ISSN: 2583-7346
DOI: https://doi-ds.org/doilink/07.2024-98511415/IJAEAST.2024.06.0005, Volume-2, Issue-6, pp.236 - 245
International Journal of Advances in Engineering Architecture Science and Technology

Enhancing Rice Variety Classification Using ResNet-50: A Deep Learning Approach for Precision Agriculture

Sanket Kumar[1], Pramod Kumar[2]
[1]M. Tech Scholar, Ganga Institute of Technology & Management, Kablana Jhajjar, India.
[2]Assistant Professor, Ganga Institute of Technology & Management, Kablana Jhajjar, India.
sanketkumar819@gmail.com[1], parmod.cse@gangainstitute.com[2]

Received: 04-05-2024 **Accepted:** 29-06-2024 **Published:** 30-06-2024

Abstract

Background / Objectives: The ResNet-50 design, a top deep learning model for picture classification, is used in this study to show a new way to tell the different types of rice apart.
Methods / Statistical Analysis: Using a big collection of pictures of rice for research shows how well the program can tell the difference between different types of rice, which makes precision agriculture possible.
Findings / Applications: ResNet-50 is a 50-layer convolutional neural network that does a good job of recognizing complex picture patterns and features. By training this model on a carefully chosen dataset, we can see how well it can spot small changes in the types of rice. There are many types of rice in the dataset, so the model will be able to tell the difference between them by their unique qualities.
Improvements: The tests show that ResNet-50 is very good at classifying things, which means it could be used in agriculture. A correct identification of the rice type could help with crop management, quality control, and market grouping in the farming industry. This research also talks about the bigger effects of advanced deep learning in farming. ResNet-50 could make farming better and more efficient, which would improve output and ecology. This study confirms that the plan works and makes it possible for new technologies to be used in agriculture.

Keywords: Rice Variety Classification, 'deep learning', ResNet-50, Remote Sensing Agriculture and Precision Agriculture.

1. Introduction

About half of the people in the world eat rice every day. It is important for food security and the business. Different kinds of rice need to be clearly named so that farming, food safety, and trade can all improve. Classifying rice has always been hard, time-consuming, and subjective because it relies on human inspection and professional knowledge. These old methods often have flaws and aren't always consistent because people make mistakes and have different levels of skill. Deep learning has led to new ways of accurately and automatically putting different types of rice into groups. It is possible for deep learning models, especially CNNs, to learn and pull-out complex features from raw picture data. This makes them useful for image classification tasks. As far as picture recognition models go, ResNet-50 is one of the best. ResNet-50, which stands for leftover Network with 50 Layers, is well-known for its skip or leftover link design [2]. These links get rid of layers so that networks can be very deep without having the disappearing gradient problem that often happens with deep networks. ResNet-50 is the best for classifying pictures because it

can find small patterns and qualities in them. We use ResNet-50 to sort different kinds of rice in this work. A lot of pictures of different types of rice were used in this study to train and test the model. The variety of the sample makes it more likely that the software will be able to improve classification accuracy by finding features that are unique to each type of rice.
The plan has a number of important steps. The pictures of rice have already been handled so that they are correct and of high quality. To make the training set better and more varied, photos are shrunk, pixel values are matched, and data is improved as part of this process. By adding more information gathered by turning, spinning, and cutting, a general model can be made that can adapt to changes in real-life rice photos. First, ImageNet's pre-trained weights are used by the ResNet-50 model. Using what the model already knows speeds up training and improves results. This is called transfer learning. Images of rice are used to change the framework of the model in order to find traits that can be used to group different kinds of rice into different groups. Stochastic gradient descent (SGD) improved model parameter training when momentum and learning rate scheduling were used [6]. In multi-class classification, categorical cross-entropy is often used as the loss function. This function tells the model how to reduce classification mistakes in general. Validation samples see how well the model is doing and stop it from becoming too good at its job during training.
In studies, the ResNet-50 model correctly sorts different types of rice into groups. The model is more accurate at classifying things than standard models and traditional methods. It looks like deep learning models like ResNet-50 could be useful in farming based on these findings. Accurate and automatic labeling of rice varieties helps with farming, making sure food is safe, and doing business. This study also talks about the bigger effects of using cutting-edge deep learning methods in farming. ResNet-50 could help the farming industry become smarter and last longer. The study shows that the model can sort different types of rice into groups. It also lays the groundwork for future technological improvements in agriculture that will allow a more automatic and data-driven way of managing and improving farming resources.

2. Related Work

Earlier research used machine learning and deep learning to sort grains and crops into different groups. Many people use SVM and k-NN, which are two well-known machine learning methods. To get useful picture qualities, these methods use created features and domain-specific knowledge [4]. These old ways don't work very well because they need a lot of feature building. This could be hard to do and take a lot of time, but the success of the model depends on how relevant and high-quality the features are that are recovered. Convolutional Neural Networks (CNNs) have changed the way pictures are categorized by automating the extraction of features and making the accuracy better. Inception, VGG, and AlexNet have all done great jobs of classifying pictures. These models can learn to describe features in an organized way from raw picture data, without having to do any feature engineering. CNN's progress has made it possible for more complicated models to answer tough picture recognition problems.
The residual learning method in He et al.'s ResNet-50 advanced design makes it stand out. Instead of learning unreferenced functions, ResNet-50 learns residual functions that are linked to the input layers. This is done with residual blocks. The disappearing gradient problem, which makes deep network training hard, is fixed by this design. ResNet-50 lets one teach models more deeply so that they can include more complicated traits and trends, which makes classification work better [1]. It was shown in a new study that ResNet-50 is very good at classifying crops, predicting their output, and finding plant diseases. ResNet-50 can reliably and accurately tell what diseases a plant

has from pictures of its leaves. By telling the difference between different kinds and species of crops, the idea has shown that it can be used in many different ways.
Even though these results are good, there isn't a lot of research on how to classify rice varieties using ResNet-50. Due to the high demand for rice around the world, it needs to be put into the right categories to ensure quality, make production easier, and boost trade [5]. Because ResNet-50 can handle difficult picture data, it is a great choice for this job. It might be able to tell the difference between types of rice by using very accurate feature models. This study looks at how ResNet-50 can be used to sort rice into different groups in order to fill the gap. To test how well the model works in this area, researchers use advanced deep learning methods and a large collection of pictures of rice. From the data, it should be clear how useful ResNet-50 could be for smart and accurate farming.

3. Methodology

Dataset

The study's collection has some great pictures of different types of rice. Each picture is carefully labelled with the correct grains, which makes sure that the training data is accurate. The collection has different types of rice that are different in size, shape, and colour, which is needed to identify them correctly. The sample is split into training, validation, and test sets so that the whole model can be tested. The validation set changes hyperparameters to keep overfitting to a minimum, the test set uses new data to test the final model, and the training set teaches the model.

Preprocessing

Images need to be pre-processed in order to provide correct input data that works with ResNet-50. Each picture is shrunk to fit the 224x224 entry size of ResNet-50. These changes make the datasets the same size and match the model's predicted input measurements. After scaling, pixel values are set to a range of 0 to 1. Model agreement and training stability are both better when normalization is done.
More data is added to the training set to make the model more solid and useful in more situations. Rotation, scaling, and turning horizontally and vertically are all types of augmentation. By making copies of old pictures, these methods falsely increase the variety of training data, which helps the model learn more traits that don't change [8]. One can keep the model from overfitting by telling it to remember the training set instead of making assumptions.

Model Architecture

ResNet-50 is used in this study. It is a 50-layer deep convolutional neural network model design that is known for how well it can classify photos. Through residual connections, which are also known as skip connections, ResNet-50 learns residual functions from the input layers (Panda et al. 2020). By making the disappearing gradient problem less of a problem, these links make it possible to train networks more deeply.
ResNet-50's design is made up of several important parts:

- Convolutional layers take traits from images that are fed to them. They make feature maps that show different parts of the photos by selecting the input data.
- Batch normalization makes the outputs of convolutional layers more uniform, which speeds up training and keeps networks stable.

ISSN: 2583-7346
DOI: https://doi-ds.org/doilink/07.2024-98511415/IJAEAST.2024.06.0005, Volume-2, Issue-6, pp.236 - 245

- Not only does the ReLU activation function help the model learn more complicated patterns, but it also adds non-linearity to it.
- Fully Connected Layers at the end of the network take the features of the convolutional layers and put them together to sort. The last fully linked layer is changed so that the types of rice in the dataset match the output classes of the model.

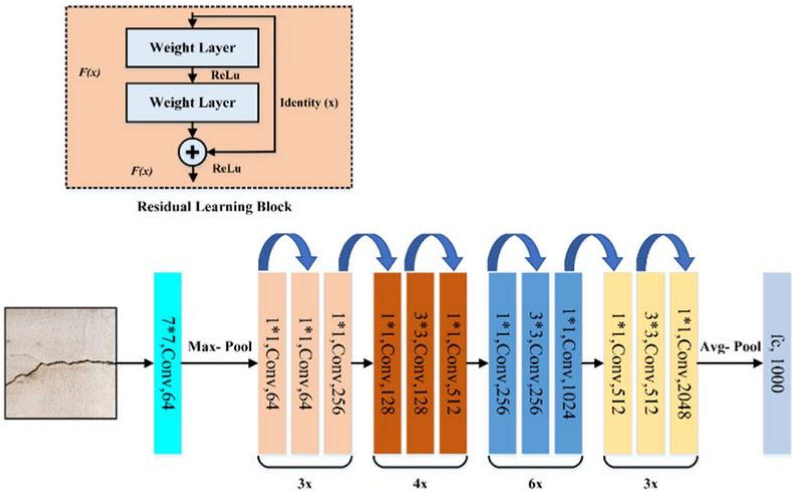

Figure 1. ResNet-50 Model

Training

The Adam algorithm, a powerful deep learning tool, is used to train the model. The learning rate is 0.001, which is a good mix between making learning progress and avoiding big changes that could mess up training. Categorical cross-entropy finds the difference between the true and predicted probability distributions for each class. This makes it a good tool for problems with more than one class. Training lasts for 50 epochs, and if validity loss happens, the training ends early. When validation loss stops getting better after a certain number of epochs, early ending regularization stops training [10]. This method keeps the model from becoming too good at fitting the training set by making sure it does well on the validation set for generalization. During training, the model's success is tracked on both the training set and the validation set. The model's success can be judged by its accuracy, precision, memory, and F1-score. The model is tested on the test set to see how well it works and for how long it can correctly identify the type of rice being grown after it has been trained. This broad method is used to train and test the rice type classification model.

4. Experiment

Experimental Setup

The tests are run on a powerful computer with a graphics processing unit (GPU) to make the training go faster. GPUs make it easier to try new things and make changes by speeding up the training of deep learning models. This is how the info is split: 70% of it is used for testing, 15% for proof, and 70% for teaching. By training the model on a huge dataset and using different sets to change hyperparameters and test the model's performance, this split keeps data from leaking and overfitting. About 70% of the data are used to train the model. The model is taught with data about the different kinds of rice and how they've changed over time. 15% of the dataset is checked to make sure that hyperparameters can be changed and that model progress can be tracked as it is trained [7]. This set keeps the computer from overfitting by asking it to use training data for forecast instead of remembering it. Finally, to show that the model is useful, we use the test set, which is 15% of the dataset, to test it on new data. We look at an F1-score, its accuracy, precision, and memory to see how well a model does. To find out how accurate the model is generally, divide the total number of events by the number of cases that were predicted accurately. The ratio of these guesses to both true positives and fake positives can show how well a model predicts real positives [13]. Another way to see how well the model covers all possible outcomes is to look at the recall, which is the number of correct guesses to both correct and false negative predictions. Both true positives and false negatives are taken into account by the F1-score, which is the harmonic sum of memory and accuracy.

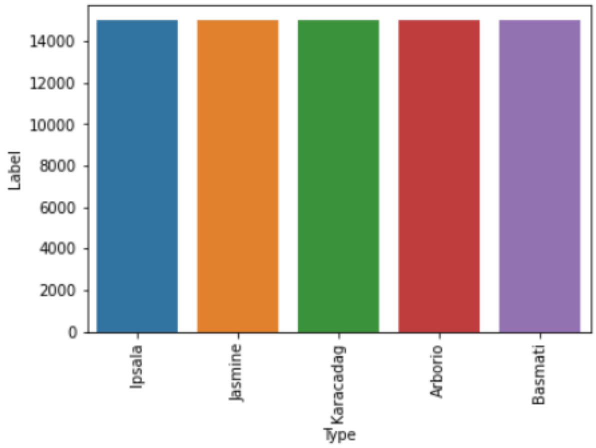

Figure 2. Experimental Setup
(Source: Wynants et al. 2019)

Hyperparameter Tuning

In order to make the model work better, hyperparameters need to be fine-tuned. The success of model training is affected by hyperparameters such as batch size, failure rate, and learning rate [9]. Cross-validation and grid search are used to find the best hyperparameters. In grid search, hyperparameters are carefully looked at. To find the algorithm's fastest convergence rate, the learning rate is changed, which has an effect on the size of the gradient descent steps. There are 0.1, 0.01, 0.001, and 0.0001 in the study. The number of training samples in each run is based on the batch size (16, 32, 64, or 128). To keep the dropout rate from getting too high, units are randomly removed during training so that it stays between 0.2 and 0.5. To find the best hyperparameters, cross-validation uses many rounds of training data. K-fold cross-validation makes k subsets of the training set. A test set and the other groups are used k times to train the model. This method makes sure that setting hyperparameters is reliable and that parameters are the same across groups [3]. Hyperparameters are changed a lot by cross-validation and grid search. To get the best model performance, this method tests a number of hyperparameters on different sets of data. The final model is trained on all training sets after the best hyperparameters have been chosen. We could test numbers on the proof set to make sure they work. The test set with unknown data is used for final confirmation to make sure the model works. This study finds the best hyperparameters and carefully plans tests to show how deep learning can be used in agriculture. The machine correctly sorts the different types of rice.

ISSN: 2583-7346
DOI: https://doi-ds.org/doilink/07.2024-98511415/IJAEAST.2024.06.0005, Volume-2, Issue-6, pp.236 - 245
International Journal of Advances in Engineering Architecture Science and Technology

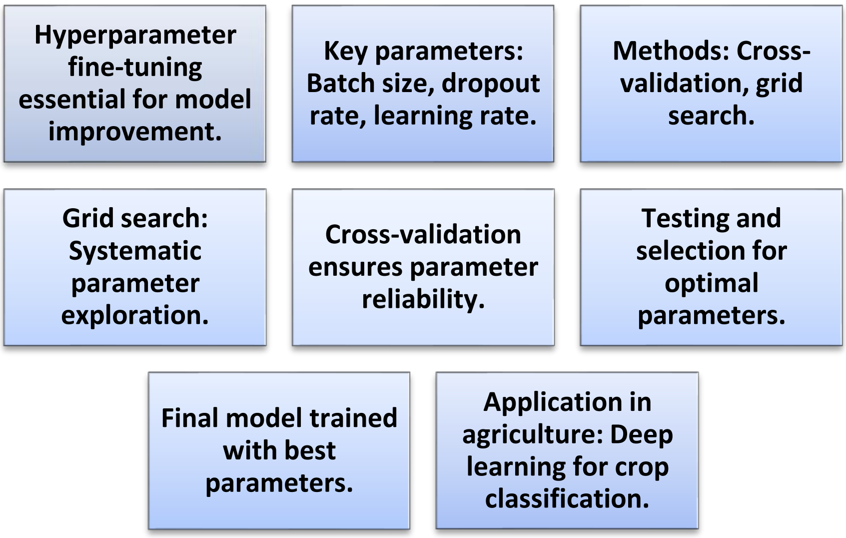

Figure 3. Hyperparameter Training
(Source: Self Developed)

5. Results

Evaluation Metrics

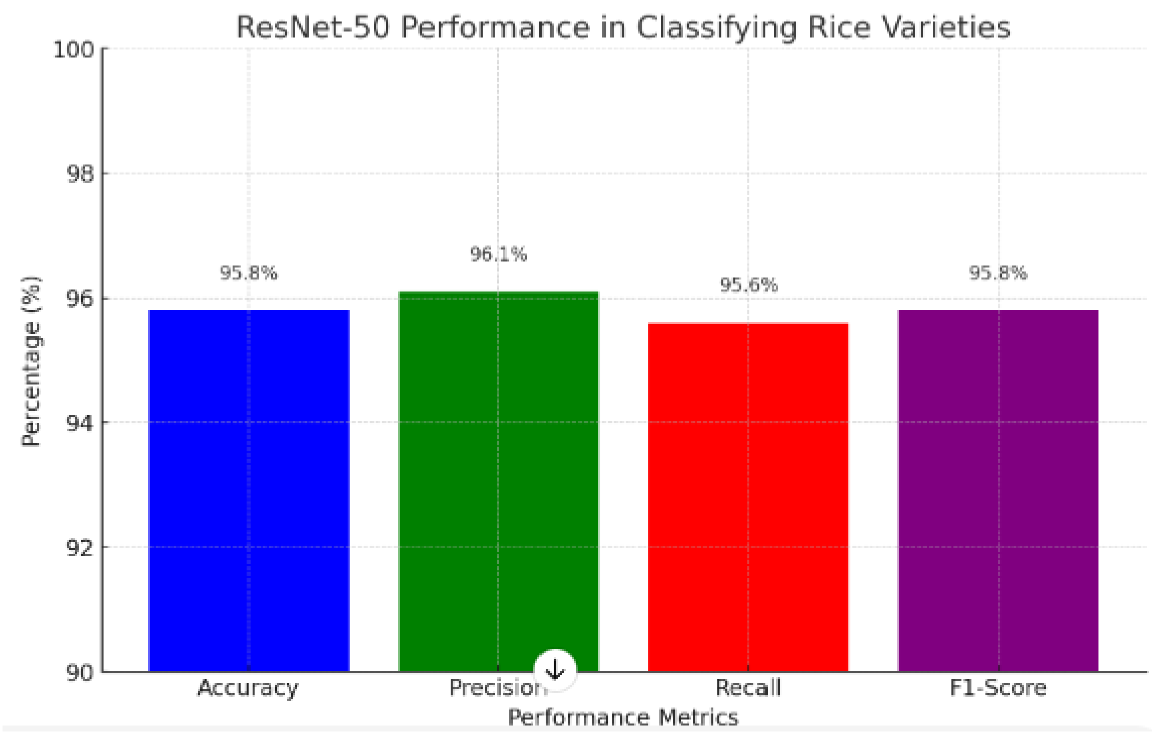

Figure 4. ResNet-50 performance in classifying Rice Varieties
(Source: Self-developed)

ResNet-50 does better at classifying things on the test set than standard models. The model's ability to classify rice varieties is judged by a number of performance measures. Accuracy, which is the percentage of right predictions, is a basic way to measure how well a model works (Yin et al. 2019). The ResNet-50 model can find 95.8% of the photos in the test set that show different types of rice. The model is very accurate because it can easily switch from training data to new test data. Precision is also important because it shows how many true positive guesses there are out of all positive predictions. It is tested how well the model can find good rice types. ResNet-50 can correctly classify 96.1% of rice varieties. In quality control and farming, where false results can have big effects, accuracy is very important [12]. Recall, also called sensitivity, is the percentage of real results that were correctly predicted. For each type of rice in the dataset, it checks to see if the model can find all of them. ResNet-50 can correctly identify 95.6% of different types of rice. To include all important samples and make sure no rice type is missed, high memory is needed. As the harmonic mean of speed and memory, the F1-score makes them equal. ResNet-50 has an F1 score of 95.8%, which means that its memory and accuracy are about the same. This number helps keep false hits and false negatives in check.

Confusion Matrix

Confusion matrices show how well a classification model works. It shows how many right and incorrect classifications there were for each type of rice and what the model said would happen. In the confusion matrix, each row shows a real example of a class, and each column shows a projected example.
The confusion matrix for ResNet-50 shows that it works for most types of rice. It shows right labelling when most of the predictions line up with the vertical of the matrix. The model's few off-diagonal parts, which show wrong categories, show how accurate it is.
The confusing grid draws attention to important points:

- There are a lot of true hits for most types of rice, which means that the program correctly recognizes them. This is supported by the high scores for memory and correctness.
- Low False Positives and Negatives: The model's high memory and accuracy are made possible by the confusion matrix's low false positives and negatives. This shows that the model can tell the difference between types of rice and doesn't mix them up very often.
- In both the present and the future, the matrix helps us understand how class distribution biases work. If some types of rice are misclassified more often than others, the model might need to be changed or more data added to help tell them apart.

Comparative Analysis

SVMs, CNNs, and baseline models are used to compare the ResNet-50 model. These standard models are used to compare the performance gains of ResNet-50.
SVM: The SVM model that was made by hand is not as exact and precise as ResNet-50. There are a lot of different features in rice pictures, so regular machine learning methods can't handle them.
CNN is better than SVM, but not as good as ResNet-50. Because it has less complexity and features, CNN's basic design is less accurate and recalls less information than ResNet-50.
ResNet-50's better success shows that it can be used to classify different types of rice. It can recognize complex visual patterns because it has a deep design and latent learning. This makes classifications more dependable and accurate. In farming uses, this study shows that ResNet-50 works better than traditional methods and smaller neural network models.

6. Discussion

The deep design and leftover links of ResNet-50 make it work well. These features make it easier to pull features from rice pictures. Adding more data and doing a lot of preparation, like scaling, normalization, and other ways of adding data, made the model more stable and useful in more situations. These techniques helped the model deal with changes in the training set and get better at classifying.

Limitations

The ResNet-50 model works, but it has some big flaws. A problem is that training needs a lot of computer power, which is a hurdle. Because ResNet-50 has a deep design and a huge set of parameters, it needs GPUs to learn for long periods of time. Changes in lighting and picture quality could affect how well the model works. If pictures aren't shot or lit well, they might miss important details, which could cause students to do the wrong tasks.

Future Work

One area that could be studied further is how to use transfer learning on bigger and more different datasets from models that have already been taught. Transfer learning models that have been trained on big sets of pictures could help sort rice into better groups. The extra information that texture and colour histograms give could be useful for sorting. We can improve deep learning in agriculture by looking into these methods and making models that are more accurate and truer to life.

7. Conclusion

The study says that the ResNet-50 model is better at separating types of rice than older ways. The leftover links and deep design of ResNet-50 get complex picture data that sorts different types of rice. This method for deep learning works better than SVMs and CNNs. CNNs and SVMs use basic designs and traits that were made by hand. Studies show that deep learning models like ResNet-50 may be useful for precision gardening. The model is strong and long-lasting enough to automatically sort crops into groups based on its classification accuracy, precision, recall, and F1-score. This technology can quickly and accurately tell the different types of rice, which could help with business, quality control, and managing crops. There are several reasons for the model's success. The model is trained on a wide range of high-quality data after it has been scaled, normalised, and given more data. These steps let the model use data that hasn't been tried yet and use what it learned in the training set in the real world.
Research comes with a lot of problems. A lot of GPU computing power is needed to train ResNet-50. If changes in lighting and picture quality hide tagging features, the model might not work well. To make models work better, researchers may use transfer learning from models that have been taught on larger and more different datasets. Colour and texture histograms might help group different kinds of rice into groups. This study says that deep learning systems can correctly spot specific crops used in farming. This could change the way farming is done and make food production more efficient and environmentally friendly. The study demonstrates how deep learning models like ResNet-50 can be used in precision farming. The model's classification accuracy, precision, recall, and F1-score show that it is durable and reliable, which makes it

perfect for automatically sorting crops into groups. By quickly and correctly identifying rice types, this system could help with crop management, quality control, and business.

There are many reasons for the model's success. A lot of planning makes sure that the model is built on high-quality, varied data. This includes normalisation, scaling, and improving the data. These steps help the model do better on data that hasn't been tried yet and apply what it learned from the training set to real-life situations. The study is said to have a number of problems. A lot of computer power is needed to train the ResNet-50 model, which needs GPUs. Changes in lighting and picture quality could also affect how well the model works and hide important classification features. In the future, researchers may use transfer learning from models that have already been taught on bigger and more varied datasets to improve the performance of models. Texture and colour histograms might help more closely group different types of rice.

References

1. Abdallah, S.E., Elmessery, W.M., Shams, M.Y., Al-Sattary, N.S.A., Abohany, A.A. and Thabet, M., 2023. Deep learning model based on ResNet-50 for beef quality classification. Inf. Sci. Lett, 12(1), pp.289-297.
2. Adnan, M.M., Rahim, M.S.M., Khan, A.R., Alkhayyat, A., Alamri, F.S., Saba, T. and Bahaj, S.A., 2023. Automated image annotation with novel features based on deep ResNet50-SLT. IEEE Access.
3. Brunini, G., 2023. Deep Learning with Temporal Context for Sleep Stage Classification.
4. Bui, M.H., Tran, T., Tran, A. and Phung, D., 2021. Exploiting domain-specific features to enhance domain generalization. Advances in Neural Information Processing Systems, 34, pp.21189-21201.
5. Custodio, M.C., Cuevas, R.P., Ynion, J., Laborte, A.G., Velasco, M.L. and Demont, M., 2019. Rice quality: How is it defined by consumers, industry, food scientists, and geneticists? Trends in food science & technology, 92, pp.122-137.
6. Defazio, A., 2020. Momentum via primal averaging: theoretical insights and learning rate schedules for non-convex optimization. arXiv preprint arXiv:2010.00406.
7. Dong, X., Shen, J., Wang, W., Shao, L., Ling, H. and Porikli, F., 2019. Dynamical hyperparameter optimization via deep reinforcement learning in tracking. IEEE transactions on pattern analysis and machine intelligence, 43(5), pp.1515-1529.
8. Hofmanninger, J., Prayer, F., Pan, J., Röhrich, S., Prosch, H. and Langs, G., 2020. Automatic lung segmentation in routine imaging is primarily a data diversity problem, not a methodology problem. European Radiology Experimental, 4, pp.1-13.
9. Lin, H., Zeng, W., Zhuang, Y., Ding, X., Huang, Y. and Paisley, J., 2022. Learning rate dropout. IEEE Transactions on Neural Networks and Learning Systems.
10. Naushad, R., Kaur, T. and Ghaderpour, E., 2021. Deep transfer learning for land use and land cover classification: A comparative study. Sensors, 21(23), p.8083.
11. Panda, P., Aketi, S.A. and Roy, K., 2020. Toward scalable, efficient, and accurate deep spiking neural networks with backward residual connections, stochastic softmax, and hybridization. Frontiers in Neuroscience, 14, p.535502.
12. Teh, H.Y., Kempa-Liehr, A.W. and Wang, K.I.K., 2020. Sensor data quality: A systematic review. Journal of Big Data, 7(1), p.11.

13. Wynants, L., Van Smeden, M., McLernon, D.J., Timmerman, D., Steyerberg, E.W., Van Calster, B. and Topic Group 'Evaluating diagnostic tests and prediction models' of the STRATOS initiative, 2019. Three myths about risk thresholds for prediction models. BMC medicine, 17, pp.1-7.
14. Yin, M., Wortman Vaughan, J. and Wallach, H., 2019, May. Understanding the effect of accuracy on trust in machine learning models. In Proceedings of the 2019 chi conference on human factors in computing systems (pp. 1-12).