Blanco Rocío Rocha,
Computer Engineer, Ph.D.,
and professor in the field of Business Organization,
as well as Vice-Dean of Business, Business Law,
and Distance Learning, and Coordinator
of the Master's Degree in Business and Information Technologie;
University of Cantabria;
Ortega Angel Cobo,
Mathematician, Ph.D., and Full Professor
in the area of Applied Mathematics and Computer Science;
University of Cantabria;
Albor Lina,
MBA, and Master's in Business
and Information Technologies and Project Management;
Deduce Data Solutions

DEEP LEARNING PREDICTIVE MODELS FOR BUSINESS MANAGEMENT IN AGRICULTURE: ANALYSING AVOCADO PRICE TRENDS IN COLOMBIAN WHOLESALE MARKETS

Abstract. This study explores the application of deep learning predictive models to analyse the evolution of avocado prices in Colombian wholesale markets. The research aims to enhance decision-making processes in the agricultural sector by leveraging advanced machine learning techniques. The application of models based on recurrent neural networks has proven to be an efficient alternative for prediction processes in time series.

Introduction. Anticipating problems and detecting opportunities is one of the keys to the success of any organization. In this regard, predictive models are of enormous importance in changing contexts. The exploration of sequential data over time allows for the identification of patterns, trends, and the prediction of future events, thereby improving the decision-making process. It is equally important that both strategic and operational decisions be guided by data and information (DDDM – Data Driven Decision Making).

The agricultural sector faces significant challenges in maintaining productivity and sustainability. Companies in primary sectors, such as agriculture, especially require techniques that allow them to predict medium- and long-term key indicators such as crop production levels, the demand for agricultural products, and the prices of inputs and the sale of their products. Forecasting of agricultural variables at national scale is essential for developing policies [1]. Predictive models, particularly those based on deep learning [2], offer promising solutions for forecasting market trends and prices. This paper focuses on the application of these models to predict avocado prices in Colombian wholesale markets, providing insights that can aid in strategic planning and decision-making. The case of avocado production in Colombia has been considered as an example due to the importance of the agricultural sector in this country, and the role of this product in its agricultural exports. Predicting the price evolution of agricultural products is a particularly complex problem due to the variety of factors that can influence it (environmental, climatic, regulatory, economic, etc.). Furthermore, in recent years, there has been significant growth in the prices of certain agricultural products that has broken the price trends of previous years.

Methodology: Recurrent Neuronal Networks. To address this complex problem, we have opted to use Deep Learning (DL) techniques, an area of artificial intelligence that seeks to enable computers to identify patterns in data and make predictions [3]. The goal of Machine Learning is to learn from experience, considering data analysis as an essential element for developing intelligent systems. DL can be seen as a scalable form of ML that allows for the exploration of larger datasets, automating much of the feature extraction process, and reducing the manual human intervention required. Classical ML relies more on human intervention to learn.

The basic computational unit on which most DL techniques are based is the artificial neuron. This is a computational model inspired by the functioning of biological neurons. An artificial neuron receives numerical input values, which, when weighted by determined weights, generate a linear activity of the form $\sum_{i=1}^{n} w_i x_i + b$, upon which an activation function is applied to produce a numerical output that can then act as input for other neurons, thus forming an entire Artificial Neural Network (ANN). The weights w_i and the bias b of each neural unit must be adjusted based on training data, seeking to minimize the error between the known actual output values for that training data and the values returned by the model.

The simplest model of an ANN is the well-known multilayer perceptron, in which neurons are organized into layers, and information propagates in one direction only. These networks are known as feedforward networks. With enough neurons, the multilayer perceptron is a universal approximator for any function and are suitable for single-step predictions. However, ANN are not very efficient when dealing with sequences, especially time series data. For this type of problem, Recurrent Neural Networks (RNN) are more appropriate, they are ideal for multi-step predictions due to their ability to capture temporal dependencies. RNN provide high predictive capacity and accuracy compared to conventional statistical methods [4].

In RNN neurons receive their previous outputs as inputs, thereby creating recurrent connections. These connections provide memory to the network and improve predictive processes when previous outputs can influence future outputs. Among the most well-known examples of RNNs are SimpleRNNs, LSTM (Long Short-Term Memory) models, and GRUs (Gated Recurrent Units). Figure 1 shows the difference between a basic neuron and a recurrent one.

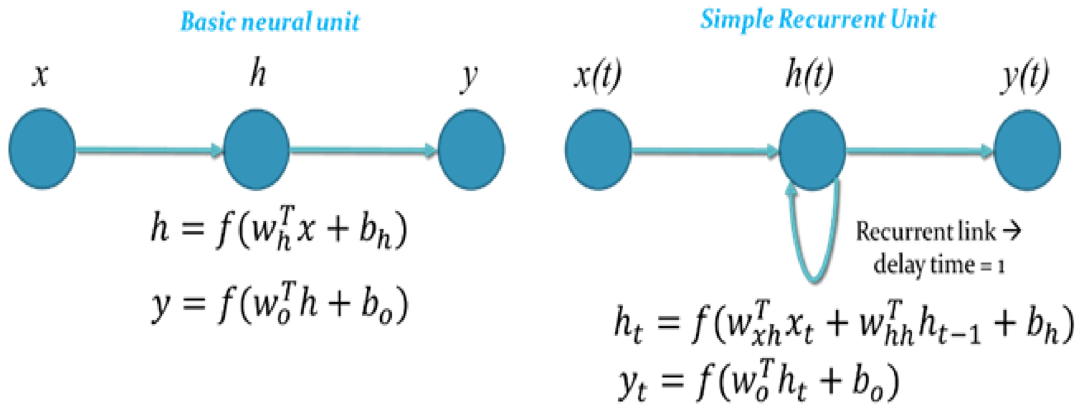

Fig. 1. Differences between a basic neuron and a recurrent neuron

Networks based on SimpleRNN take advantage of the structure of the data to create simpler models, resulting in computational savings. However, they are vulnerable to the vanishing gradient effect, leading to learning difficulties, and they also struggle to learn from inputs that are too far back in time. In this regard, LSTM and GRU networks incorporate mechanisms that allow them to find long-distance patterns, making them more useful for long-term predictions in time series [5]. GRU networks were proposed in [6] as a simplified version of LSTMs, with fewer parameters. In these networks, the output *h(t)* still depends on *x(t)* and *h(t−1),* but a mechanism is also integrated that allows them to remember or forget components of *h(t−1)*. This is achieved by including two gates with binary classifiers (logistic regression neurons): the update gate and the reset gate. The network also addresses the vanishing gradient problem, which affects the learning of long-term dependencies.

Fortunately, there are software libraries that allow users to apply this type of DL model without being experts in the underlying algorithms. In this work, we have used the TensorFlow package, an open-source library developed by Google for building and training ML and DL models. It was released in 2015 and is widely used for developing artificial intelligence (AI) applications.

Data collection and construction of an autoregressive model. To analyse the possibilities of using GRU networks in predicting the price of agricultural products, the weekly price sequence of the Hass variety avocado in the wholesale market of Bogotá, D.C., Corabastos, Colombia was obtained. The series corresponds to the average weekly prices between June 23, 2012, and April 13, 2024, with a total of 609 data points. Figure 2 shows the time series divided into two subsets; the first, corresponding to 80% of the data, will be used for training the network. The price range of the series is between 2610 and 11611 Colombian pesos, with an average value during the period of 4529.35 pesos per kilogram of avocado. In order to scale the data, the prices were considered in thousands of Colombian pesos (1000 COL $ are equivalent to approximately 0.21 euros).

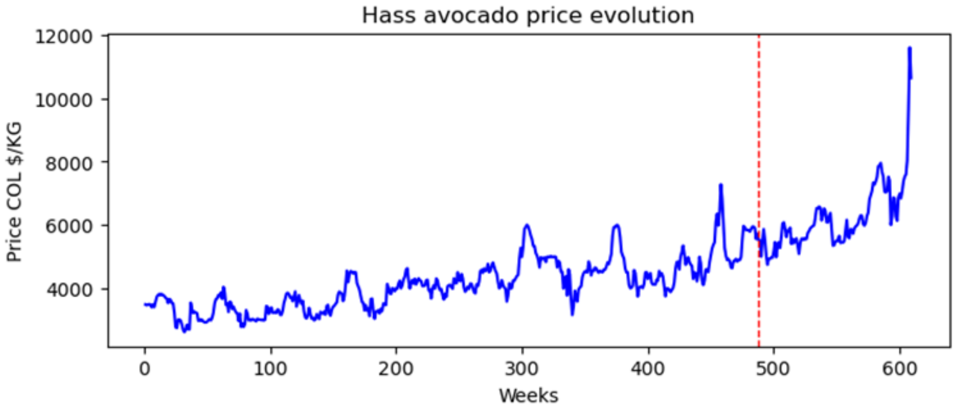

Fig. 2. Evolution of Hass avocado prices in the wholesale market of Bogotá from June 2012 to April 2024. The dashed line separates the data used for training (left) and for model validation (right)

To train the network, a dataset was generated consisting of inputs corresponding to the weekly prices of a complete year (52 weeks), with the associated output being the price of the following week. That is, the data would be in the form $(\boldsymbol{x}_i, y_i)$ where $\boldsymbol{x}_i = (p_i, p_{i+1}, \cdots, p_{i+52})$ and $y_i = p_{i+53}$. 80% of the data was reserved for training and the remaining 20% for validation. The goal of the network is to be capable of making a prediction $\widehat{p_t}$ of the medium-term prices. In this autoregressive model, two types of predictions can be made: one-step prediction $\widehat{p_t} = w_0 + w_1 p_{t-1} + w_2 p_{t-2} + \cdots + w_{52} p_{t-52}$ or a multi-step prediction

$$\widehat{p_t} = w_0 + w_1 \widehat{p_{t-1}} + w_2 \widehat{p_{t-2}} + \cdots + w_{52} \widehat{p_{t-52}}$$

where, for the price prediction, the predictions of previous weeks are used instead of the actual data.

From a practical point of view, multi-step predictions are more useful, although they are logically often less accurate. *For training the network, two models based on recurrent networks were considered:*

- *SimpleRNN*: composed of a layer of 20 simple recurrent neurons, which received 52 input values (prices for one year), and the outputs were sent to a neuron that generated the output value. The *ReLU* activation function was used with a learning rate of 0.05 in an Adam optimization algorithm.
- *GRU*: a network with 52 input values to 100 GRU neurons and a basic output neuron. The learning rate of the algorithm was reduced to 0.001.

In the following section, the results of applying both models are summarized.

Application of RNN models for price prediction. In the case of the SimpleRNN model, it is observed that the one-step predictions yield reasonable results (Figure 3 shows an example). However, it is not very useful to have a prediction model that generates only a prediction for the immediately following week. The mean prediction error with this model was 0.96 for the training data and 1.09 for the test data.

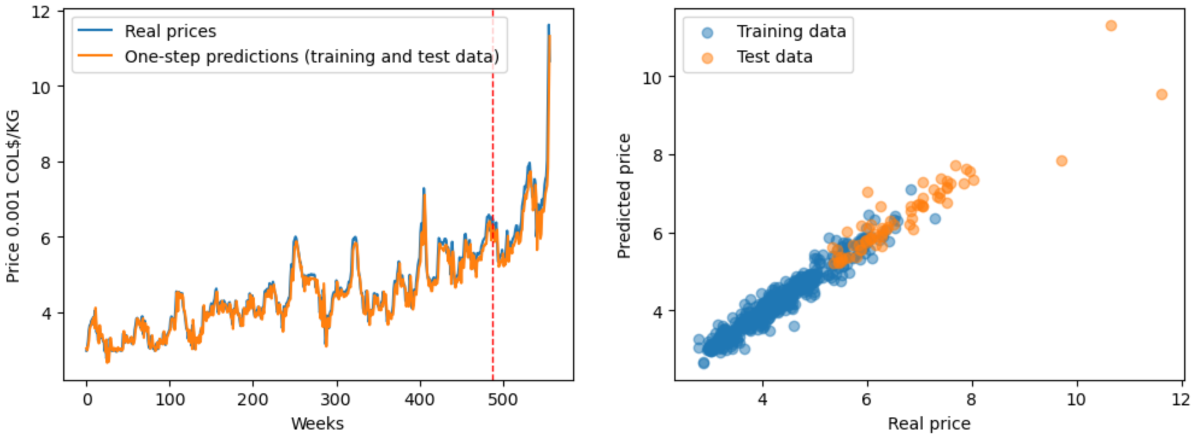

Fig. 3. One-step predictions of the SimpleRNN model

Nevertheless, when a multi-step prediction is made with the SimpleRNN model, it becomes apparent that it is not capable of capturing the medium-term trend. Figure 4 shows the multi-step predictions for the test data from both the SimpleRNN and the GRU models. As can be seen, in the first 10 weeks, the SimpleRNN seems to capture the trend well, but after that, the prediction errors become significant. In the case of the GRU model, it is observed to capture the overall trend better.

It seems obvious that in a multi-step prediction model, the errors increase the wider the time horizon for which the prediction is desired. To verify this, the GRU model was used to generate predictions across the entire dataset with annual (52 weeks), semi-annual (26 weeks), quarterly (12 weeks), or monthly (4 weeks) time horizons. In this case, the actual data from previous periods was used to try to predict these time horizons. Figure 5 graphically illustrates how prediction errors significantly decrease when the prediction time horizon is reduced. The mean errors were 0.60 for one-year predictions, 0.51 for semi-annual predictions, 0.40 for quarterly predictions, and 0.25 for monthly predictions.

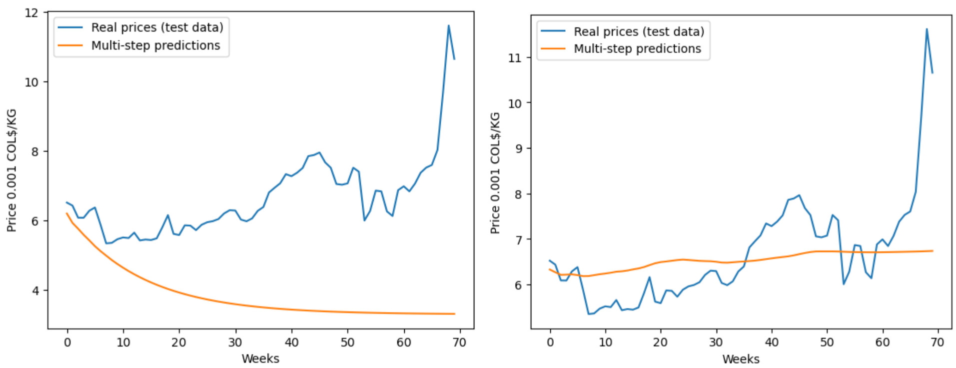

Fig. 4. Multi-step predictions on the test data for the SimpleRNN model (left) and GRU model (right)

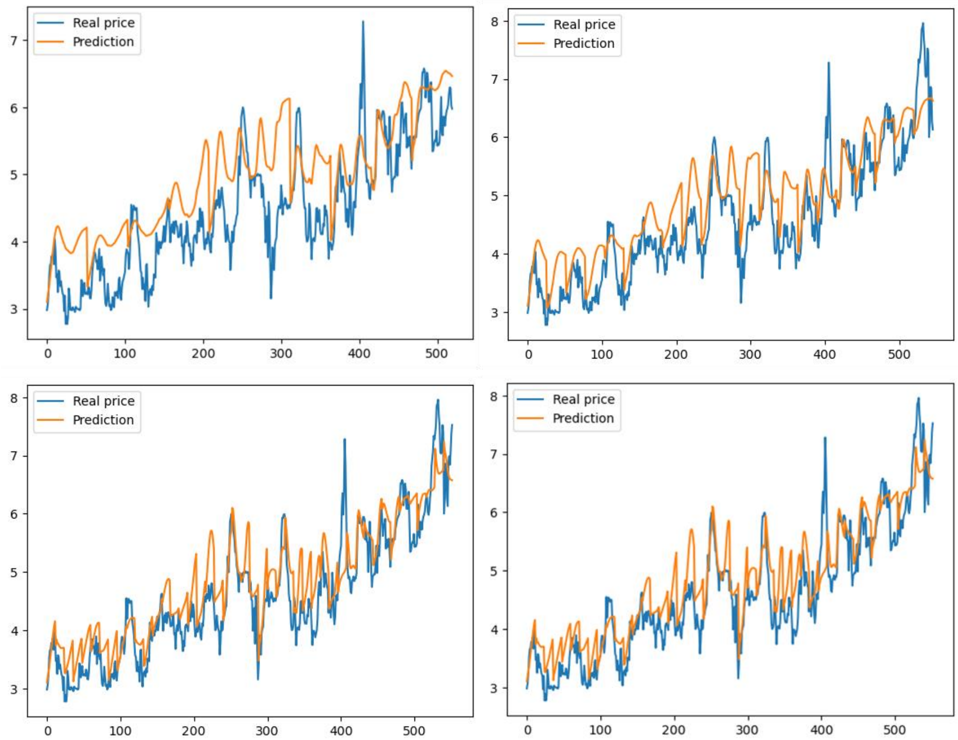

Fig. 5. Multi-step predictions on the complete dataset for annual, semi-annual, quarterly, and monthly prediction horizons (from left to right and from top to bottom)

Finally, 25 independent runs of the complete process (training the network and multi-step prediction with different forecasting periods) were conducted, yielding average values for different quality indicators shown in Tables 1 and 2. Regarding execution times, it should be noted that the experiment was conducted on a simple personal computer with a 12th Gen Intel(R) Core(TM) i7-12700 2.10 GHz processor, Windows 10 Enterprise operating system, and using TensorFlow on Spyder IDE 5.5.1 with Python 3.12.4 64-bit.

Analysis of Results and Conclusions. The predictive DL models have proven to be effective for short to medium-term predictions They demonstrated good accuracy in forecasting avocado prices. SimpleRNN was effective for short-term predictions, while GRU model provided robust multi-step forecasts, maintaining low error rates across multiple time steps. The models successfully identified patterns and trends in the price data, offering valuable insights for market analysis. The prediction errors obtained were smaller when the required prediction time frame was shorter. In any case, these are considered reasonable errors given the complexity of the problem. As a potential improvement to the models, the incorporation of other variables related to crop production and climatic factors could be considered. However, the results of the simple model analysed can still be considered satisfactory

The application of deep learning models in the agricultural sector shows significant potential for enhancing predictive accuracy and supporting sustainable development. By accurately forecasting market trends, these models can help stakeholders make informed decisions, anticipate market fluctuations, and optimize resource allocation. This study demonstrates the effectiveness of deep learning models in predicting avocado prices in Colombian wholesale markets. The findings highlight the potential of these models to improve decision-making processes and support sustainable agricultural practices. Future research could expand the application of these models to other agricultural products and markets, further enhancing their utility and impact.

Table 1

Average execution times over 25 independent runs

Time (in seconds)	*Training time*	*Prediction time (one year ahead)*	*Prediction time (one semester ahead)*	*Prediction time (one quarter ahead)*	*Prediction time (one month ahead)*
RNN	7.0779	18.5339	19.2980	19.5623	19.7772
GRU	69.4162	18.5904	19.6062	20.0458	20.1080

Table 2

Mean errors, standard deviations, and Root Mean Square Errors averaged over 25 independent runs for each prediction time horizon

	Predictions one year ahead			*Predictions one semester ahead*		
	Mean error	*σ*	*RMSE*	*Mean error*	*σ*	*RMSE*
RNN	1.6236	0.9826	1.9122	0.8897	0.6821	1.1237
GRU	0.4786	0.3833	0.6142	0.4181	0.3730	0.5609
	Predictions one quarter ahead			*Predictions one month ahead*		
	Mean error	*σ*	*RMSE*	*Mean error*	*σ*	*RMSE*
RNN	1.0748	1.8793	2.2399	0.4151	0.3807	0.5689
GRU	0.3450	0.3148	0.4679	0.2498	0.3230	0.4086

Acknowledgments. This work has been carried out with the support of the CYTED network: AI4AGROIB “*Red iberoamericana para el desarrollo de sistemas inteligentes para la agricultura*”, from the Ibero-American Program for Science and Technology for Development.

References

1. Cesarini, L., Gonçalves, R., Martina, M., Romao, X., Monteleone, B., Lobo, F., & Figueiredo, R. (2024). Comparison of deep learning models for milk production forecasting at national scale. *Computers and Electronics in Agriculture* 221, 1-11, 108933.

2. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* 521, 436–444 (2015).

3. Shinde, P. P., & Shah, S. (2018). A review of machine learning and deep learning applications. In 018 Fourth International Conference on Computing Communication Control and Automation (ICCUBEA), Pune, India, 2018, pp. 1-6.

4. Zeebaree, D. Q., Abdulazeez, A. M., Abdullrhman, L. M., Hasan, D. A., & Kareem, O. S. (2021). The prediction process based on deep recurrent neural networks: a review. *Asian Journal of Research in Computer Science*, 11(2), 29-45

5. Chandra, N., Ahuja, L., Khatri, S. K., & Monga, H. (2021). Utilizing gated recurrent units to retain long term dependencies with recurrent neural network in text classification. *J. Inf. Syst. Telecommun*, 2, 89.

6. Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing* (EMNLP), pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.

Dehtiar Nataliia,
Ph.D. in Economics,
Senior Lecturer,
Department of Economics and Entrepreneurship,
Faculty of Economics and Management;
Chen Xiao,
Master Student,
Faculty of Economics and Management;
Sumy National Agrarian University

THE REGIONAL ECONOMIC EFFECTS OF BEIJING-TIANJIN-HEBEI INTEGRATION: ECONOMIC RESILIENCE AND DEVELOPMENT PROSPECTS

Contemporary China is developing regional integration as a strategy to promote economic growth, reduce inequality and poverty, and remove barriers to the flow of capital, goods and services, human capital, and innovative ideas. In response to the growing challenges and disproportion of socio-economic, ecological development of Northern China, during the last decade a successful strategy of regional integration "Beijing-Tianjin-Hebei" (BTH) was implemented.

The Beijing-Tianjin-Hebei integration strategy aims to foster economic resilience, enhance regional connectivity, and stimulate sustainable development across one of China's most economically significant regions. The BTH region received the conditional name "third pole" of the Chinese economy (after such regions as "Yangtze River Delta Region" and "Pearl River Delta Region") [1]. This BTH region includes the city of Beijing, the city of Tianjin and 11 prefecture-level cities of Hebei province, and also unites more than 120 million people [2].

It should be noted that the implementation of such an integration strategy made it possible to ensure the resilience of the BTH region in conditions of high pressure