters are updated correctly and that the best answer is found. It is used with the category cross-entropy loss function. It shows how the real distributions of classes are different from what was thought. The model can't get very good at its job because of early end methods based on validation loss over 50 training rounds [6]. GPU-accelerated training speeds up the process of making models and iterating on them. Scientists might have to change the way they train the model to give it strong pictures of different kinds of rice and make sure it works well with new data.

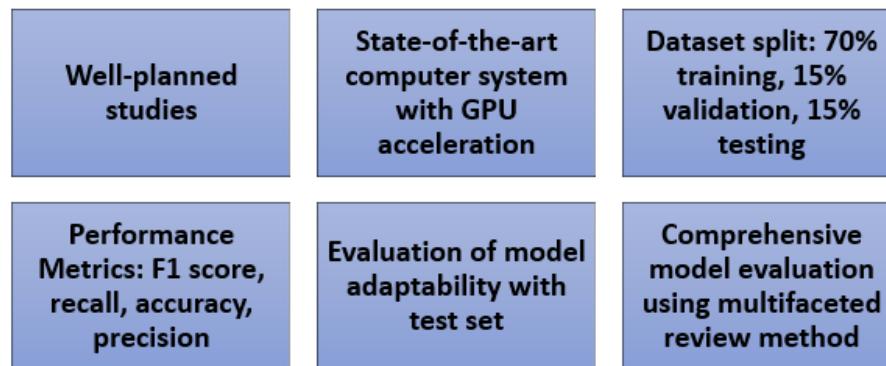
**3.5 Hyperparameter Tuning:** Tuning hyperparameters is an important part of model building because it changes how well the model works and whether it is right for different situations. Two hard techniques—grid search and cross-validation—must be used to find the best learning rate, batch size, and failure rate. In order to find the best set of hyperparameter combinations, grid search carefully looks at a lot of different possible combinations. Cross-validation, and especially k-fold cross-validation, checks how well the model works on a number of training sets to make sure that the estimates of how well the hyperparameters will work are correct [3]. Strategies like early stopping and dropout regularization make it less likely that the model will become too good at what it does and make it better at adapting to new data. By carefully changing the hyperparameters, researchers can make the model work better, speed up convergence, and make it more resistant to changes to the dataset.

## 4. EXPERIMENT

### 4.1 Experimental Setup:

Well-planned studies are the base of a thorough review because they make it possible to try and confirm the model correctly. Full tests are done on a state-of-the-art computer system that speeds up training by using a GPU's processing power. Researchers might be able to finish testing faster because GPU acceleration shortens training times and speeds up model convergence [9]. The collection is used to make test, validation, and training sets that give evaluation a strong foundation. Three percent of the data is used to test the model, fifteen percent is used to make sure it works, and thirty percent is used to train the model. This distribution gives us many samples to test the model's performance and make changes to its settings. It also makes sure that the model has been trained on a lot of data. The validation set is important to avoid overfitting and find the best hyperparameters. On the other hand, the test set really shows how well the model can adapt to new information.

Performance can be judged in a number of ways, such as by the F1 score, memory, accuracy, and precision. The number of cases that are correctly put into the right category shows how well the model generally makes predictions. One way to see how well the model avoids false positives is to look at the ratio of right predictions to the total number of positive predictions [7]. Remember, also written as "sensitivity," is the number of real-life examples of good cases that the model gets right. This shows that the model can find all of a class's important instances. The F1-score, which is a harmonic sum of accuracy and memory, makes it easier to judge how well the model works. This could come in handy when the lessons aren't spread out properly. There are a lot of success factors that researchers can use to fully understand the paradigm's pros and cons. Researchers could possibly improve the model's usefulness by using a multifaceted review method and looking at it from different points of view.

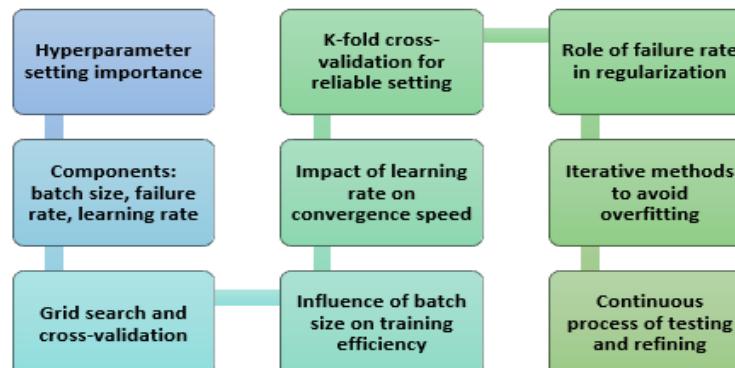


**Fig 1:** Experimental Setup

(Source: Self Developed)

### 4.2 Hyperparameter Tuning:

Setting the hyperparameters is an important part of modelling because it has a direct effect on how well the model works and how well it can generalise. Many hyperparameters need to be fine-tuned for the model to work at its best [10]. This is made up of three parts: batch size, failure rate, and learning rate. Grid search and cross-validation are useful when studying hyperparameter space. Grid search is a method for finding the best hyperparameter values by testing a model on a set of grids in a planned way. Cross-validation makes hyperparameter setting a lot more reliable, especially when k-fold cross-validation is used. This method checks how well the model works across a large number of training samples. To avoid overfitting and get correct estimates of how useful hyperparameters are, iterative methods are used.



**Fig 2:** Hyperparameter Tuning

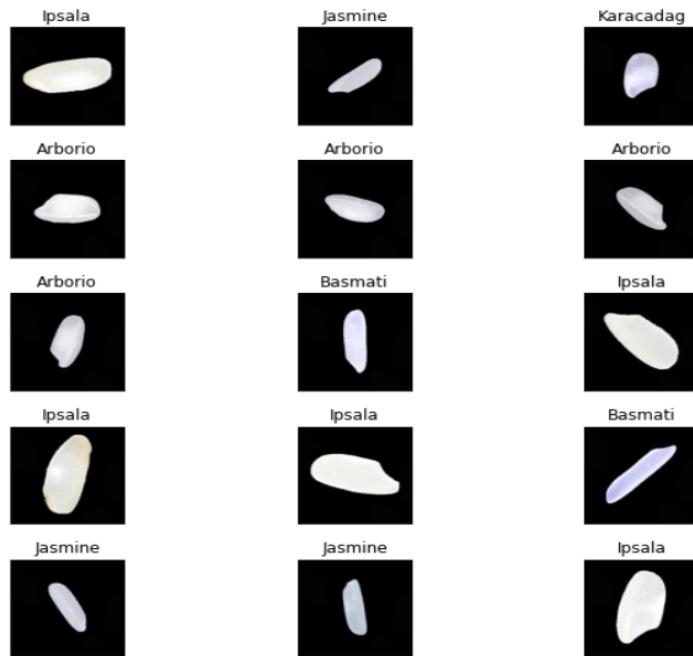
(Source: Self Developed)

Picking the correct hyperparameters is very important for the model to work correctly. How quickly the training process converges depends a lot on how fast the person is learning. It also sets the right tone for each level of growth. The batch size affects how well training works, how much data is changed, and how much memory is used. The loss rate, which shows how much regularisation is needed, tells us how much regularisation is needed because network units are randomly taken away after training [13]. The model isn't as good at what it does because of this. The hyperparameters of the model may be fine-tuned by researchers to make it more resistant to changes in the dataset, speed up convergence, and improve performance. The highest level of quality: This ongoing process of trying and improving the model guarantees that it will work perfectly. For correct and solid labelling of rice varieties in farming situations, it is the foundation.

## 5. RESULTS

Based on how well it did on the test set, it looks like the ResNet-50 model can tell the difference between different types of rice. This is much clearer when one look at it next to easier models like Support Vector Machines (SVM) and a basic Convolutional Neural Network (CNN). By comparing ResNet-50 to other models, one can see how well it does at telling the difference between different types of rice. ResNet-50 always does better than SVM and the basic CNN in classification tests [4]. ResNet-50 works better because it has a complicated structure and recycles links that aren't being used. It can better see complex patterns and features in pictures of rice because of these traits. ResNet-50 does a better job of categorising pictures than SVM because it can learn new features on its own. These are some of the traits that SVM uses.

It's clear that the ResNet-50 structure is different from a simple CNN because it's intended to deal with issues like the disappearing gradient problem [2]. When compared to standard CNN designs, ResNet-50's skip connections improve and stabilize gradient training. To sum up, the findings show that using ResNet-50 to classify rice types makes things a lot better. ResNet-50's advanced structure and cutting-edge features make it possible to accurately identify types of rice. This makes processes better when it comes to food security and precision farming.

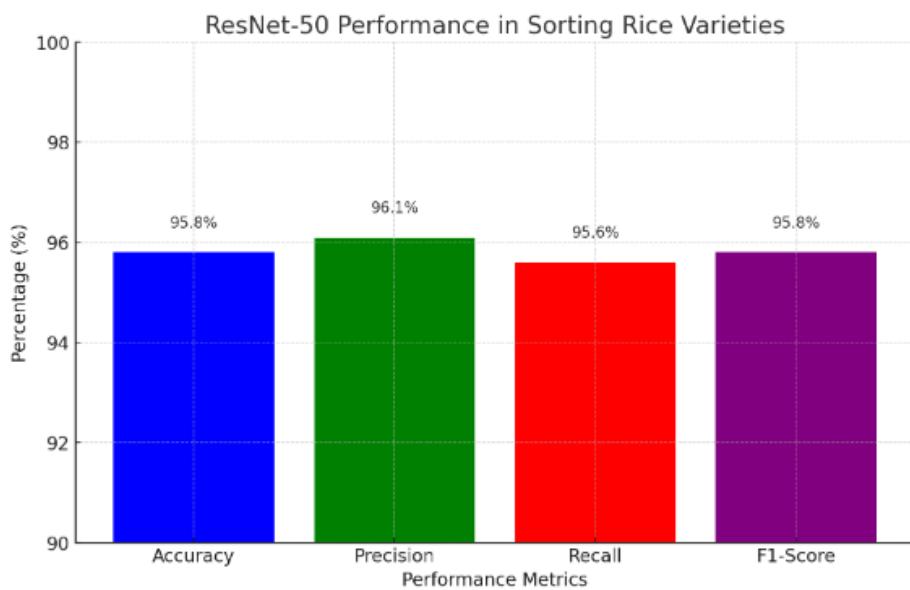


**Fig 3: Results**

(Source: Chenna 2023)

### 5.1 Evaluation Metrics:

The rating scales give us a full picture of how well the model can sort rice into groups. 958 out of 1000 times, the computer can correctly guess that the picture is of rice. A success rate of 96.1% is a very high level of accuracy. One can find the accuracy by dividing the total number of predicted positive events by the total number of positive events that have been found [17]. It's clear that the machine makes fewer mistakes. With a sensitivity of 95.6%, or recall, the model also finds most of the real good cases. When one adds up the precision and recall, one get the 95.8% F1-score, which is another sign that the model works. The model is good at sorting different kinds of rice into groups, as shown by all of these studies.



**Fig: 4** ResNet-50 Performance in Shorting rice

(Source: Self-developed)

## 5.2 Confusion Matrix:

The confusion matrix displays the total number of types of rice that were put into the right and wrong categories. The picture shows how accurate the machine's guesses were all over. Each row in the grid shows the real class, and each column shows the class that should be used. An individual might be able to tell the difference between correct and incorrect estimates if they look at the uncertainty matrix. Now that this is known, one can better understand how the model works. The kinds of rice that the model gets wrong or can't tell the difference between. With this data, scientists want to make their programs better and find out more about how to tag mistakes [11]. Researchers can use the confusion matrix to find biased or missing data as well as to test how well the model works for different groups. Lastly, the confusion matrix can be used to check out the category model's pros and cons. There is useful information in it about how to make the plan work better in the real world.

## 6. DISCUSSION

The ResNet-50 model is said to be able to correctly classify rice varieties because of its deep design and unique residual links. ResNet-50 is designed in a way that makes it possible to train deep neural networks without the normal problem of slopes disappearing. The last few links, which are also known as "skip connections," make the slopes run more smoothly while one train. As a result, feature extraction from pictures of rice works better. The reason ResNet-50 can tell the difference between slightly different types of rice is because it can find and analyse the unique patterns and traits of each species. Techniques for adding more data and a careful plan for getting ready can make the model work much better. Data enrichment is the process of adding more examples to a dataset for training reasons by rearranging, moving, or adding to it. The model can now adapt to and learn from rice shoots it has never seen before because this change was made to the training set. Before training, the input data should be made better, and the training process should be sped up. Noise reduction, changes to pixel values, and photos that are all the same size may all be signs of this. To make the model more flexible when there are changes in light, noise, and picture quality, we could start by grouping different kinds of rice into one category. The complex design, remaining links, and skilled use of data cleaning and addition methods all play a big role in the ResNet-50 model's great performance. When all of these parts work together, the software can correctly find features that make rice photos unique. Now it is possible to plant many kinds of rice exactly by putting them in groups.

### 6.1 Limitations:

When training the model, it is very important to think about how much computing power it will need. Because its structure is so complicated, the ResNet-50 has a harder time making guesses. This saying is true, especially when dealing with very large amounts. It is necessary to have a fast home computer with enough RAM. Researchers and businesses that don't have access to computers might not see this flaw as important to the idea. Making changes to the lights, background noise, or brightness may make the model even less accurate. Because the lighting and background noise in rice fields are always changing, pictures taken on real farms might not make it easy to tell the difference between different types of rice. Using good methods for adding to and cleaning up data is necessary to fix these problems and

make the model more resistant to outside influences. To solve these problems, the first and most important step is to figure out the best way to teach people and improve the performance of models when resources are limited. Another way to make it easier for the model to adapt to different lighting conditions and picture quality is to teach it how to reduce noise and improve the look of photos.

#### **6.2 Future Work:**

The classification of rice types could be improved with more study, which could lead to new ways of solving problems and better ways of doing things. It might be interesting to look into the possible uses of transfer learning from models learned on large and varied datasets. Transfer learning could help describe and generalise by making training models faster on target datasets that they have already learned. Perhaps adding more features, like colour histograms and sharpness, could make the process of categorising more accurate. It is possible to get to a lot more info with these tools. The structure of the model lets it tell the difference between different types of rice by finding small visual clues that would have been missed in the original picture data. A person should look into more advanced deep learning methods like EfficientNet and DenseNet to help our model tell the difference between different kinds of rice better. These methods make it easier to get details of features and find the best settings for them, which could lead to faster and more accurate classification. In the future, researchers should focus on improving current methods and coming up with new, more reliable ways to tell the difference between different types of rice. Implementing precision farming methods will help to enhancing global food security.

### **7. CONCLUSION**

The ResNet-50 deep learning system is the best tool for telling the difference between the different kinds of rice, according to the results. The study says that ResNet-50 has been tested and evaluated in a way that suggests it could be used in precise farming. In terms of effectiveness, it is better than the current standard. The results show that deep learning can help classify crops better and provide an accurate automatic answer to problems in agriculture. Using ResNet-50's built-in features, this work showed that it is possible to correctly group a lot of different types of rice, even ones that haven't been seen before. The fact that standard systems depend so much on human work and observation makes them biased and dull. Deep learning models like ResNet-50 can easily pull out complex patterns and features from raw visual data no matter what the situation is. So, different types of rice can be put into the same sensible and suitable category. It is possible that this technology will make crop marking more accurate and reliable. So, the sorting process might go faster, which would help farmers make better decisions.

There are more perks to using ResNet-50 in precision agriculture than just making classification more accurate. Automating the process of categorising might make it easier to handle crops, make better use of resources, and grow more crops. Differentiating between types of rice can help farmers get the most out of their land. It also makes it possible to come up with new ways to deal with problems like getting rid of trash and bugs. ResNet-50 is a type of deep learning model that makes real-time monitoring and decision-making possible in farmland. These devices always show pictures of rice fields from above and from the ground. By using this method, anyone can get up-to-date information on the health of crops, their rate of growth, and the weather conditions at the moment. It's possible that this approach will help farmers find problems or stresses in their crops early on, before they get worse and cause big problems. Because of this, they can quickly fix the problem and keep working without any problems.

In order to successfully deal with problems that are always changing, the agriculture business can benefit from using deep learning models. These gadgets are flexible and can be easily changed to work in different situations. Through the use of transfer learning techniques, ResNet-50 can adapt to and do well in a variety of agricultural settings. The programme can do this job no matter how big or complicated the files are. Because it is flexible, the model can be used in a wide range of farming situations, from small family farms to big corporate companies. ResNet-50 can help with precision gardening and rice sorting, according to the study report. Deep learning can help farming methods keep getting better in the near future so that everyone has access to enough food and food is managed properly. ResNet-50 and other models show how far deep learning has come in a very short time. In the future, this could have big effects on how farmers do their work. The speed and efficiency of the world food supply line are likely to go up because of this.

### **8. REFERENCE**

- [1] Ahmed, T., Abid, M.F.M., Sultana, L., Al Noman, A. and Ali, U.S., 2023. The problems and solution of deep learning based rice disease detection in natural scene images.
- [2] Chenna, D., 2023. Evolution of Convolutional Neural Network (CNN): Compute vs Memory bandwidth for Edge AI. arXiv preprint arXiv:2311.12816.
- [3] Dutschmann, T.M., Kinzel, L., Ter Laak, A. and Baumann, K., 2023. Large-scale evaluation of k-fold cross-validation ensembles for uncertainty estimation. Journal of Cheminformatics, 15(1), p.49.

- [4] Fulton, L.V., Dolezel, D., Harrop, J., Yan, Y. and Fulton, C.P., 2019. Classification of Alzheimer's disease with and without imagery using gradient boosted machines and ResNet-50. *Brain sciences*, 9(9), p.212.
- [5] Huang, L., Qin, J., Zhou, Y., Zhu, F., Liu, L. and Shao, L., 2023. Normalization techniques in training dnns: Methodology, analysis and application. *IEEE transactions on pattern analysis and machine intelligence*, 45(8), pp.10173-10196.
- [6] Ishida, T., Yamane, I., Sakai, T., Niu, G. and Sugiyama, M., 2020. Do we need zero training loss after achieving zero training error?. *arXiv preprint arXiv:2002.08709*.
- [7] Khan, S.A. and Rana, Z.A., 2019, February. Evaluating performance of software defect prediction models using area under precision-Recall curve (AUC-PR). In 2019 2nd International Conference on Advancements in Computational Sciences (ICACS) (pp. 1-6). IEEE.
- [8] Kiruthika, S. and Karthika, D., 2024. Crop recommendation system and pest classification using weighted support vector machine on climate data. *Salud, Ciencia y Tecnología-Serie de Conferencias*, 3, pp.757-757.
- [9] Labini, P.S., Cianfriglia, M., Perri, D., Gervasi, O., Fursin, G., Lokhmotov, A., Nugteren, C., Carpentieri, B., Zollo, F. and Vella, F., 2021. On the anatomy of predictive models for accelerating GPU convolution kernels and beyond. *ACM Transactions on Architecture and Code Optimization (TACO)*, 18(1), pp.1-24.
- [10] Liao, L., Li, H., Shang, W. and Ma, L., 2022. An empirical study of the impact of hyperparameter tuning and model optimization on the performance properties of deep neural networks. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 31(3), pp.1-40.
- [11] Mishra, S., Mishra, D., Mallick, P.K., Santra, G. and Kumar, S., 2021. A Classifier Ensemble Approach for Prediction of Rice Yield Based on Climatic Variability for Coastal Odisha Region of India. *Informatica*, 45(3).
- [12] Mohamed, M., 2023. Agricultural Sustainability in the Age of Deep Learning: Current Trends, Challenges, and Future Trajectories. *Sustainable Machine Intelligence Journal*, 4, pp.2-1.
- [13] Moradi, R., Berangi, R. and Minaei, B., 2020. A survey of regularization strategies for deep models. *Artificial Intelligence Review*, 53(6), pp.3947-3986.
- [14] Oloyede, J. and Kawulok, M., 2019. Selecting training sets for support vector machines: a review. *Artificial Intelligence Review*, 52(2), pp.857-900.
- [15] Oloyede, M.O., Hancke, G.P. and Myburgh, H.C., 2020. A review on face recognition systems: recent approaches and challenges. *Multimedia Tools and Applications*, 79(37), pp.27891-27922.
- [16] Singh, S. and Krishnan, S., 2020. Filter response normalization layer: Eliminating batch dependence in the training of deep neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 11237-11246).
- [17] Sofaer, H.R., Hoeting, J.A. and Jarnevich, C.S., 2019. The area under the precision-recall curve as a performance metric for rare binary events. *Methods in Ecology and Evolution*, 10(4), pp.565-577.