SCHOLARLY DEBUT
by Scholarly Review

Summer 2026

Investigating to what extent learning hours affect students' GPA through a comparison of linear regression and neural-network machine learning models

By Jinmo Wu

AUTHOR BIO

Jinmo Wu is a Year 13 student at Macleans College in Auckland. Passionate about mathematics and physics, he is interested in using data analysis and quantitative methods to explore relationships between real-world variables. His recent work focuses on statistical interpretation of academic performance data and the application of mathematical modeling in educational contexts.

ABSTRACT

This paper investigates that to what extent students' study time can be used to predict their Grade Point Average (GPA) through the application of machine learning algorithms. As the relevance of predicting academic performance continues to rise in the field of education, the use of data analysis offers a practical means of exploring the relationship between learning behavior and academic performance. A publicly accessible dataset with study time and corresponding GPA data was employed for analysis. Two models were employed and compared for predictive purposes: linear regression and a neural network model. Linear regression was employed as a control model to establish the linear relationship between study time and GPA, while the neural network model was employed to explore the possibility of non-linear relationships within the data. Standardization of the data was conducted before model training, and the performance of the models was assessed using Mean Absolute Error (MAE) and Mean Relative Error (MRE), as well as graphical assessment using scatter plots. The results show that although linear regression establishes a straightforward and easy-to-interpret relationship, the neural network model performs better in terms of predictive accuracy by exploring more complex relationships.

Keywords: *machine learning, linear regression, neural networks, GPA Prediction, Educational Data Mining*

SCHOLARLY DEBUT
by Scholarly Review

Summer 2026

INTRODUCTION

The number of applications submitted to institutions of higher learning has been increasing each year, thereby increasing the workload for the admissions team. For the 2023-2024 academic year, the Common App (an organisation that helps students apply to institutions of higher learning) reported a 11% increase in the number of submitted applications compared to the previous year. In the UK, the UCAS system (an organisation in the UK that reviews the applications submitted by students) reported that the number of entrants aged 18 years and younger rose from 594,940 to 600,660 by January 2025. On average, each member of the admissions team currently reviews 400 applications per admission cycle.

On the other hand, the importance of university planning and admission results for students and parents continues to grow. It's a fact that the majority of present-day students are in need of an early forecast of their possible grades in subjects to decide whether they can apply or need to prepare further. Because of the need for such a tool in an increasing number of people, it is an increasingly valuable tool in contemporary times for the development of a program predicting the academic performance of students using their study time. Therefore, this paper presents two different types of machine learning models to predict the results of students with their study hours each day to help relieve the burden.

METHODS

This study applies two machine-learning approaches—linear regression and a neural network—to predict students' GPA based on their study hours. The aim is to evaluate how well each model performs and whether more advanced AI techniques offer a meaningful improvement.

Dataset

The data for this project includes the study hours and GPAs for students. Study hours are considered the independent variable, and GPAs are considered the dependent variables that need prediction. Prior to modeling, data standardization with the StandardScaler function takes place. The StandardScaler function allows the data to be modified to have a resultant mean and standard deviation of 0 and 1, respectively. This step can be referred to as standardization. This method works by subtracting the mean from data points and dividing by the standard deviation, standard deviation for each variable that needs to be standardized. Data standardization helps such that data points exist on equal scales.

The source of the data used in this research work is a publicly available set of data published on the Kaggle platform. The name of the set of data is GPA Study Hours. This set of data provides the study hours per week of undergraduate students along with their Grade Point Average. The set of data comprises 193 observations with two variables: study_hours and GPA. The study_hours variable denotes the number of study hours per week the students perform, while the gpa variable denotes the performance level of the students.

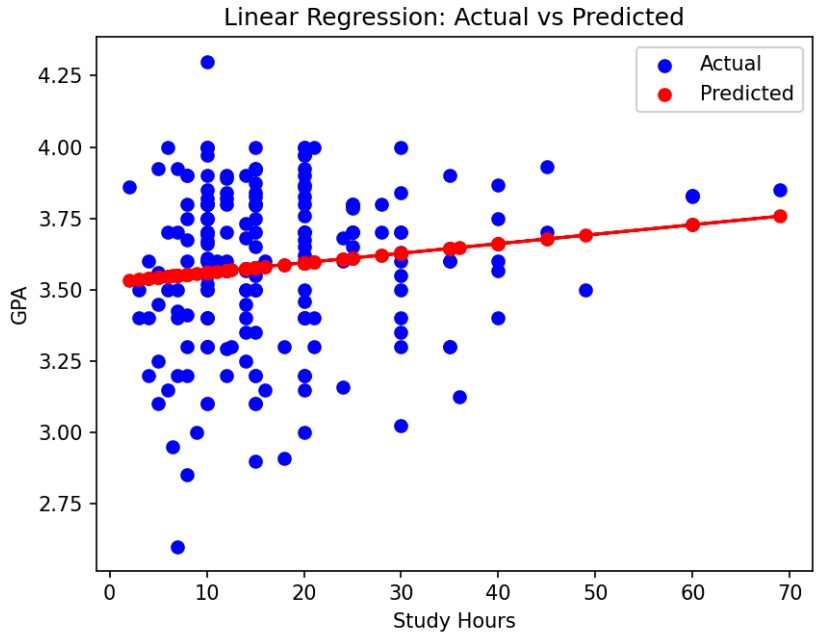
This particular dataset has been selected due to it being a practical and simple example of how the time spent studying can be correlated with results. Even though there are only two variables and no additional information regarding participants, it can still be considered adequate for representing the

capability for prediction for machine-learning algorithms. Also, being free and readily available, it makes it easier for others to test and reproduce the analysis. All information has been handled and processed following ethics guidelines and without using any personal information.

Machine Learning Models

The first model implemented is a linear regression model. Linear regression assumes a direct and proportional relationship between study hours and GPA. The model learns a single weight and bias that define the best-fitting straight line through the data. After training, the predictions are transformed back to the original GPA scale through inverse scaling. Model performance is assessed using Mean Absolute Error (MAE), with the formula $MAE = \frac{1}{n}\sum_{i=1}^{n}\left|y - \hat{y}_i\right|$, which measures how far the predictions deviate from the real GPA values, and Mean Relative Error (MRE), with the formula $MAE = \frac{1}{n}\sum_{i=1}^{n}\frac{\left|y_i - \hat{y}_i\right|}{y_i}$, which is the percentage of the MAE to the true value. Scatter plots, Figure 1 and Figure 2, are also created to visually illustrate the difference between actual and predicted results. Although linear regression is simple and easy to interpret, it may not capture more complex patterns that exist within real-world educational data.

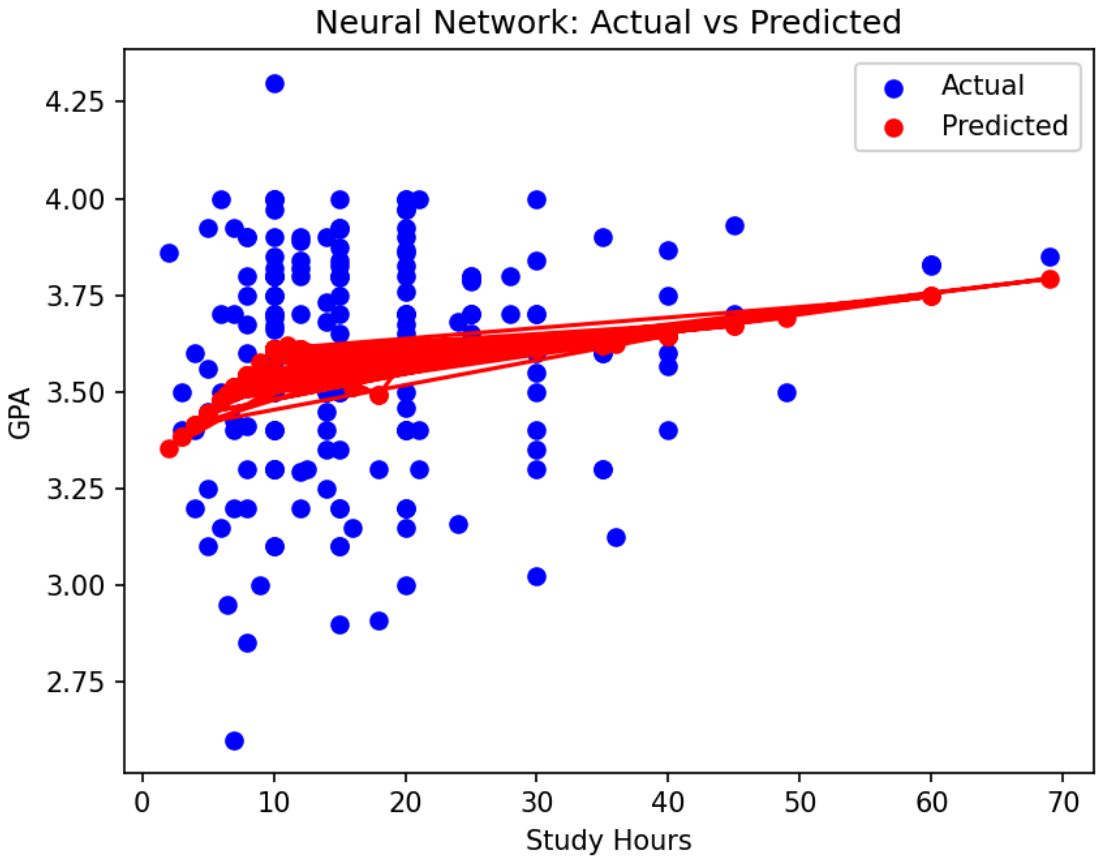
Figure 1
Linear regression scatter plot



Note. Created by author, 2026.

The second model is a neural network. Neural networks are designed to learn non-linear relationships by passing information through multiple layers of interconnected neurons. In this project, the neural network includes three hidden layers with 32, 16, and 8 neurons, all using the ReLU activation function, which works when the input x is positive, ReLU returns x, and when the input x is zero or negative, ReLU returns 0. The output layer uses a linear activation function to support continuous predictions. The model is trained for 300 epochs using the Adam optimiser and mean squared error (MSE) as the loss function. Both training loss and validation loss are recorded to monitor the model's learning process and identify potential overfitting. After training, the predictions are converted back to the original scale, and the model's accuracy is evaluated using the same MAE and MRE metrics used for linear regression.

Figure 2
Neural network scatter plot



Note. Created by author, 2026.

A comparison is then made between the two models. Linear regression offers simplicity and interpretability but is limited to modelling straight-line relationships. In contrast, the neural network is capable of learning more complex and non-linear patterns, which may lead to more accurate GPA predictions. By examining the error values and visual results from both models, this study determines which approach provides a more reliable method for supporting academic performance prediction.

RESULTS

Performance of the Linear Regression Model

First, a linear regression formula is utilized to predict the GPA of the students given their study hours. Analyzed data indicate that as study hours ascend, there is a corresponding increase in predicted GPAs, indicating a linear relationship. This verifies that there is a positive correlation between study hours and academic achievement.

However, if we observe the scatter plot (Figure 1) of the actual GPAs and the predicted values in the linear regression model, there appear to be some variations in the data. This is mainly because some students who have a similar number of study hours end up with different GPAs. These variations cannot be included in the linear regression model since the rate of change remains the same.

Performance of the model was also determined using MAE and MRE. Both measures indicate that although the average error in predictions is fairly stable, the model fails to offer very accurate predictions. However, linear regression is still a workable baseline model which is inflexible.

Performance of the Neural Network Model

The neural network model portrays enhanced performance relative to a linear regression model. This is evident from the fact that the predictions for GPA lie closer to real GPA data, and the scatter plot (Figure 2) portrays a set of points that are clumped around the actual values. This indicates better mapping performed by a neural network for hours studied and GPA.

In the training phase, both the training loss and the validation loss decrease gradually for 300 epochs. This shows that the model extracts relevant patterns from the data without any strong overfitting. Multiple hidden layers in the network help capture non-linear relationships that cannot otherwise be presented in a linear fashion.

In relation to the quantitative analysis for the results, the MAE and MRE for the neural network are lower compared to the linear regression model. This implies that the forecasted values by the neural network are closer to the GPA compared to the linear regression model. It is clear that in cases where the relation between the variables is not strictly linear, the following models are better suited for academic performance prediction.

Comparison of Model Performance

A comparison of both models reveals distinct variations in their accuracy of prediction. In linear regression, there is simplicity associated with the model. However, linear regression faces problems when there is no strict linear relation between the data points. The neural network is more adaptable regarding variations within the data and is associated with better predictions.

The above results indicate that there are some non-linear relations between study hours and GPA. Hence, a learning model with the capability of learning from non-linear relationships, such as a neural network model, would be a better choice for tasks such as the above.

DISCUSSION

Implications for Students and Academic Planning

On the basis of the results of this study, it has been seen that machine learning algorithms have a great use in providing support to students to decide upon their educational path. Using a prediction of GPA on the basis of study hours, a student is able to have a better insight into a possible performance in studies so far, taking into consideration whether more work is needed.

These predictive models may also alleviate stress associated with having realistic expectations. Students will not have to wonder or guess what happens in their academic journey because there will be predictive models based on data insights.

Benefits and Limitations of Using AI in Education

An important advantage of applying AI technology to education is that it is capable of processing data effectively and estimating some hidden patterns, which might not be possible through human analysis. Perhaps one of the most important advantages of using neural networks is that they generate more precise predictions than previous models. Nonetheless, this study will also present certain limitations. In this study, only one variable is used as input in the study hours feature. Yet, there are many more factors that contribute to academic outcomes. Therefore, predictions are not entirely correct indications of academic achievement.

Fairness and bias should be taken into consideration while using AI. The dataset used in the study fails to include details about the students' backgrounds, learning environments, or personal situations. The created model could not generalize to all students. In the real-life context at the stage of admission, transparency, fairness, and ethical use of the collected data should be ensured. AI should not replace human decision-making. It should be used to support decision-making.

CONCLUSION

Summary of Findings and Overall Conclusions

In this paper, the application of machine learning algorithms for the prediction of the GPA of students using the number of study hours was investigated. In this case, two algorithms are applied: linear

regression and the neural network approach. It was observed that although the linear regression method is straightforward in the sense that it establishes a baseline model for prediction, the accuracy attained by the neural network approach is better as it does not require the prediction to be linear.

For improved predictions, it is advisable for future studies to include other predictor variables like scores, difficulty levels, or participation in additional activities. Another area for improvement is to have access to a broader range of data for enhanced reliability.

In conclusion, it appears that machine learning has the potential to aid students and institutions during the college admission process. It should be noted that machine learning can be an effective predictive tool if it is designed and implemented properly.

REFERENCES

1. Brownlee, J. (2019). *Machine learning mastery with Python*. Machine Learning Mastery. https://machinelearningmastery.com/
2. Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer. https://www.springer.com/gp/book/9780387310732
3. Chollet, F. (2018). *Deep learning with Python*. Manning Publications. https://www.manning.com/books/deep-learning-with-python
4. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press. https://www.deeplearningbook.org/
5. Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer. https://doi.org/10.1007/978-0-387-84858-7
6. James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). *An introduction to statistical learning: With applications in R* (2nd ed.). Springer. https://www.statlearning.com/
7. Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *International Conference on Learning Representations*. https://arxiv.org/abs/1412.6980
8. Ng, A. (2018). *Machine learning yearning: Technical strategy for AI engineers*. Deeplearning.ai. https://www.deeplearning.ai
9. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830. https://jmlr.org/papers/v12/pedregosa11a.html
10. Romero, C., & Ventura, S. (2010). Educational data mining: A review of the state of the art. *IEEE Transactions on Systems, Man, and Cybernetics, Part C, 40*(6), 601–618. https://doi.org/10.1109/TSMCC.2010.2053532
11. Scikit-learn developers. (n.d.). *StandardScaler*. Scikit-learn documentation. https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.StandardScaler.html
12. Shivam, S. (n.d.). *GPA v/s study hours dataset* [Data set]. Kaggle. https://www.kaggle.com/code/shiivvvaam/gpa-v-s-study-hours/input
13. Willmott, C. J., & Matsuura, K. (2005). Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance. *Climate Research, 30*(1), 79–82. https://doi.org/10.3354/cr030079

14. Zhang, Y., & Yang, Q. (2018). A survey on multi-task learning in deep neural networks. *IEEE Transactions on Knowledge and Data Engineering, 34*(12), 5589–5608. https://doi.org/10.1109/TKDE.2021.3054719