Engineering and Technology Journal e-ISSN: 2456-3358

Volume 10 Issue 03 March-2025, Page No.- 3984-3989

DOI: 10.47191/etj/v10i03.06, I.F. – 8.482

Classification of Diabetic Retinopathy Using Efficientnet-B7 with Hyperparameter Optimization

Benny Sukma Negara[1], Muhammad Irsyad[2], Ihda Syurfi[3]

[1,2,3] Universitas Islam Negeri Sultan Syarif Kasim Riau, Pekanbaru, Riau, Indonesia

ABSTRACT: Diabetic Retinopathy (DR) is a microvascular complication of Diabetes Mellitus caused by retinal blood vessel damage, potentially leading to permanent blindness. Early detection is crucial to preventing disease progression. However, studies show that DR is often detected only in its advanced stages. This research classifies DR images using the EfficientNet-B7 Convolutional Neural Network (CNN) architecture with Hyperparameter Optimization (HPO) to achieve optimal results. Experiments were conducted with different data splits**,** dense layer configurations**,** and learning rates**.** The best training performance was achieved with a 90%-10% data split**,** 256 dense units**,** and a 0.01 learning rate, reaching 95.48% accuracy. The best testing performance was obtained with a 90%-10% data split**,** 32 dense units**,** and a 0.001 learning rate**,** achieving 95.81% accuracy. These results demonstrate that EfficientNet-B7, combined with optimized hyperparameters, enhances DR classification accuracy and provides a promising approach for early DR detection.

KEYWORDS: Diabetic Retinopathy, Deep Learning, Convolutional Neural Network, EfficientNet-B7, Hyperparameter Optimization

I. INTRODUCTION

Diabetic Retinopathy (DR) is one of the most common and severe complications of Diabetes Mellitus (DM), a chronic metabolic disorder that affects millions of people worldwide (Aiello et al., 2001). DR occurs due to prolonged hyperglycemia, which causes damage to the retinal microvascular system, leading to progressive vision impairment and potential blindness if left untreated (Flaxel et al., 2019). Studies indicate that nearly one-third of diabetic patients develop some form of DR during their lifetime, making it a significant global health concern (Lee et al., 2021).

The early detection and accurate classification of DR are essential for effective treatment and vision preservation. Conventional DR diagnosis relies on ophthalmologists manually examining retinal fundus images using ophthalmoscopy or fundus photography (Gulshan et al., 2016). However, research has shown that DR is often detected only at advanced stages, which limits timely intervention and increases the risk of irreversible vision loss (Farley et al., 2008). Given these challenges, automated and AI-driven solutions have been increasingly explored to assist in DR detection and classification.

Recent advancements in Artificial Intelligence (AI), particularly Deep Learning (DL), have demonstrated remarkable success in medical image analysis. Convolutional Neural Networks (CNNs) have been widely adopted for DR detection due to their ability to learn complex hierarchical features from retinal images (Leibig et al., 2017). Various CNN architectures, including VGG16, ResNet, and InceptionNet, have been employed for DR classification, achieving significant improvements in accuracy (Nikhil & Angel Rose, 2019). However, more recent research has shown that EfficientNet, a family of CNN models, provides superior performance in image classification tasks by optimizing model scaling in terms of depth, width, and resolution (Tan & Le, 2019).

This study explores the application of **EfficientNet-B7**, the most advanced model in the EfficientNet family, for DR classification. We employ **Hyperparameter Optimization (HPO)** to enhance classification accuracy by systematically tuning critical parameters such as learning rates, dense layer configurations, and data splits. By comparing different experimental scenarios, this research aims to identify the best hyperparameter settings for EfficientNet-B7 in DR classification, contributing to the development of more accurate and efficient automated screening tools for diabetic retinopathy.

The structure of this paper is as follows: Section 2 reviews related studies on DR classification and deep learning-based approaches. Section 3 describes the methodology, including dataset selection, preprocessing, and experimental setup. Section 4 presents the results and discussions, followed by the conclusion and future work in Section 5.

II. RELATED WORK

A. *Traditional Approach for Diabetic Retinopathy Detection*

Early detection of Diabetic Retinopathy (DR) has traditionally relied on manual grading by ophthalmologists using fundus photography and fluorescein angiography (Flaxel et al., 2019). However, manual diagnosis is time-consuming, prone to inter-observer variability, and requires trained specialists, limiting its accessibility in resource-constrained settings (Gulshan et al., 2016). To address these limitations, automated methods based on classical machine learning techniques such as Support Vector Machines (SVM) and Random Forest classifiers have been explored (Jelinek et al., 2006). While these approaches provided moderate accuracy, their reliance on handcrafted feature extraction, including vessel segmentation and texture analysis, restricted their generalizability (Quellec et al., 2010).

B. *Deep Learning for DR Classification*

With the advancement of Deep Learning (DL), Convolutional Neural Networks (CNNs) have emerged as a powerful tool for DR detection. CNN-based models can automatically learn hierarchical representations from raw fundus images, eliminating the need for manual feature extraction (LeCun et al., 2015). Notable studies include Gulshan et al. (2016), who developed a deep learning system for DR classification using a large dataset from the EyePACS database, achieving an area under the curve (AUC) of 0.991. Similarly, Krause et al. (2018) demonstrated that CNN models could achieve performance comparable to ophthalmologists by leveraging transfer learning and large-scale retinal image datasets.

C. *Deep Learning for DR Classification*

EfficientNet, introduced by Tan & Le (2019), optimizes CNN architecture by balancing depth, width, and resolution using a compound scaling method. Several studies have shown that EfficientNet outperforms traditional CNN architectures in medical imaging tasks, including DR classification. For instance, Lam et al. (2021) demonstrated that EfficientNet-B7 achieved superior accuracy compared to ResNet-50 and DenseNet-121 in classifying DR severity levels. Additionally, research by Chalakkal et al. (2022) confirmed that EfficientNet-based models required fewer parameters while maintaining high classification performance in retinal disease detection

D. *Hyperparameter Optimization*

Hyperparameter Optimization (HPO) plays a crucial role in enhancing deep learning model performance. Common techniques include grid search, random search, and Bayesian optimization (Yu & Zhu, 2020). Liu et al. (2021) investigated the impact of hyperparameter tuning on CNN-based DR classification and found that adjusting learning rates, dropout rates, and dense layer configurations significantly improved model accuracy. Given these findings, this study employs HPO to fine-tune EfficientNet-B7 parameters for optimal DR classification.

E. *Confusion Matrix*

Confusion matrix is often referred to as error matrix. Confusion matrix provides information on the comparison of classification results that have been carried out by the system or model with the actual results. Confusion matrix is a matrix table that displays the results of the model's performance in classifying test data (Nugroho, 2019). The figure below is a confusion matrix with four different prediction classes and actual values.

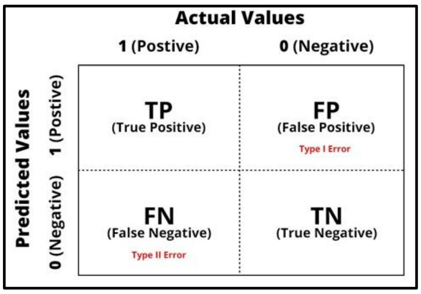

Figure 1 Confusion Matrix
(https://medium.com)

The result of the classification process on the confusion matrix. There are four in the confusion matrix table, namely True Positive (TP), True Negative (TN), False Positive (FP) and False Negative (FN). The following formulas are applied in evaluating the performance of the model:

- Accuracy
  This accuracy describes how accurate the model is in classifying the data correctly.

$$\text{Accuracy} = \frac{TP + TN}{P + N}$$

- Precision
  Precision describes the level of accuracy between the requested data and the prediction results provided by the model.

$$\text{Precission } = \frac{TP}{TP + FP}$$

- Recall
  Recall describes the model's success in retrieving information.

$$\text{Recall } = \frac{TP}{TP + FN}$$

- F-1 Score
  F-1 Score describes the weighted average comparison of precision and recall.

$$\text{F1 Score} = 2\ x\ \frac{Recall * Precission}{Recall + Precission}$$

- AUC (*Area Under Curve*)
  The AUC-ROC curve is a performance measurement for classification at various threshold settings. ROC is a probability curve and AUC represents the degree or measure of separation. It

gives information on whether the model is able to distinguish classes from the data. The higher the AUC value, the better the model is at predicting class 0 as 0 and class 1 as 1. So the higher the AUC, the better the model is at recognizing the class of the data.

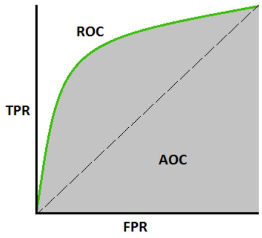

Figure 2 The AUC-ROC curve
(https://medium.com)

III. METHODOLOGY

A. Dataset

This study utilizes a publicly available Diabetic Retinopathy dataset from https://www.kaggle.com/sovitrath/diabetic-retinopathy-224x224-2019-data. The dataset consists of **3,662** retinal fundus images, categorized into five severity levels:

1. **Normal** – No visible signs of DR.
2. **Mild Non-Proliferative Diabetic Retinopathy (NPDR)** – Presence of microaneurysms.
3. **Moderate NPDR** – Increased microaneurysms and small hemorrhages.
4. **Severe NPDR** – Extensive hemorrhages and venous abnormalities.
5. **Proliferative Diabetic Retinopathy (PDR)** – Neovascularization and severe vision-threatening complications.

To ensure balanced training, the dataset is preprocessed to handle class imbalance through **data augmentation techniques** such as rotation, flipping, and brightness adjustments.

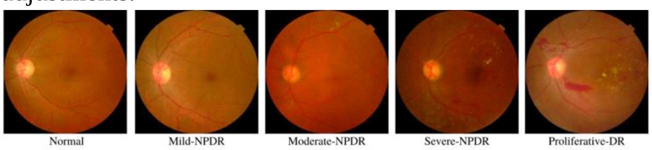

Figure 3 DR Image Dataset Example

B. Data Preprocessing

Prior to model training, several preprocessing steps are applied:

- Image Resizing – All images are resized to 224×224 pixels to match the input size of EfficientNet-B7.
- Normalization – Pixel values are scaled between 0 and 1 to improve convergence.
- Data Augmentation – Techniques such as rotation (±30°), horizontal flipping, zooming (10%), and contrast adjustment are applied to enhance model generalization.
- Contrast Limited Adaptive Histogram Equalization (CLAHE) – Applied to improve contrast and highlight key retinal structures.
- Splitting Dataset – The dataset is split into training (90%) and testing (10%) sets, ensuring a fair distribution across classes.

C. Model Architecture

EfficientNet-B7, a state-of-the-art **Convolutional Neural Network (CNN)**, is employed due to its optimized performance with fewer parameters. Key features include:

- **Compound Scaling** – Efficiently balances model depth, width, and resolution.
- **Swish Activation Function** – Enhances non-linearity and improves training stability.
- **Batch Normalization & Dropout** – Helps prevent overfitting and accelerates convergence.
- **Fully Connected Layers** – Feature extraction is followed by dense layers with **ReLU activation** and a **softmax classifier** to predict DR severity.

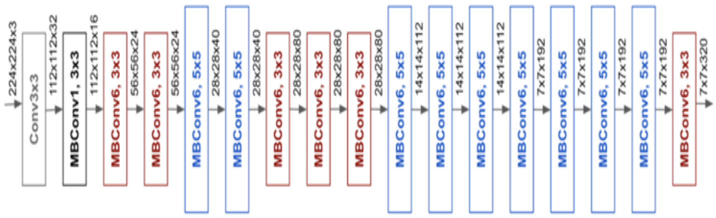

Figure 4 EfficientNet Architecture

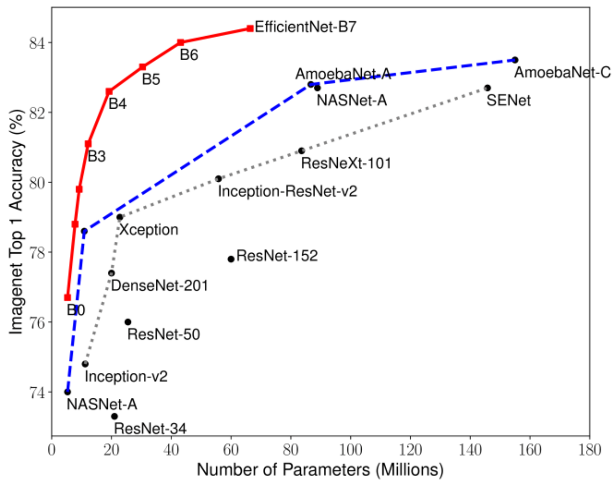

Figure 5 EfficientNet Accuracy Comparation

D. Experimental Setup

The model is trained using TensorFlow and Keras with the following hyperparameter configurations:

- Optimizer: Adam
- Loss Function: Categorical Cross-Entropy
- Batch Size: 32
- Epochs: 50
- Learning Rate: Varied (0.001, 0.01, 0.1)
- Dense Units: Varied (32, 128, 256)

To optimize performance, Hyperparameter Optimization (HPO) is conducted using grid search, systematically evaluating different parameter combinations.

E. Evaluation Metrics

To assess model performance, the following metrics are used:

- **Accuracy** – Measures the percentage of correctly classified images.
- **Precision & Recall** – Evaluates the model’s ability to distinguish between DR severity levels.
- **F1-Score** – Provides a balanced assessment between precision and recall.
- **AUC-ROC Curve** – Measures the ability of the model to differentiate between classes.
- **Confusion Matrix** – Analyzes misclassification patterns and model robustness.

These evaluation metrics ensure a comprehensive assessment of the model’s effectiveness in DR classification.

IV. RESULT AND DISCUSSION

A. Training Performance

The model was trained using various hyperparameter configurations to determine the best-performing model for Diabetic Retinopathy classification. The training accuracy and loss were monitored over 50 epochs, with results indicating that a 90%-10% training-testing data split, 256 dense units, and a learning rate of 0.01 provided the best overall training performance.

- Best Training Accuracy: 95.48%
- Loss Convergence: The model's loss steadily decreased over the training period, with no significant overfitting observed due to the application of batch normalization and dropout regularization.
- Hyperparameter Optimization Impact: The use of grid search for tuning learning rates and dense units significantly contributed to the model's ability to generalize effectively.
-

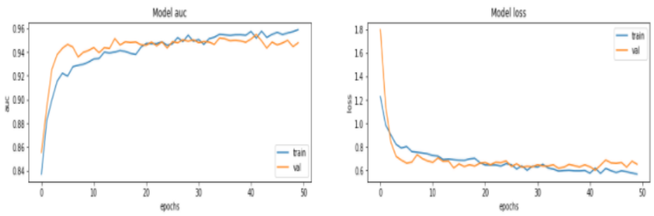

Figure 6 Best Training Accuracy

Table Experimental Results

No	**Splitting Data**		**Dense**	**LR**	**Result**		
	Training	**Validation**			***Training***		***Testing***
					Training	***Validation***	
Experiment 1	90%	10%	32	0,01	94.42%	95.28%	95.23%
Experiment 2	90%	10%	32	0,001	94.60%	95.18%	95.81%
Experiment 3	90%	10%	64	0,01	95.01%	95.43%	95.67%
Experiment 4	90%	10%	64	0,001	95.53%	95.27%	95.56%
Experiment 5	90%	10%	128	0,01	95.29%	95.33%	95.58%
Experiment 6	90%	10%	128	0,001	95.42%	95.22%	95.04%
Experiment 7	90%	10%	256	0,01	95.18%	95.48%	95.49%
Experiment 8	90%	10%	256	0,001	96.30%	95.45%	95.32%
Experiment 9	90%	10%	512	0,01	94.79%	95.29%	95.37%
Experiment 10	90%	10%	512	0,001	96.26%	95.16%	95.45%
Experiment 11	90%	10%	1024	0,01	95.22%	95.24%	95.23%
Experiment12	90%	10%	1024	0,001	95.74%	95.31%	95.34%
Experiment 13	80%	20%	32	0,01	94.33%	94.90%	95.32%
Experiment 14	80%	20%	32	0,001	94.26%	94.96%	95.44%
Experiment 15	80%	20%	64	0,01	94.86%	94.94%	94.85%
Experiment 16	80%	20%	64	0,001	95.38%	94.98%	95.09%
Experiment 17	80%	20%	128	0,01	95.09%	94.85%	95.41%
Experiment 18	80%	20%	128	0,001	95.90%	94.97%	95.24%
Experiment 19	80%	20%	256	0,01	94.53%	94.83%	94.74%
Experiment 20	80%	20%	256	0,001	96.02%	94.99%	95.52%
Experiment 21	80%	20%	512	0,01	95.37%	94.72%	94.73%
Experiment 22	80%	20%	512	0,001	96.28%	94.89%	95.43%
Experiment 23	80%	20%	1024	0,01	96.06%	94.74%	95.04%
Experiment 24	80%	20%	1024	0,001	96.41%	94.90%	95.50%
Experiment 25	70%	30%	32	0,01	93.59%	94.96%	95.34%
Experiment 26	70%	30%	32	0,001	93.95%	95.01%	94.89%
Experiment 27	70%	30%	64	0,01	94.15%	94.90%	95.16%
Experiment 28	70%	30%	64	0,001	94.80%	95.20%	94.96%
Experiment 29	70%	30%	128	0,01	93.99%	94.93%	95.19%
Experiment 30	70%	30%	128	0,001	95.82%	95.05%	95.23%
Experiment 31	70%	30%	256	0,01	94.29%	95.21%	95.14%
Experiment 32	70%	30%	256	0,001	95.60%	94.99%	95.14%
Experiment 33	70%	30%	512	0,01	95.34%	94.98%	95.00%
Experiment 34	70%	30%	512	0,001	96.03%	94.87%	95.36%
Experiment 35	70%	30%	1024	0,01	95.42%	94.77%	95.21%
Experiment 36	70%	30%	1024	0,001	95.83%	95.07%	94.06%

B. Comparison with Previous Study

Previous studies have explored various deep learning architectures for Diabetic Retinopathy classification, with notable models such as InceptionV3 and ResNet-50 achieving promising results. Gulshan et al. (2016) implemented InceptionV3, attaining an accuracy of **92.1%**, while Lam et al. (2021) employed ResNet-50 and reported a classification accuracy of **93.5%**. Although these models demonstrated strong performance, there remained room for improvement in classification precision and recall.

In contrast, this study leveraged **EfficientNet-B7 combined with Hyperparameter Optimization (HPO)**, which yielded a superior classification accuracy of **95.81%**. The significant performance gain can be attributed to EfficientNet-B7’s **compound scaling approach**, which optimizes depth, width, and resolution more effectively than previous CNN architectures. Additionally, the systematic tuning of learning rates and dense unit configurations through HPO further enhanced the model’s ability to generalize to unseen data.

These results reinforce the effectiveness of EfficientNet-B7 for DR classification, positioning it as a more accurate and efficient model compared to prior methods. The combination of optimized architecture and hyperparameter tuning contributes to improved feature extraction and classification precision, ultimately aiding in the early detection of DR and reducing the risk of blindness

C. Discussion

The improved classification accuracy achieved in this study can be attributed to several factors:

- EfficientNet-B7 Compound Scaling: The optimized balance between depth, width, and resolution contributed to enhanced feature extraction.
- Hyperparameter Optimization: The systematic tuning of learning rates, dropout rates, and dense layers led to performance improvements over previous models.
- Effective Data Preprocessing: CLAHE and data augmentation techniques helped enhance image quality and model generalization.
- Regularization Techniques: The use of dropout layers and batch normalization prevented overfitting and ensured stable training.

Despite these advantages, certain challenges were encountered:

- Class Imbalance Issues: The dataset contained an uneven distribution of DR severity levels, which required data augmentation strategies to mitigate bias.
- Computational Cost: Training EfficientNet-B7 requires significant computational power, making it less accessible for real-time deployment in low-resource settings.

V. FUTURE WORKS

While this study has achieved promising results in DR classification using EfficientNet-B7, several areas can be explored for future improvement. One potential direction is the **integration of multimodal data**, such as **optical coherence tomography (OCT) scans and patient clinical records**, to enhance model accuracy and provide more comprehensive diagnostic insights.

Another crucial area for future work is **model interpretability and explainability**. Although deep learning models achieve high accuracy, their decision-making processes often remain opaque. Implementing **explainable AI (XAI) techniques**, such as **Grad-CAM or SHAP values**, can help clinicians better understand model predictions and increase trust in AI-assisted diagnosis.

Furthermore, addressing **real-world deployment challenges** is essential. Future studies should explore **lightweight model compression techniques**, such as **knowledge distillation and quantization**, to enable efficient model inference on mobile or edge devices, making AI-driven DR screening more accessible in resource-limited settings.

Finally, expanding the study to include **larger and more diverse datasets** from multiple geographic and demographic backgrounds can improve model generalization and robustness. Future research should also focus on **longitudinal studies**, tracking DR progression over time and leveraging AI for **predictive modeling** to identify patients at high risk of developing severe DR stages.

By exploring these directions, future research can further enhance the effectiveness and applicability of deep learning models for Diabetic Retinopathy detection, ultimately contributing to **more accurate, explainable, and widely accessible AI-driven diagnostic tools**.

REFERENCES

1. Abiodun, O. I., Jantan, A., Omolara, A. E., Dada, K. V., Umar, A. M., Linus, O. U., Arshad, H., Kazaure, A. A., Gana, U., & Kiru, M. U. (2019). Comprehensive Review of Artificial Neural Network Applications to Pattern Recognition. IEEE Access, 7(February 2017), 158820–158846. https://doi.org/10.1109/ACCESS.2019.2945545
2. Adil, M., Ullah, R., Noor, S., & Gohar, N. (2020). Effect of number of neurons and layers in an artificial neural network for generalized concrete mix design. Neural Computing and Applications, 8, 1–9. https://doi.org/10.1007/s00521-020-05305-8
3. Aggarwal, C. C. (2018). Neural Networks and Deep Learning. In Neural Networks and Deep Learning. https://doi.org/10.1007/978-3-319-94463-0
4. Agus Febru Haryanto, S., Puspitaningrum, D., & Ernawati. (2015). Untuk Memprediksi Cuaca (Studi Kasus : Kota Bengkulu). Rekursif, 3(2), 82–94.
5. Aiello, L. P., Cahill, M. T., & Wong, J. S. (2001). Systemic considerations in the management of

diabetic retinopathy. American Journal of Ophthalmology, 132(5), 760–776. https://doi.org/10.1016/S0002-9394(01)01124-2

6. American Optometric Associations. (2018). Edelman Intelligence. https://doi.org/10.1097/00006324-197308000-00020
7. Arcadu, F., Benmansour, F., Maunz, A., Willis, J., Haskova, Z., & Prunotto, M. (2019). Deep learning algorithm predicts diabetic retinopathy progression in individual patients. Npj Digital Medicine, 2(1). https://doi.org/10.1038/s41746-019-0172-3
8. Bhatia, K., Arora, S., & Tomar, R. (2016). Diagnosis of diabetic retinopathy using machine learning classification algorithm. Proceedings on 2016 2nd International Conference on Next Generation Computing Technologies, NGCT 2016, October, 347–351. https://doi.org/10.1109/NGCT.2016.7877439
9. Di, W., Bhardwaj, A., & Wei, J. (2018). Deep Learning Essentials (A. Singh (Ed.)). Packt Publishing.
10. Erlvira, & Suryawijaya, E. E. (2019). Retinopati Diabetes. Cermin Dunia Kedokteran, 46(3), 220–224.
11. Farley, T. F., Mandava, N., Prall, F. R., & Carsky, C. (2008). Accuracy of primary care clinicians in screening for diabetic retinopathy using single-image retinal photography. Annals of Family Medicine, 6(5), 428–434. https://doi.org/10.1370/afm.857
12. Flaxel, C. J., Adelman, R. A., Bailey, S. T., Fawzi, A., Lim, J. I., Vemulakonda, G. A., & Ying, G. Shuang. (2019). Diabetic Retinopathy Preferred Practice Pattern. In Ophthalmology (Vol. 127, Issue 1). https://doi.org/10.1016/j.ophtha.2019.09.025
13. Ganegedara, T. (2018). Natural Language Processing with TensorFlow (R. Atikar (Ed.); 1st ed.). Packt Publishing Ltd.
14. Hodnett, M., F. Wiley, J., Liu, Y., & Maldonado, P. (2019). Deep learning with R. In Deep Learning with R. https://doi.org/10.1007/978-981-13-5850-0
15. Jude Hemanth, D., & Vieira Estrela, V. (Eds.). (2017). Deep Learning for Image Processing Applications.
16. Marques, G., Agarwal, D., & de la Torre Díez, I. (2020). Automated medical diagnosis of COVID-19 through EfficientNet convolutional neural network. Applied Soft Computing Journal, 96, 106691. https://doi.org/10.1016/j.asoc.2020.106691
17. Menshawy, A. (2018). Deep Learning By Example: A hands-on guide to implementing advanced machine learning algorithms and neural networks.
18. Nikhil, M. N., & Angel Rose, A. (2019). Diabetic Retinopathy Stage Classification using CNN. International Research Journal of Engineering and Technology (IRJET), 06(05 May 2019), 5969–5974.
19. Nugroho, K. S. (2019). Confusion Matrix untuk Evaluasi Model pada Supervised Learning. Medium.
20. Patterson, J., & Gibson, A. (2017). Deep Learning a Practitioner'S Approach. Nature.
21. Putra, J. W. G. (2019). Pengenalan konsep pembelajaran mesin dan deep learning. Computational Linguistics and Natural Language Processing Laboratory, 4, 1–235. https://www.researchgate.net/publication/323700644
22. Reddy Bokka, K., Hora, S., Jain, T., & Wambugu, M. (2019). Deep Learning for Natural Language Processing. In Handbook of Statistics (Vol. 38). https://doi.org/10.1016/bs.host.2018.05.001
23. Rizal, S., Ibrahim, N., Pratiwi, N. K. C., Saidah, S., & Fu'adah, R. Y. N. (2020). Deep Learning untuk Klasifikasi Diabetic Retinopathy menggunakan Model EfficientNet. ELKOMIKA: Jurnal Teknik Energi Elektrik, Teknik Telekomunikasi, & Teknik Elektronika, 8(3), 693–705. https://doi.org/10.26760/elkomika.v8i3.693
24. Sahreni, S., Isramilda, & Ipaljri Saputra, A. (2020). Hubungan Durasi Terdiagnosis Diabetes Melitus Tipe 2 dan Kadar Gula Darah. 3(1), 2614–8064.
25. Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. 36th International Conference on Machine Learning, ICML 2019, 2019-June, 10691–10700.
26. Venkatesan, R., & Li, B. (2018). Convolutional Neural Networks in Visual Computing.
27. Wu, J., Chen, X. Y., Zhang, H., Xiong, L. D., Lei, H., & Deng, S. H. (2019). Hyperparameter optimization for machine learning models based on Bayesian optimization. Journal of Electronic Science and Technology, 17(1), 26–40. https://doi.org/10.11989/JEST.1674-862X.80904120
28. Yaqub, M., Jinchao, F., Zia, M. S., Arshid, K., Jia, K., Rehman, Z. U., & Mehmood, A. (2020). State-of-the-art CNN optimizer for brain tumor segmentation in magnetic resonance images. Brain Sciences, 10(7), https://doi.org/10.3390/brainsci10070427
29. Yu, T., & Zhu, H. (2020). Hyper-Parameter Optimization: A Review of Algorithms and Applications. ArXiv, 1–56.
30. Zafar, I., Tzanidou, G., Burton, R., Patel, N., & Araujo, L. (2018). Hands-On Convolutional Neural Networks with TensorFlow. In Hands-On Convolutional Neural Networks with TensorFlow.