Financial Distress Prediction with Optimized Artificial Neural Networks: A Comparative Study of Optimization Algorithms

Ni Wayan Dewinta Ayuni[1], I Made Wijana[2],
I Made Dwijendra Sulastra[3], Ni Nengah Lasmini[4],
and Agus Adi Putrawan[5]

[1,2,3] Accounting Department, Politeknik Negeri Bali, Bali, Indonesia
[4] Information Technology Department, Politeknik Negeri Bali, Bali, Indonesia
dewintaayuni@pnb.ac.id

Abstract. Deep learning, a prominent branch of machine learning, has emerged as a powerful predictive modeling approach known for its high accuracy. One of its most essential components is the Artificial Neural Network (ANN), which emulates the human brain's neural structure, comprising thousands of interconnected neurons. Previous studies have highlighted the effectiveness of ANN in predicting financial distress, a condition indicating a company's financial difficulties and often serving as an early warning sign of potential bankruptcy. However, the predictive performance of ANN models is highly influenced by the optimization algorithms. Various optimization algorithms have been developed to improve ANN performance, including Stochastic Gradient Descent (SGD), Root Mean Square Propagation (RMSProp), Adaptive Gradient Descent (AdaGrad), Adadelta, and Adaptive Moment Estimation (Adam) optimizer. The application of these optimizers can yield varying model performance depending on the specific dataset or case. This study aims to apply and compare different ANN optimization algorithms to determine the most suitable one for enhancing the predictive performance of financial distress models in the Indonesian property and real estate sector. The findings reveal that the Adam optimizer outperforms other optimization algorithms, achieving the highest accuracy and superior AUC in predicting financial distress among companies in this sector.

Keywords: Artificial Neural Networks, Financial Distress, Optimization Algorithms

1 Introduction

Financial distress refers to a condition of financial difficulty experienced by a company, which serves as an early warning sign of potential bankruptcy. This condition represents a critical issue that must be anticipated and avoided. Corporate bankruptcy has significant and far-reaching impacts, affecting various stakeholders, including investors, management, employees, creditors, and the country's overall economic stability (Brenes et al., 2022). Therefore, predicting financial distress is a vital objective that enables companies to anticipate potential bankruptcy and take appropriate preventive measures (Liashenko et al., 2023). Traditional statistical approaches, such

A. A. N. G. Sapteka et al. (eds.), *Proceedings of the International Conference on Sustainable Green Tourism Applied Science - Engineering Applied Science 2025 (ICOSTAS-EAS 2025)*, Advances in Engineering Research 280,
https://doi.org/10.2991/978-94-6463-878-3_23

as discriminant analysis and logistic regression, have been widely used to predict corporate financial distress (Letkovsky et al., 2023). However, the limitations of these models in capturing the complex and nonlinear patterns inherent in financial data present challenges that must be addressed (Peralungal & Natchimuthu, 2024). In response, researchers in the era of big data and advancements in artificial intelligence (AI) have turned to deep learning as a superior solution for financial distress prediction due to its high accuracy (Brenes et al., 2022).

Artificial Neural Networks (ANN), also known as neural networks, are the fundamental framework of deep learning. In ANN, stimuli are received by the input layer, processed through one or more hidden layers containing activation functions, and ultimately passed to the output layer. These activation functions assign weights to the inputs and determine the output passed along the network (Sharma et al., 2020). Previous studies have highlighted the advantages of ANN over other machine learning methods in the context of financial distress prediction. For example, Marso and Merouni (2020) found that the Neural Network model outperformed Logistic Regression in predicting financial distress within Poland's manufacturing sector. Awalia and Kristanti (2023) employed ANN to predict financial distress in the banking sector and reported a high accuracy rate of 87%. Aydin et al. (2022) also concluded that ANN provided higher predictive accuracy compared to Decision Tree methods when applied to manufacturing, trade, and service sectors. Additionally, their study noted that ANN requires no assumptions, supports both linear and complex nonlinear structures, offers strong predictive power, and is user-friendly.

Despite these advantages, the performance of ANN models is highly dependent on the training process, including parameter tuning, hyperparameter selection, convergence, and optimization efficiency. Various optimization algorithms have been developed for training ANN models, such as Stochastic Gradient Descent (SGD), Root Mean Square Propagation (RMSProp), Adaptive Gradient Descent (AdaGrad), Adadelta, and Adaptive Moment Estimation (Adam) (Défossez et al., 2020; Irfan et al., 2023; Mandasari et al., 2023; Mehmood et al., 2023). Stochastic Gradient Descent (SGD) is an optimizer that updates model parameters based on the gradient derived from a random subset of the dataset rather than the entire dataset. SGD is known for its iteration speed and suitability for large datasets. However, it tends to experience oscillations (sharp fluctuations in parameter updates) and may result in unstable convergence (Haque et al., 2022). Mehmood et al. (2023) further noted that SGD requires a large number of iterations and may exhibit degraded performance on large-scale datasets. RMSProp (Root Mean Square Propagation) was developed to address the weaknesses of standard SGD, particularly concerning oscillations and convergence instability. RMSProp adaptively adjusts the learning rate for each parameter based on the moving average of squared gradients. However, its efficiency decreases in simpler optimization tasks (Qorich & Ouazzani, 2024).

AdaGrad (Adaptive Gradient Descent) is another optimizer that adapts the learning rate for each parameter based on the accumulated history of gradients. While effective in some contexts, AdaGrad can cause the learning rate to diminish excessively, potentially halting the training process prematurely (Qorich & Ouazzani, 2024). Adadelta, an extension of AdaGrad, attempts to overcome this issue by using an exponentially decaying average of previous gradients, though at the cost of increased computational complexity (Al-Khazzar et al., 2024). Adam (Adaptive Moment Estimation) has emerged as one of the most popular optimizers in deep learning, combining the advantages of SGD with momentum and RMSProp to achieve more stable and faster convergence during weight updates (Al-Khazzar et al., 2024). Given the various strengths and limitations of these optimization algorithms, their effectiveness may vary depending on the dataset or specific case under consideration. Therefore, this study applies five optimization algorithms—SGD, RMSProp, AdaGrad, Adadelta, and Adam Optimizer—to the task of financial distress prediction, aiming to identify which optimizer yields the best model performance. This study contributes to the growing body of literature on financial distress prediction by providing a comprehensive empirical comparison of five optimization algorithms—SGD, RMSProp, AdaGrad, Adadelta, and Adam—in training Artificial Neural Networks. Methodologically, the integration of Principal Component Analysis (PCA), five-fold cross-validation, and L2 regularization enhances the robustness and generalizability of the predictive model. Practically, the study provides decision-makers and financial analysts with a data-driven recommendation for selecting suitable optimization algorithms to detect financial distress early, thereby supporting risk mitigation efforts in volatile business sectors.

2 Methodology

The dataset used in this study was derived from the financial statements of companies in the property and real estate sector listed on the Indonesia Stock Exchange (IDX) for the period of 2018 to 2022. The input features consisted of 18 financial ratios, which were reduced to 5 key variables through Principal Component Analysis (PCA) for feature selection. These selected features are: Return on Assets (ROA), Earnings Per Share (EPS), Book Value Per Share (BVPS), Operating Profit Margin (OPM), and Net Profit Margin (NPM). Financial distress classification was defined based on the presence of one or more of the following criteria: (a) Negative working capital; (b) Negative operating income; or (c) Negative net income. To ensure the predictive nature of the model, financial distress was classified using a lag of t+1 years relative to the input features. The Artificial Neural Network (ANN) architecture employed in this study consists of 1 input layer with 5 neurons, hidden layers with 32 neurons in the first layer and 16 neurons in the second, and 1 output layer with 1 neuron for binary classification. The ReLU (Rectified Linear Unit) activation function was applied between the input and hidden layers (Lederer, 2021), while the Sigmoid activation function was used between the final hidden layer and the output layer (Suklabaidya & Das, 2023). To address the class imbalance between financially distressed and non-distressed companies, the L2 kernel regularization was implemented (Mai et al., 2019; Putra et al., 2023; Smiti & Soui, 2020).

Data splitting for training and testing was performed using 5-fold cross-validation, aiming to reduce the risk of overfitting and to ensure robustness in model evaluation. The optimization algorithms evaluated in this study include: Stochastic Gradient

Descent (SGD), Root Mean Square Propagation (RMSProp), Adadelta, Adagrad, and Adam Optimizer. The comparative performance of these optimizers was evaluated using the following metrics: accuracy, F1-score, precision, and recall, which were calculated using standard classification metric formulas.

Table 1. Confusion Matrix

Category		Actual class	
		Positive	Negative
Prediction Class	Positive	TP (True Positive)	FP (False Positive)
	Negative	FN (False Negative)	TN (True Negative)

$$accuracy = \frac{TP+TN}{P+N} \tag{1}$$

$$precision = \frac{TP}{TP+FP} \tag{2}$$

$$recall = \frac{TP}{TP+FN} = \frac{TP}{P} \tag{3}$$

$$F1 = \frac{2 \times precision \times recall}{precision+recall} \tag{4}$$

In addition, the Receiver Operating Characteristic (ROC) curve was also used as a performance evaluation metric. The ROC curve is a graphical representation that illustrates the trade-off between sensitivity and specificity (Kristanti & Dhaniswara, 2023). It serves as a visual representation of the performance of a binary classification model across various classification thresholds. The x-axis represents 1 - specificity, also known as the False Positive Rate (FPR), while the y-axis represents sensitivity, also referred to as the True Positive Rate (TPR).

Figure 1. ROC Curve

The area under the ROC curve is referred to as the AUC (Area Under the Curve), which can be calculated using the following formula (Kuizinienė et al., 2024):

$$AUC = \frac{1+TPR-FPR}{2} \tag{5}$$

where TPR denotes the True Positive Rate, and FPR refers to the False Positive Rate, which is defined as FPR=1−recall. A predictive model with a high AUC value is considered to have good classification performance.

3 Result and Discussion

3.1 Result

To assess the impact of optimization algorithms on the performance of the ANN model in predicting financial distress, five optimizers were compared: Adam, SGD, AdaGrad, RMSProp, and Adadelta. The comparison utilized key classification metrics: F1-score, accuracy, precision, recall, and AUC, as shown in Table 2 and Figure 2.

Table 2 The Comparison of Optimization Algorithms in ANN

Optimizer	Learning Rate	F1	Accuracy	Precision	Recall	AUC
Adam	0.010	0.616775*	0.734462*	0.625000*	0.532143	0.749462*
SGD	0.100	0.598730	0.656615	0.408991	0.710714*	0.723669
AdaGrad	0.010	0.575621	0.617846	0.280845	0.360714	0.607634
RMSProp	0.100	0.558421	0.623692	0.358281	0.582143	0.725741
Adadelta	0.001	0.436140	0.633231	0.417835	0.485714	0.633584

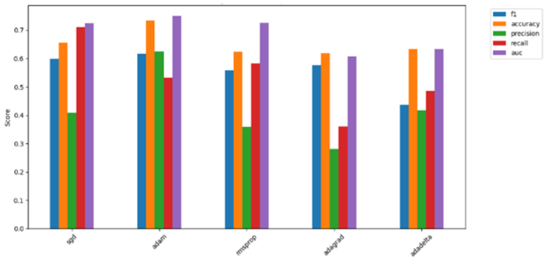

Figure 2. The Comparison of Optimization Algorithms in ANN

Table 2 and Figure 2 show that Adam Optimizer gives the highest score for F1, accuracy, precision, and AUC for training data in the ANN model. The SGD Optimizer gave the highest recall value. The confusion matrix for testing data using the best model of the ANN with the Adam optimizer is displayed in Figure 3(a). To evaluate the performance of the Artificial Neural Network (ANN) model in predicting financial distress, this study employed the Receiver Operating Characteristic (ROC) curve along with its corresponding Area Under the Curve (AUC) value. The ROC curve provides a comprehensive graphical representation of the trade-off between the True Positive Rate (TPR) and the False Positive Rate (FPR) at various classification thresholds. ROC Curve was displayed in Figure 3 (b). As illustrated in Figure 3 (b), the ROC curve for the ANN model in this study yielded an AUC value of 0.64.

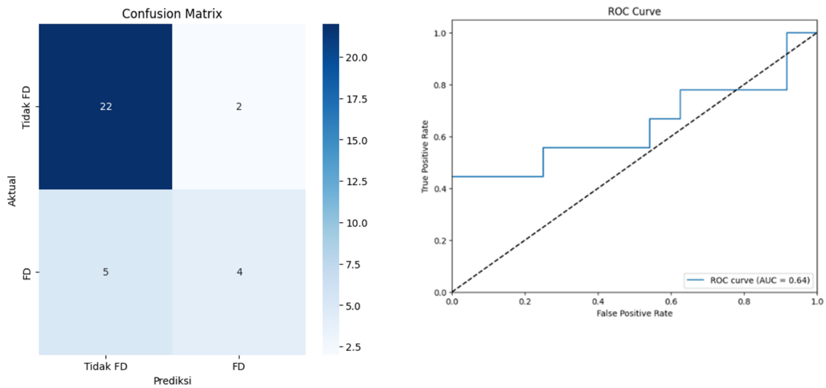

(a) (b)

Figure 3. (a) Confusion Matrix for the Testing Data; **(b)** ROC Curve for the Testing Data

3.2 Discussion

From the results above, it can be said that Adam Optimizer achieved the best overall performance in the ANN model for predicting financial distress in Indonesia's property and real estate sector. Adam produced the highest accuracy (73.45%), F1-score (0.6168), precision (0.6250), and AUC (0.7495). This superiority can be attributed to Adam's ability to combine the advantages of momentum and RMSProp, enabling it to adjust the learning rate for each parameter adaptively (Al-Khazzar et al., 2024). In complex and volatile financial data, such as that of the property sector, this adaptive capability enables the model to capture nonlinear patterns without losing convergence stability. The relatively high AUC also reflects Adam's stronger discriminative ability to distinguish between distressed and non-distressed firms, even though it remains at a moderate level (AUC = 0.7495).

The SGD Optimizer stood out in terms of recall (0.7107), meaning it detected a higher proportion of distressed firms compared to other optimizers. In the context of risk management, high recall is beneficial for early warning systems as it minimizes the risk of false negatives. However, SGD's lower precision (0.4089) indicates a higher risk of false positives, which could lead to overly conservative decision-making or excessive allocation of resources to companies that are actually healthy. This outcome aligns with the nature of SGD, which updates weights based on random mini-batches, often causing parameter oscillations when handling imbalanced datasets (Haque et al., 2022).

AdaGrad produced the weakest results among the optimizers, with very low precision (0.2808), recall (0.3607), and the lowest AUC (0.6076), showing that the model not only generates many false alarms but also fails to identify a large share of truly distressed firms, resulting in poor class separation overall. These shortcomings are critical in financial distress prediction, where both false positives and false negatives carry serious consequences. The main issue lies in AdaGrad's aggressive and irreversible learning rate decay, which causes the optimizer to stop updating parameters too early and converge at suboptimal points. While this adaptive mechanism can be Effective in sparse settings, this approach led to weak generalization and an inability to capture complex decision boundaries in this context.

Precision reflects the accuracy of positive predictions. With a value of only 0.3583, RMSProp correctly identifies about 36% of the firms it classifies as distressed. The low precision suggests a high false positive rate, meaning the model frequently predicts distress where there is none. In practice, this tendency to overpredict distress could cause unnecessary concern or misallocation of resources, as healthy firms are incorrectly flagged as high-risk. Recall measures the ability to detect true positives. At 0.5821, RMSProp captures roughly 58% of actual distressed firms. This recall score is moderate: the algorithm identifies more than half of the true distressed cases but still misses a considerable proportion. The AUC represents the model's capacity to distinguish between distressed and non-distressed firms across thresholds. An AUC of 0.7257 is considered acceptable to good, showing that RMSProp achieves a fair balance in ranking positive and negative cases. However, it is still below the performance of optimizers like Adam and SGD, which likely maintained stronger generalization and better separation between classes.

Adadelta's performance indicates notable weaknesses in handling the dataset, as reflected in its lowest F1-score (0.4361) and low AUC (0.6336), which together suggest a poor balance between precision and recall, as well as a limited ability to distinguish between distressed and non-distressed firms. Although its accuracy (63.32%) might initially appear moderate, this figure is misleading due to the dataset's class imbalance, where non-distressed firms dominate the sample. In such cases, a model can achieve seemingly acceptable accuracy by predominantly predicting the majority class while still failing to capture a meaningful proportion of distressed cases. This is precisely what occurred with Adadelta, as its bias toward the majority class undermined its ability to detect financial distress, thereby limiting its practical value for early warning systems where correctly identifying distressed firms is far more critical than simply achieving high overall accuracy.

As illustrated in Figure 3, the ROC curve for the ANN model in this study yielded an AUC value of 0.64. Looking at the AUC of the best model (Adam) at 0.7495, the model's ability to rank companies by their distress risk can be considered good but not yet highly reliable (AUC $\geq$ 0.80). Factors such as class imbalance and the limited number of input variables after PCA reduction are likely contributors. Previous studies by Aydin et al. (2022) and Kristanti and Dhaniswara (2023) have demonstrated that incorporating more relevant variables and applying data balancing techniques, such as oversampling or SMOTE, can significantly enhance model performance. These findings have practical implications. First, Adam should be the default choice when training ANN models for financial distress prediction in volatile sectors such as property and real estate. Previous research has also yielded similar results, where the Adam optimizer demonstrated the best performance among others (Irfan et al., 2023; Mandasari et al., 2023; Suklabaidya & Das, 2023). Second, if the primary objective is to minimize missed detections of distressed firms, SGD can be considered, though with the trade-off of a higher false alarm rate. Third, optimization methods such as AdaGrad and Adadelta should be avoided unless extensive hyperparameter tuning is conducted.

Overall, the performance differences across optimizers underscore the importance of selecting an optimization algorithm that aligns with the dataset characteristics and prediction goals. These results also open opportunities for further experimentation with hybrid optimizers or learning rate scheduling to enhance both predictive accuracy and model discrimination capability.

4 Conclusion

From the results and discussion, it can be concluded that Adam is the most suitable optimizer for financial distress prediction using ANN in this study due to its consistent performance across all metrics and superior AUC. SGD can be considered a good alternative when high recall is prioritized, such as in conservative early warning systems. AdaGrad and Adadelta are less effective in this context and may require significant tuning or hybridization to perform competitively. RMSProp shows moderate potential but does not outperform Adam or SGD. Despite the promising results, this study is subject to several limitations. First, the dataset is restricted to publicly listed companies in the Indonesian property and real estate sector, which limits the generalizability of the findings to other industries or regions. Second, the ANN architecture was fixed and did not explore more complex configurations or alternative deep learning models, such as LSTM or CNN, which may have yielded improved performance. Future research is encouraged to expand the scope of industries, adopt more advanced neural network architectures, incorporate ensemble or hybrid modeling approaches, and conduct external validation to enhance model reliability and practical applicability.

Acknowledgment

The authors would like to express sincere gratitude to the Center for Research and Community Service (P3M), Politeknik Negeri Bali, for providing research funding support through the internal research grant scheme. This research would not have been possible without the financial assistance and institutional support from P3M Politeknik Negeri Bali.

References

Al-Khazzar, A. M., Altaweel, Z., & Hussain, J. S. (2024). Using deep neural networks in classifying electromyography signals for hand gestures. *IAES International Journal of Artificial Intelligence*, *13*(1), 217–227. https://doi.org/10.11591/ijai.v13.i1.pp217-227.

Awalia, R., & Kristanti, F. T. (2023). Prediksi financial distress menggunakan artificial neural network (ANN) pada perusahaan perbankan yang terdaftar di Bursa Efek Indonesia (BEI) Periode 2017-2021. *Jurnal Kajian Pendidikan Ekonomi dan Ilmu Ekonomi*, *7*(2), 2549–2284.

Aydin, N., Sahin, N., Deveci, M., & Pamucar, D. (2022). Prediction of financial distress of companies with artificial neural networks and decision tree models. *Machine Learning with Applications*, *10*(November), 100432. https://doi.org/10.1016/j.mlwa.2022.100432.

Brenes, R. F., Johannssen, A., & Chukhrova, N. (2022). An intelligent bankruptcy prediction model using a multilayer perceptron. *Intelligent Systems with Applications*, *16*(September). https://doi.org/10.1016/j.iswa.2022.200136.

Défossez, A., Bottou, L., Bach, F., & Usunier, N. (2020). *A simple convergence proof of Adam and AdaGrad*. http://arxiv.org/abs/2003.02395.

Haque, R., Ho, S. B., Chai, I., & Abdullah, A. (2022). Parameter and hyperparameter optimisation of a deep neural network model for personalised predictions of asthma. *Journal of Advances in Information Technology*, *13*(5), 512–517. https://doi.org/10.12720/jait.13.5.512-517.

Irfan, D., Gunawan, T. S., & Wanayumini, W. (2023). Comparison of SGD, RMSprop, and Adam optimization in animal classification using CNNs. *International Conference on Information Science and Technology Innovation (ICoSTEC)*, *2*(1), 45–51. https://doi.org/10.35842/icostec.v2i1.35.

Kristanti, F. T., & Dhaniswara, V. (2023). The accuracy of artificial neural networks and logit models in predicting the companies' financial distress. *Journal of Technology Management and Innovation*, *18*(3), 42–50. https://doi.org/10.4067/s0718-27242023000300042.

Kuizinienė, D., Savickas, P., Kunickaitė, R., Juozaitienė, R., Damaševičius, R., Maskeliūnas, R., & Krilavičius, T. (2024). A comparative study of feature selection and feature extraction methods for financial distress identification. *PeerJ Computer Science*, *10*. https://doi.org/10.7717/PEERJ-CS.1956.

Lederer, J. (2021). *Activation functions in artificial neural networks: A systematic overview*. 1–42. http://arxiv.org/abs/2101.09957.

Letkovsky, S., Jencova, S., Vasanicova, P., Gavura, S., & Bacik, R. (2023). Predicting bankruptcy using artificial intelligence: The case of the engineering industry. *Economics and Sociology*, *16*(4), 178–190. https://doi.org/10.14254/2071-789X.2023/16-4/8.

Liashenko, O., Kravets, T., & Kostovetskyi, Y. (2023). Machine learning and data balancing methods for bankruptcy prediction. *Ekonomika*, *102*(2), 28–46. https://doi.org/10.15388/Ekon.2023.102.2.2.

Mai, F., Tian, S., Lee, C., & Ma, L. (2019). Deep learning models for bankruptcy prediction using textual disclosures. *European Journal of Operational Research*, *274*(2), 743–758. https://doi.org/10.1016/j.ejor.2018.10.024.

Mandasari, S., Irfan, D., Wanayumini, W., & Rosnelly, R. (2023). Comparison of SGD, Adadelta, Adam optimization in gender classification using CNN. *JURTEKSI (Jurnal Teknologi dan Sistem Informasi)*, *9*(3), 345–354. https://doi.org/10.33330/jurteksi.v9i3.2067.

Marso, S., & El Merouani, M. (2020). Predicting financial distress using a hybrid feedforward neural network with a cuckoo search algorithm. *Procedia Computer Science*, *170*(2019), 1134–1140. https://doi.org/10.1016/j.procs.2020.03.054.

Mehmood, F., Ahmad, S., & Whangbo, T. K. (2023). An efficient optimization technique for training deep neural networks. *Mathematics*, *11*(6). https://doi.org/10.3390/math11061360.

Peralungal, A., & Natchimuthu, N. (2024). Machine learning in financial distress: A scoping review. *International Research Journal of Multidisciplinary Scope*, *5*(3), 457–474. https://doi.org/10.47857/irjms.2024.v05i03.0779.

Putra, P. L. W. E., Naufal, M., & Hidayat, E. Y. (2023). A comparative study of the MobileNet architecture optimizer for crowd prediction. *Jurnal Informatika: Jurnal Pengembangan IT*, *8*(3), 241–247. https://doi.org/10.30591/jpit.v8i3.5703.

Qorich, M., & Ouazzani, R. . (2024). Optimizer algorithms and convolutional neural networks for text classification. *IAES International Journal of Artificial Intelligence, 13*(1), 451 458. https://doi.org/10.11591/ijai.v13.i1.pp451-458.

Sharma, S., Sharma, S., & Athaiya, A. (2020). Activation functions in neural networks. *International Journal of Engineering Applied Sciences and Technology*, *04*(12), 310–316. https://doi.org/10.33564/ijeast.2020.v04i12.054.

Smiti, S., & Soui, M. (2020). Bankruptcy prediction using a deep learning approach based on borderline SMOTE. *Information Systems Frontiers*, *22*(5), 1067–1083. https://doi.org/10.1007/s10796-020-10031-6.

Suklabaidya, S., & Das, I. (2023). Performance comparison of multiple ANN optimizers on IoT-enabled sensor fire dataset. *International Arab Journal of Information Technology*, *20*(5), 767–772. https://doi.org/10.34028/iajit/20/5/9.