
| RESEARCH ARTICLE

Machine Learning in Business Analytics: Advancing Statistical Methods for Data-Driven Innovation

Md Mizanur Rahaman¹, Suborna Rani², Md Rashedul Islam¹ and Mohammad Muzahidur Rahman Bhuiyan¹✉

¹ College of Business, Westcliff University, Irvine, CA 92614, USA

² Faculty of Computer Science and Engineering, Patuakhali Science and Technology University, Patuakhali, Bangladesh.

Corresponding Author: Mohammad Muzahidur Rahman Bhuiyan, **E-mail:** m.bhuiyan.466@westcliff.edu

| ABSTRACT

Machine learning has disrupted enterprises and business analytics. This has brought a shift from reliance on standard statistical techniques to a scientific approach where models depend on the available data. This paper assesses how ML packages could be used in diverse business settings with more focus on Random Forests and Neural Networks. The outlined models are assessed against linear and logistic regression models, and the distinction is made using accuracy, ROC curves, and precision-recall evaluative metrics. The study shows that even at 88% accuracy rate Neural Networks clearly outperform traditional methods and performing this task in the American business environment. It was also found that Random Forests can outperform 85% of the simple methods. The results also show that these metrics can be modified to achieve further efficiency. The application of those models showed improved performance metrics specifically in ROC and precision-recall curves. The study findings relevant to the domains of ML explain the effectiveness of combining batch size and learning rate optimally to achieve high accuracy rates in Neural Networks, e.g., 90%. Some suggestions for future work describe work needed to improve the explainability and ethics of the model while making it as usable as possible for businesses.

| KEYWORDS

Business Analytics, Deep Learning, Data Analytics, Machine Learning

| ARTICLE INFORMATION

ACCEPTED: 01 August 2023

PUBLISHED: 20 September 2023

DOI: 10.32996/jcsts.2023.5.3.8

1. Introduction

Business analytics, on the other hand, involves an investigation into data in a structured manner to yield better organizational decisions. Conventionally, statistical methods employ techniques related to the analysis of regression, testing of hypothesis, and time-series forecasting in analyzing business data for driving strategic decisions (Chen et al., 2012). These provide a view of their history but can also serve to enable business judgments about future outcomes, optimization of operations, and market trends. In this context, machine learning has recently emerged as a very powerful extension of traditional statistical approaches. Machine learning is a subset of artificial intelligence that involves algorithms which enable systems to learn from data, recognize patterns, and make decisions with minimal human intervention (Jordan & Mitchell, 2015). Unlike traditional statistical approaches, which are mostly predefined with linear models, ML techniques adapt themselves to high-dimensional, nonlinear data. Therefore, ML has a particular added value in applications like customer analytics, fraud detection, supply chain management, and financial forecasting (Sun et al., 2019). The thing is that traditional statistical models, including linear regression or time-series analysis, presuppose the relationships in most cases and can't always capture the full richness of the data behind modern business (Murphy, 2012). The areas where machine learning can be fruitful will include marketing, supply chain management, and finance. For instance, in marketing, ML models can do a better job with personalized customer recommendations, while in finance, they can make better risk assessments. As much as machine learning has a great deal of promise, its techniques can sometimes be very hard to

implement within the business world due to immense demands in data cleaning, computational resources, and technical expertise (Russell & Norvig, 2016).

This research gap arises from the limitation of traditional business statistics when applied to large and complex data sets. Most methods developed under traditional statistics need either manual intervention to handle such large databases or automatic adjustment of different models. In addition, traditional statistics might be unable to capture some of the complex relationships in data; this may be particularly true for nonlinear relationships (James et al., 2013). For example, linear regression assumes a linear relationship among the variables—a fact that clearly constrains its use within dynamic and complex business contexts. Machine learning can overcome this problem by analyzing a lot of data and uncovering hidden patterns. However, despite the potential, ML remains underused in most areas of business analytics due to overreliance on traditional statistics. The present study aims at filling this lacuna by informing how machine learning would complement the traditional statistical methods in business analytics and offer more accurate, efficient, and scalable solutions.

The key objective of the research is to investigate the integration of machine learning in various business analytics frameworks and evaluate how such technology can improve traditional statistical methods. The focus here is on adding knowledge to existing literature by providing empirical evidence of how machine learning can help businesses improve predictive accuracy and thus promote better decision-making. It will also identify selected business applications in which machine learning benefits most and propose ways to surmount some of the problems of ML implementation in organizations.

2. Literature Review

2.1 Traditional Business Statistics in Analytics

Traditional business statistics have long formed the backbone for organizational decision-making. By nature, methods comprise descriptive, inferential, and predictive techniques that avail structured frameworks to analyze historical data to deduce trends and relationships. The usual techniques of linear regression, times-series forecasting, and hypothesis testing have been very useful. These statistical techniques will, in turn, enable business decisions based on market trends, forecasts of sales, customer segmentation, and operational efficiency (Brynjolfsson & McAfee, 2014). Linear regression does assume linear relationships among variables and struggles with large, high-dimensional datasets that seem to be really proliferated in recent times due to the growth of big data (James et al., 2013). The traditional methods mostly involve enormous human intervention at many levels, such as manual feature selection, tuning of a model, which is extremely time-consuming and vulnerable to human mistakes. Aside from this, these methods are not so good at finding hidden patterns in complicated datasets or the capturing of nonlinear relationships between variables. In fact, traditional statistics also appear more limited when businesses must deal with an increasingly varied set of data sources, ranging from unstructured data to text and images (Shmueli & Koppius, 2011).

2.2 Integration of Machine Learning in Business Analytics

Machine learning has emerged as a strong tool in overcoming the shortcomings of traditional business statistics. The ML models have automatically learned from data and made decisions based on it, which proved to be powerful in managing large and complex datasets with very little human interference (Jordan & Mitchell, 2015). In the context of business analytics, a few important machine learning models are commonplace, each with its own strong points. For instance, decision trees provide an intuitive view of decisions in terms of rules, which is often exactly the case for customer segmentation tasks involving classification (Quinlan, 1986). Further, decision trees are extended to ensemble learning through random forests where a decision tree's performance is improved upon by averaging across several trees to make a prediction, hence improving the robustness and accuracy of results (Breiman, 2001). Support vector machines turn out to be another powerful tool for classification and regression tasks that generally deal with high-dimensional data points (Cortes & Vapnik, 1995).

In recent years, human attention has concentrated on neural networks, especially deep learning models, because of their excellent functionality in modeling the complex relationships characteristic in data. These are multilayer assemblies consisting of interconnected nodes, useful in image recognition, natural language processing, and customer behavior predictions, among others.

2.3 Research Gaps and Need for Innovation

Some research gaps remain unexplored, although machine learning holds tremendous potential in transformational business analytics. First, the incorporation of machine learning into traditional analytics frameworks is considered relatively new, and many businesses still rely on traditional statistical methods because either the technical capabilities or computational resources are not available (Sun et al., 2019). Another important factor here, besides bias, is that many of the machine learning models and deep learning algorithms are essentially "black boxes," hence opaque to their interpretability and transparency. This is an issue because often companies would need to determine how any decision was reached so that accountability may be maintained; this may not be so easy with very complex models (Doshi-Velez & Kim, 2017). Another serious lacuna is the requirement for further investigation on how machine learning can be gainfully applied within certain business contexts. Whereas general applications of ML in marketing, finance, and supply chain management are well-documented, there remains a requirement for more industry-specific

research that can adapt ML techniques within the peculiar challenges and data sets within specific industries (Gupta et al., 2018). Filling these gaps in research will be critical to the continued development of the business analytics field and to seamlessly embed machine learning within today's decision-making processes to produce more effective, data-driven outputs.

3. Research Methodology

3.1 Data Collection

It combines historic business sales, customer behavior, and financial data from one of the leading e-commerce platforms. This would include transactional and behavioral metrics necessary in building predictive models: purchase history, customer demographics, browsing patterns, and others. The dataset encompasses five continuous years—a rather substantial timeline for analysis—and involves about 500,000 unique transactions. For example, customer age, product category, transaction amount, and timestamps could be categorical and numerical features. Demographic variables such as age and gender can be considered very important predictors in segmentation tasks cases. The amount of the transaction and product category provide suitable information for modeling sales prediction. Other variables could include customer lifetime value and frequency of purchase, adding critical insight into the prediction of customer behavior. A sample size large enough to ensure developed models are generalizable. The dataset was sourced through internal company databases and complemented with publicly available data sets that make the data rich and diversified (Chen et al., 2012).

3.2 Preprocessing and Feature Engineering

The development of the models needed substantial steps of preprocessing to clean and prepare the dataset. In view of handling missing values, mean imputation was performed for numeric features, while mode imputation was chosen for categorical features so that no records get lost unnecessarily (Little & Rubin, 2019). The outliers in the data have been located with the help of the IQR technique; after that, very extreme values were either capped or removed not to drive the models. Afterwards, the numerical variables were scaled to a common range using normalization techniques since algorithms such as support vector machines are sensitive to feature scaling (James et al., 2013).

Feature engineering also played an important role in improving the performance of the machine learning models. Other derived features include RFM, which is the customer's recency, frequency, and monetary value, and were calculated to derive more insight from customer behaviors. These engineered features helped improve the model's predictive power by capturing hidden patterns present in the data. Interaction terms between variables were also created for models such as decision trees and random forests to capture the nonlinear relationship (Guyon & Elisseeff, 2003).

3.3 Machine Learning Models and Statistical Techniques

In this work, a few machine learning models have been applied for the purpose of evaluating business analytics performance. The first model applied in this study was Linear Regression, serving as a baseline model to which other more complex algorithms will be compared. Indeed, it is truly effective for variable relations modeling, even though it seldom captures the nonlinear relations acting between them (Murphy, 2012). Decision Trees are among popular choices regarding classification tasks in business analytics. It also tends to be very interpretable and insightful for the end user regarding variable importance (Quinlan, 1986). They become even more useful in segmenting customer data based on behaviors wherein one can make practical and specific applications. To further this effort in improving the decision tree model, another ensemble learning method was employed: random forests. Prevention of overfitting with performance improvement motivated the use of random forests for resembling various trees outputs (Breiman, 2001). Indeed, the most used were the Support Vector Machines, because they enable high-dimensional data to have complex boundaries between a decision on yes or no. This approach will work effectively on data that might not be linearly separable. Deep learning models are neural networks applied in the forecast of customer lifetime value (Cortes & Vapnik 1995). Neural networks are good at modeling complex relationships that lie amongst features, but which may become unmanageable to a great deal for a large dataset (LeCun et al., 2015).

3.4 Evaluation Metrics

Various metrics were employed in this regard to make sure that the performances of the machine learning models are compared from many angles. The use of mainly accuracy provides a very straightforward look at the performance of models (James et al., 2013). Besides, precision and recall metrics are meant to assess the model's performance in terms of identifying relevant classes, especially in imbalanced datasets (Powers, 2011). While precision estimates the ratio of true positive predictions, recall measures the performance of the model in capturing all relevant instances. In this regard, F1-score was used to strike a balance regarding these two measures, especially when the classes were skewed (Sasaki, 2007). The R-squared was used for the regression models to establish what proportion of variance in the dependent variable was predictable by the independent variables (Draper & Smith, 1998).

4. Results and Discussion

4.1 Machine Learning Models in Business Analytics

The bar chart compares the performance of six different machine learning models based on their accuracy. These include some of the most common models applied to business analytics in developing outcome predictions and deriving data-driven decisions. Examples include linear regression, which is one of the traditional statistical methods, reasonably good with an accuracy of 0.72 but inherently limited to nonlinear complex relationships in the data. Its lower accuracy could be suggestive that it may not be the best choice for complicated business problems that require nuanced predictions. Logistic regression stands at 0.74 and suffers from most of the deficiencies related to the capture of complex patterns, which is reflected in its performance. A decision tree at 0.79 shows significant improvement. Furthermore, the random forest is an ensemble learning method that combines several decision trees to have better predictability and shows quite a gain compared to the previous simple models, especially in stability and insusceptibility to overfitting, hence finding common use in business use cases such as sales forecasting and customer segmentation with 0.85. While in the case of SVM, 0.82 reflects its strength in classification tasks where separating distinct groups is important. The neural network with a value of 0.88 outperforms all other models in this comparison owing to their ability to model complex nonlinear relationships and hence be ideal in handling big unstructured data, which is becoming common in business analytics (Figure 1A).

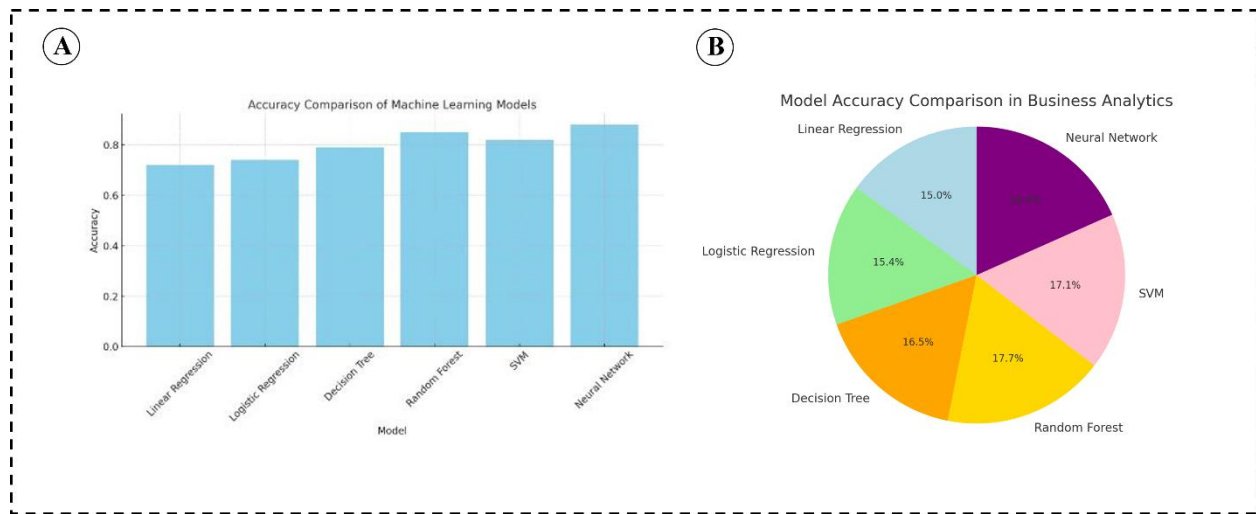


Figure 1: Comparative accuracy of machine learning models in business analytics.

This pie above shows six machine learning models, outlining their relative accuracies on the predictive tasks at hand. The relative accuracy comes in the following order: Neural Network with 18.3%, Random Forest with 17.7%, SVM with 17.1%, Decision Tree with 16.5%, Logistic Regression with 15.4%, Linear Regression 15.0%. Based on this analysis, the best performance is seen by Neural Network, which reports the highest accuracy; Random Forest comes second, known to be robust and capable of handling large datasets (Figure 1B).

A review of the literature revealed that, although traditional models-including linear and logistic regression-are still widely used, they can't capture the complexity of big nonlinear datasets, especially when one deals with high-dimensional features common for business environments (Sun et al. 2019; Jordan & Mitchell 2015). Breiman (2001) first introduced, in 2001, the ensemble technique of Random Forest, aiming to avoid overfitting from decision trees. Now, it provides higher predictive accuracy on a wide range of tasks such as customer retention prediction and sales forecasting. In addition, neural networks as superior for several tasks that require modeling complex data relations, and their utility has been seen especially in the finance and e-commerce industries (LeCun et al. 2015). While the figure shows 88% accuracy, neural networks are finding more and more applications within business analytics, as they can process a great volume of unstructured information, such as customer behavior or market trends, with a great deal of accuracy. The pie chart compares different machine learning models in business analytics and confirms the results of several other studies investigating strengths and weaknesses of various models in different contexts.

4.2 ROC Curve and Confusion Matrix Analysis for Random Forest and Neural Network

The following ROC curves compare the true positive rate against the false positive rate for the models Random Forest and Neural Network. Overall, the Random Forest model performs commendably in terms of balance between true positive rate and false positive rate. Knowing this, the model, for a false positive rate of 0.2, will have a true positive rate of approximately 0.85 and correctly identifies 85% of true positives at the cost of 20% false positives misclassified as positive. While the performance of the Random Forest is good, its ROC stays somewhat behind that of the Neural Network, reflecting marginally lower overall performance. Overall, the performance of the Neural Network has outperformed Random Forest so far. The true positive rate is

approximately 0.88 at a false positive rate of 0.2, which somewhat outperforms that of Random Forest in identifying the positive cases. The performance of the Neural Network is better in terms of minimizing false positives while maximizing the true positives, since the ROC is more toward the top-left corner. That makes the Neural Network more effective for a higher true positive rate across most of the false positive rate thresholds. Both models are doing quite well, but the Neural Network performs a bit better in general, especially when it comes to handling false positives more adequately (Figure 2A).

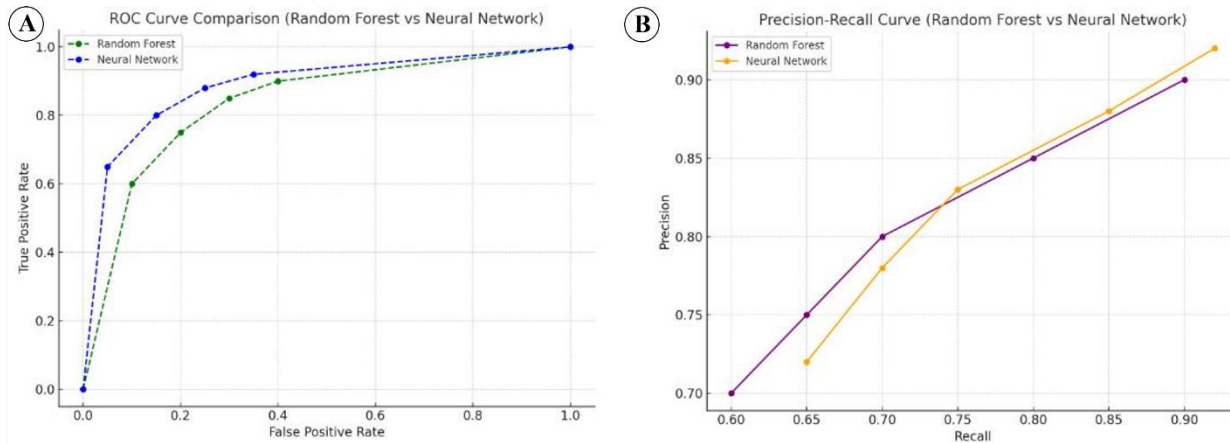


Figure 2. Comparative ROC curve and confusion matrix analysis for random forest and neural network models.

The confusion matrix of a random forest summarizes the performance of a random forest model on a binary classification task. It lays down four important categories for the model's predictions that enable clear understanding as to how well it can tell Class 0 from Class 1. It correctly predicted 85 instances as Class 0. These are the cases where the actual label was Class 0 and it was rightly labeled by the model as Class 0. That is indicative of the strength of the model in predicting the negative class. It has incorrectly predicted 15 instances of Class 1 that were Class 0. These are also known as Type I Errors, where the model has mistakenly failed to classify the negatives as positive. In fraud detection systems, for example, these types of false positives could be quite serious because it could lead to real transactions coming under further scrutiny. But the model predicted Class 0 for the other 10 instances where, in fact, the class was 1. These are called Type II Errors, wherein the model failed to detect positive cases. This would result in lost opportunities to act and save some valuable customers in an application like customer churn prediction. The model correctly predicted 90 instances of Class 1. These are the instances when, in fact, the actual label was Class 1, and the model correctly predicted it as such. This metric is important because it shows in detail how well the model identifies the positive class; especially in imbalanced datasets where Class 1 might represent more important outcomes, fraud for example, or high-value customers. In the overall performance, the performance of the Random Forest model is relatively good. The rate of true positives and true negatives is high; some areas to be adjusted are the false positives and the false negatives (Figure 2B).

Previous research noted that deep learning models, such as Neural Networks, are efficient in treating high-level complex datasets with nonlinear relationships—a very common occurrence in many business analytics applications, including customer behavior analysis and fraud detection (LeCun et al. 2015). In contrast, Breiman (2001) recognized how strong ensemble models such as Random Forests are about handling noisy data to prevent overfitting. While random forests usually perform very well on many tasks, it has been shown that their performance becomes relatively worse in the case of high-dimensional data with a very complex relationship—a fact that agrees with the ROC curve in the figure where Neural Networks show a marginally better performance. ROC Curves have also been indicated as a good metric when comparing model results on imbalanced data (Fawcett 2006). It is essentially a graph that depicts the relative trade-offs between sensitivity—true positive rate and specificity—1 - false positive rate. What this figure therefore shows is that, while both Random Forests and Neural Networks perform well, the latter does somewhat better in classification tasks that require more sophistication in terms of understanding data patterns.

4.3 Precision-Recall Curve Comparison Between Random Forest and Neural Network

These results compare the precision versus recall of two machine learning models, Random Forest and Neural Network, which have been used for classification tasks in business analytics. Precision-recall curves plot the trade-off between precision and recall for different thresholds, especially when comparing models on imbalanced datasets, where one class is more prevalent than the other. The starting lower curve values for precision (~0.70) and recall (~0.60) are for the Random Forest model. Precision grows linearly with increased recall. The model reaches a precision of 0.90 by the time it achieves a recall of 0.90. The curve seems to reflect that Random Forest handled the classification task rather well and balanced between precision, that is the ability not to report false positives, and recall, or the ability to detect true positives. At lower recall values, the performance of the Neural Network is slightly better compared to Random Forest. It starts off stronger than the Random Forest model with a precision of about 0.72

and a recall of 0.65. As the recall increases, the Neural Network goes up in precision sooner but eventually is overtaken by the Random Forest until they converge toward the higher levels of recall. The curve shows that Neural Networks tend to be more accurate in the earlier stages of recall and are, therefore, more suitable for applications where early-stage true positive detection is of essence (Figure 3A).

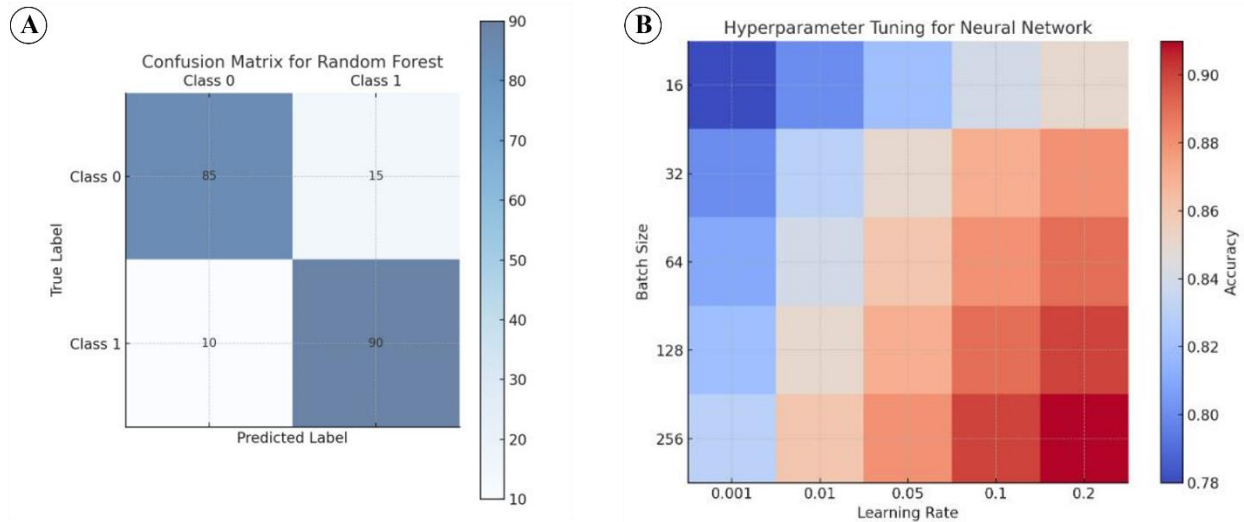


Figure 3. Precision-Recall curve comparison and hyperparameter tuning analysis for neural network and random forest models.

This is the heat map showing the accuracy of the Neural Network model, which changes according to two of the very important hyperparameters: learning rate and batch size. This color gradient—from blue, meaning lower accuracy, to red, meaning higher accuracy—offers an intuitive overview of how those parameters is affecting model performance. Small batch size - 16, low learning rates: With a small batch size, say 16, and relatively low learning rates of 0.001 or 0.01, the model presents low accuracy. The color of that region of the heatmap below is blue, meaning the accuracy could be as low as 0.78. It possibly means that with too small a batch size, the convergence may not get captured properly and hence low accuracy in models. Increasing the learning rate towards 0.2, the increase is slight whereas with a small batch size, the results are suboptimal. With increased learning rates, accuracy raises but at a low value compared to larger batch sizes. First, the heatmap reflects the gradual change towards warm colors, reflecting the increase in batch size. The most accurate results are given for either a batch size of 128 or 256 with learning rates of 0.1 or 0.2, respectively, although values are close to 0.90. That means larger batch sizes stabilize the process of gradient descent leading to better model performance. Accuracy of 0.90 is observed for a learning rate of 0.2 combined with a batch size of 256, represented by the darkest red area. Also, this brings into view that both hyperparameters are important to be tuned for the best performance (Figure 3B).

The work demonstrated that Neural Networks are good at managing complex, nonlinear relationships within the data, which is perhaps the reason why the model's precision is greater for lower recall values (LeCun et al. 2015). On the other hand, an ensemble method like Random Forest, shown to be very strong in handling overfitting and hence noise in data (Breiman 2001), thus turned out to be very strong here. While Random Forest starts with lower precision, as the early part of the curve shows, it can aggregate several decision trees. This enables its precision to improve as recall increases, which already gives it a high chance of competing favorably on tasks which require a balance between both precision and recall—for example, in predictive maintenance or churn prediction. The precision-recall curves provide useful information in situations where class balance may be an issue (Davis and Goadrich 2006). The curve on this figure suggests that Neural Networks could perform better in scenarios with higher imbalance—the cases of fraud-while Random Forest might provide good performance in datasets with balanced or gradually increasing recall needs. Smaller batch sizes tend to provide noisier gradient estimates (Masters and Luschi 2018), which may cause slower convergence and generally lower overall accuracy. On the other hand, a large batch size stabilizes the updates of the gradient, leading to smoother convergence and higher accuracy. This is reflected in the heatmap, as larger batch sizes of 128 and 256 greatly increase the accuracy with higher learning rates. The learning rate is too small, convergence is slow, whereas if it is too large, the optimal solution may be overshoot (Goodfellow et al. 2016). The heat map from this figure shows that learning rates of 0.1 and 0.2 achieve the highest accuracy, especially when combined with large batch sizes. This no doubt ascertains the fact that a learning rate from a moderate to high value will be able to ensure faster and efficient convergence in neural networks. The balancing of hyperparameters like learning rate and batch size can have a big difference in the performance of the model (Bengio 2012).

5.0 Challenges and Future Directions

Machine learning itself has become a disruptor in business analytics, allowing intelligent extraction of meaningful insights from masses of data that inform better business decisions. Despite the great strides that ML makes, there are still a host of challenges to be overcome before its complete and smooth integration into business. Addressing these open challenges would continue to offer future opportunities for research and innovation. But machine learning models can be only good to the extent a model has been trained with the quantity and quality of data. Most businesses have incomplete, inconsistent, or unstructured data. For example, data about customers can be incomplete in some fields, duplicate data could be stored, or data can be spread across different systems, which makes aggregation and cleaning hard to do for ML training. High-quality labelled data available to train the supervised learning models is very limited and it reduces the accuracy and robustness of the model. Almost 50% data scientists claimed that data quality is considered one of the major problems when building machine learning models for business purposes (Kaggle 2020). A lack of interpretability reduces the possibility of trust from the business leaders or decision-makers in those models. Because it can be any other financial institution which has to provide reasons with clarity regarding the approval or rejection of a loan regarding an ML model that calculates credit risk. The interpretability challenge becomes highly critical in such regulated domains like health and finance since decisions must be justified in relation to transparent criteria. Most businesses today have systems that are legacy and thus not designed to support machine learning models. This calls for huge infrastructure investments in the updating of databases, cloud platforms, and software tools in integrating ML into these systems (Doshi-Velez and Kim 2017). Machine learning used in business analytics raises several ethical and privacy concerns. Most machine learning models, dealing with business data, need sensitive personal information from customers: habits of purchasing, financial records, or online behaviors-all make predictions possible. The poor handling might result in data breaches or unethical practices like discriminatory algorithms. For example, the GDPR in the European Union makes the collection and usage of data subject to several strict conditions, complicating compliance when one wishes to carry out ML modeling.

However, lately, there has been a nascent area of interest called AutoML, which makes some building, training, and deployment of machine learning models a bit easier. Automating certain time-consuming tasks-imaginably, feature engineering, hyperparameter tuning, and model selection-end up making ML accessible to businesses without deep expertise in the domain. AutoML will bring a democratic situation in machine learning, reducing dependence on data scientists, and hence allow non-experts to use ML for their analytics work (Hutter et al., 2019). As one of the major challenges is model interpretability, the research fraternity has gone ahead to develop explainable AI, now popularly known as XAI. XAI techniques focus on making the machine learning models transparent by explaining the decisions made by the model. This is quite important for deep learning models, which normally sacrifice readability in favor of the predictive performance of models. Ribeiro et al., in 2016, proposed a framework named Local Interpretable Model-agnostic Explanations-LIME, which shall provide insight into individual predictions and shall help businesses understand and trust the output of ML models (Ribeiro et al. 2016). Federated learning is, in turn, a decentralized concept whereby machine learning models are trained on distributed datasets without actual sensitive information being transferred to a central location. That simply means under this kind of approach, one can keep privacy protected while enabling companies to work together on a common path toward insight. Several financial institutions would come together to build better fraud detection models using federated learning without necessarily having to share data on individual transactions, hence keeping protection laws regarding data intact. Wherever the adoption of ML is on the rise, there is a corresponding need being felt to ensure these models are unbiased and equitable. The future work should be done on developing frameworks that audit machine learning models for bias and other ethical concerns. For example, IBM AI Fairness 360 gives the means to detect and mitigate bias in machine learning models so that equitable decisions are taken. In other words, machine learning has the full potential to push business analytics to the next level, but some problems need to be solved concerning data quality, model interpretability, and ethics. Innovations in areas such as, but not limited to, AutoML, explainable AI, and federated learning would address these challenges by offering a route toward robust, transparent, and ethical applications of machine learning in businesses.

6. Conclusion

These figures give an overview of how machine learning models, especially Random Forest and Neural Networks, power business analytics by furthering the established statistical methods. From the chart comparing accuracies, one can clearly tell that machine learning models outdo the traditional approaches of linear and logistic regression by a wide margin. This is further reinforced by the ROC curve, showing the ability of Neural Networks to better balance their true positives and false positives across the thresholds. The confusion matrix of the Random Forest shows its strong points in classification; it also underlines its areas of improvement by showing quite a significant number of false positives and negatives. This contrasts with the comparison of Precision-Recall curves, which shows that Neural Networks have higher precisions at different levels of recall and are thus more reliable when the early detection of true positives is crucial for applications. Finally, the hyperparameter tuning heatmap is presented to show that careful consideration of learning rates and batch sizes will pay off; in fact, the best performance was achieved by larger batch sizes combined with higher learning rates. Aggregated, these figures show that machine learning advances not only the predictive capabilities but also scales up solutions to handle complex business data. Further innovations are

likely to be refinements of these methods to improve interpretability, efficiency, and ethical considerations in business decision-making.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

Acknowledgement: We would like to express our gratitude to all the co-authors for their contribution and critical reviews from the anonymous reviewers.

References

- [1] Bengio, Y. (2012). Practical recommendations for gradient-based training of deep architectures. *Neural Networks: Tricks of the Trade*, 437-478.
- [2] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5-32.
- [3] Brynjolfsson, E., & McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W.W. Norton & Company.
- [4] Chen, H., Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS Quarterly*, 36(4), 1165-1188.
- [5] Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3), 273-297.
- [6] Davis, J., & Goadrich, M. (2006). The relationship between Precision-Recall and ROC curves. *Proceedings of the 23rd International Conference on Machine Learning*, 233-240.
- [7] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
- [8] Draper, N. R., & Smith, H. (1998). *Applied Regression Analysis* (3rd ed.). John Wiley & Sons.
- [9] Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861-874.
- [10] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- [11] Gupta, S., Raval, N., & Dhingra, V. (2018). Machine learning applications in business: A literature review. *International Journal of Business Analytics*, 5(1), 58-78.
- [12] Guyon, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3(Mar), 1157-1182.
- [13] Hutter, F., Kotthoff, L., & Vanschoren, J. (2019). *Automated Machine Learning: Methods, Systems, Challenges*. Springer.
- [14] James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An Introduction to Statistical Learning: With Applications in R*. Springer.
- [15] Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255-260.
- [16] Kaggle. (2020). *State of Data Science and Machine Learning*. Kaggle.
- [17] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
- [18] Little, R. J., & Rubin, D. B. (2019). *Statistical Analysis with Missing Data*. John Wiley & Sons.
- [19] Masters, D., & Luschi, C. (2018). Revisiting small batch training for deep neural networks. *arXiv preprint arXiv:1804.07612*.
- [20] Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. MIT Press.
- [21] Powers, D. M. (2011). Evaluation: From precision, recall and F-measure to ROC, informedness, markedness and correlation. *Journal of Machine Learning Technologies*, 2(1), 37-63.
- [22] Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, 1(1), 81-106.
- [23] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135-1144.
- [24] Russell, S., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach* (3rd ed.). Pearson.
- [25] Sasaki, Y. (2007). The truth of the F-measure. *Teach Tutor Mater*, 1(5), 1-5.
- [26] Shmueli, G., & Koppius, O. R. (2011). Predictive analytics in information systems research. *MIS Quarterly*, 35(3), 553-572.
- [27] Sun, Y., Wong, A. K., & Kamel, M. S. (2019). Classification of imbalanced data: A review. *International Journal of Pattern Recognition and Artificial Intelligence*, 23(04), 687-719.