Predicting Vendor Performance in Data-Scarce Environments: A Hybrid Deep Learning Approach

Nacereddine Bouriche[1] [*]and Messaoud Djeddou[2]

1 University Mohamed Khider, department of Management, ECOGES laboratory, Biskra 07000, Algeria.

2 University Mohamed Khider, LARHYSS laboratory, Biskra 07000, Algeria.
nacereddine.bouriche@univ-biskra.dz

Abstract. In the era of Industry 4.0, the imperative for proactive vendor risk management is predicated upon the precise forecasting of future performance. Nevertheless, the deployment of sophisticated Deep Learning (DL) architectures such as Long Short-Term Memory (LSTM) networks remains markedly constrained in developing economies owing to a profound paucity of historical time-series data. This study interrogates the cold start phenomenon by advancing an original Generative-Predictive Framework that fuses Variational Auto-Encoders (VAEs) with LSTM networks. Applied to a leading entity within the Algerian glass manufacturing sector (Mediterranean Float Glass: MFG), the framework utilises VAEs to reconstruct latent historical performance trajectories (2015– 2019) from a limited dataset (2020–2023), thereby affording robust LSTM training for long-horizon forecasting (2026). To substantiate the methodology, a visual trajectory analysis is conducted for three representative vendor archetypes: Stable, Stagnant, and Deteriorating. The findings demonstrate that this hybrid approach efficaciously alleviates data scarcity, producing risk-aware forecasts that empower the buyer (MFG) to transcend reactive evaluation in favour of proactive, strategically informed sourcing.

Keywords: Vendor Performance Forecasting, Data Scarcity, Deep Learning, LSTM, Variational Auto-Encoder (VAE), MENA Region.

1 Introduction

Globalized trade and the rapid digitization of supply networks are predominantly reshaping the industrial environment. In such a volatile atmosphere, the resilience of a manufacturing enterprise depends less on internal efficiency and more on the predictive health of its upstream supply chain [1]. Vendor management, formerly regarded as a reactive administrative task, has thus evolved into a strategic imperative, wherein the capacity to anticipate fluctuations in performance determines long-term competitive advantage [2].

D. Agti et al. (eds.), *Proceedings of the International Conference on Artificial Intelligence Applications in Business Administration in MENA Region (ICAIABA 2026)*, Advances in Economics, Business and Management Research 393,
https://doi.org/10.2991/978-94-6239-711-8_35

Whilst Industry 4.0 promises a revolution in "Smart Procurement", the practical realization of this vision remains inconsistent. Recent academic research has successfully employed advanced deep learning architectures—such as Long Short-Term Memory (LSTM) networks [3] and Genetic Algorithms [4] to optimize procurement flows. Nonetheless, these models often function under the assumption of "Big Data" accessibility. In developing economies, particularly in the Middle East and North Africa (MENA) area, industrial stakeholders encounter a significant "Digital Divide". Organizations in these locations often exhibit exceptional operational capabilities but are hindered by a heritage of disjointed, non-digitized record-keeping.

This engenders a systemic "Cold-Start" problem: practitioners are unable to benefit from state-of-the-art artificial intelligence because their data windows are insufficient to capture intricate temporal dependencies [5]. To the best of our knowledge, no extant research has addressed this specific limitation by integrating generative models with time-series forecasting in the context of vendor forecasting performance.

This research bridges a critical gap, offering a seminal methodological advance within the field. We introduce a Hybrid Generative-Predictive Framework, uniting Variational Auto-Encoders (VAE) with LSTM networks. When applied to a case study of Mediterranean Float Glass (MFG) in Algeria [6], this approach facilitates the synthetic reconstruction of latent historical trends (2015–2019) from limited real-life datasets. This innovation, for the first time, enables robust long-horizon forecasting under conditions of data scarcity, empowering the buyer to progress from reactive evaluation to proactive, strategically informed sourcing decisions.

2 Literature Review

2.1 From Measurement to Prediction

The scholarly discourse on vendor evaluation has evolved through three discernible phases. The initial phase was characterised by a predominance of transactional, cost-centric models. The subsequent phase witnessed the advent of Multi-Criteria Decision Making (MCDM), broadening the evaluative lens to encompass a more comprehensive array of metrics. Notably, recent studies have demonstrated the efficacy of integrating holistic measurement instruments especially with the Analytic Hierarchy Process (AHP) and TOPSIS [7] to assess vendor performance. Whilst such frameworks have addressed the challenge of measurement, they remain intrinsically retrospective serving to diagnose historical deficiencies rather than anticipate prospective risks [8].

The third and prevailing era is distinguished by the ascendancy of Predictive Analytics. As observed by Gidiagba et al. [5], Machine Learning (ML) now enables organisations to advance from mere "Evaluation" to the more sophisticated domain of "Forecasting". Nonetheless, the extant body of research remains primarily confined to binary classification tasks. For example, Steinberg et al. [9] employed ML to predict rudimentary "late" versus "on-time" delivery outcomes, whilst Achatbi [10] concentrated on reliability scoring. The present study endeavours to transcend these limitations by forecasting granular, multidimensional performance scores rather than simplistic categorical labels.

2.2 The Challenge of Small Data in Deep Learning

A principal impediment to the widespread adoption of Deep Learning within industrial domains is the acute scarcity of consistent time-series data. Neural networks necessitate extensive datasets to minimise error; when deployed on "Small Data", they are frequently susceptible to overfitting. Oluwatomisin et al. [11] underscore that forecasting variability becomes unreliable in the absence of adequate historical depth. Likewise, Presciuttini et al. [12] contend that the integration of emergent technologies mandates robust data foundations to underpin sustainable performance.

While Shidpour et al. [13] recently proposed frameworks for vendor selection, and Gabellini et al. [21] investigated probabilistic supply risk forecasting, the specific challenge of generating missing historical data remains underexplored. This identifies a critical need for a generative intermediary that can expand the dataset before the forecasting phase begins.

2.3 Generative AI as a Methodological Solution

To address the challenge of insufficient historical data, this study adopts a novel statis- tical approach, inspired by recent advances in the field. Mohammed [22] contends that rigorous adaptation of analytical techniques to the peculiarities of the available evidence can significantly strengthen risk evaluation. Here, we implement the framework of Variational Auto-Encoders (VAE), as originally articulated by Kingma and Welling [23]. Rather than relying on elementary linear interpolation, this method enables the creation of plausible historical records that retain the subtle, non-linear patterns observed in vendor performance. By translating limited observations into a condensed latent representation, this methodology captures the underlying volatility and permits the reconstruction of a credible operational narrative.

3 Methodology

The proposed framework operates in two sequential phases to transform the cold-start dataset into a robust forecasting tool.

3.1 Data Acquisition and Pre-processing

The empirical dataset was attained from the procurement archives of Mediterranean Float Glass (MFG), a leading joint-stock company (SPA) operating within the Algerian industrial sector. The raw data comprised annual performance evaluations for thirteen strategic vendors across five main criteria. Before feeding this data into the neural networks, we performed rigorous preprocessing. All performance scores were normalised to a uniform scale. This scaling is essential in deep learning to prevent criteria with larger numerical ranges from disproportionately influencing the model's weight updates compared to smaller-weighted criteria.

3.2 Phase I: Generative Augmentation (VAE)

To address the cold-start problem inherent in sparse historical data, we employed a Variational Autoencoder (VAE) [14]. Unlike conventional imputation methods that merely interpolate missing values, the VAE learns the underlying probability distribution governing vendor performance patterns [15]. This generative approach ensures that synthesized historical records maintain statistical consistency with observed operational behaviour.

The VAE architecture comprises two principal components working in tandem. The encoder network compresses the available sparse data spanning 2020–2023 into a lower-dimensional latent space, a compact probabilistic representation that encapsulates the fundamental characteristics of vendor behaviour [16]. Subsequently, the decoder network samples from the learned latent distribution to reconstruct a complete time series spanning the whole nine-year period from 2015 to 2023. By optimizing the evidence lower bound (ELBO), which minimizes reconstruction error while regularizing the latent space [14,17], the model generates statistically plausible data for the missing years (2015–2019) that align with the vendor's documented operational profile.

3.3 Phase II: Long-Term Forecasting with LSTM Networks

Having established a comprehensive nine-year historical baseline through VAE-based augmentation, we implemented a Long Short-Term Memory (LSTM) network for multi-horizon forecasting [18]. LSTMs address the vanishing gradient problem that limits standard Recurrent Neural Networks (RNNs) by employing a sophisticated gating architecture that electively maintains pertinent information across extended temporal sequences [19]. The LSTM cell adjusts information propagation through three specific gates: (A) The forget gate dictates which features of the historical state vector represent transient noise and should be attenuated; (B) The input gate assesses incoming observations such as abrupt price fluctuations or delivery delays and assesses their significance for long-term trend modelling; and (C) The output gate synthesizes the filtered memory state to create forecasts for the target horizon spanning 2024–2026 [20].

4 Empirical Case Study

4.1 Industrial Context: Mediterranean Float Glass (MFG)

The framework undergoes rigorous evaluation within Mediterranean Float Glass (MFG), a leading Algerian joint-stock company (SPA: Société par Actions), a subsidiary of the Cevital Group. MFG stands as a leading force in Algerian industry, distinguished by its expertise in high-performance flat glass [14]. The company's remarkable logistical achievements notably, its status as the world's first exporter of the "Jumbo" glass format by sea exert significant demands upon its upstream supply chain. Even minor disruptions in vendor performance may precipitate systemic risks, threatening the Just-in-Time logistics that underpin MFG's relationships with international partners across Europe and Latin America.

4.2 Performance Evaluation Criteria

MFG undertakes vendor evaluation by means of a five-dimensional multi-criteria Score (maximum 100 points), crafted to afford a comprehensive portrait of each vendor's operational standing. Table 1.

Table 1. Vendor performance evaluation criteria and weighting at MFG.

Criterion	Weight (Max)	Operational Definition & Strategic Value
Product/Service Quality	25 pts	Assesses technical compliance with MFG's HSE standards, production durability, and the efficacy of warranty execution.
Meeting Deadlines	25 pts	Evaluates the vendor's ability to adhere to delivery schedules and maintain quantity accuracy,, vital for maritime logistics.
Customer Service	25 pts	Measures responsiveness to technical claims, cooperation in new product development, and alignment with corporate values.
Price Evolution	20 pts	Monitors price stability over the fiscal year, cost-reduction initiatives, and the transparency of cost-breakdown structures.
Payment Term	5 pts	Evaluates the flexibility of financiak arrangements and credit terms, contributing to MFG's overall cash-flow resilience.

4.3 Data Constraints and Sampling Strategy

The principal impediment to the adoption of advanced analytical methods at MFG lies in the paucity of reliable digital records predating 2020. Although the company has operated continuously since 2007, the systematic recording of comprehensive performance metrics was only instituted in recent years. For the purposes of illustrating the efficacy of our methodological approach, three representative cases were drawn from the strategic raw materials portfolio. In accordance with commercial confidentiality, these profiles are referred to here as Vendor A, Vendor B, and Vendor C.

5 Analysis, Results and Discussion

This section illustrates the results of the VAE-LSTM framework. Rather than relying on static tabular data, our analysis focuses on the longitudinal trajectories visualised in the figures below. This visual approach allows us to understand better the dynamic behaviour and stability of vendor performance over the extended 2015–2026 horizon.

5.1 Phase I Analysis: Generative Historical Reconstruction

The initial phase of the study focused on validating the generative capabilities of the VAE. As observed in the dashed grey lines across all three figures (2015–2019), the VAE did not merely produce a flat average. Instead, it successfully reconstructed realistic fluctuations or "latent noise" that mirror the natural volatility of industrial SC. This reconstruction provided the LSTM network with a sufficiently deep "memory bank" to found a powerful starting point before attempting to forecast the future.

5.2 Profile A: The High-Performing Benchmark (Vendor A)

Figure 1 visualizes the trajectory of a "Stable" partner. The green solid line (2020–2023) and the subsequent red dotted forecast (2024–2026) demonstrate a relationship defined by consistency.

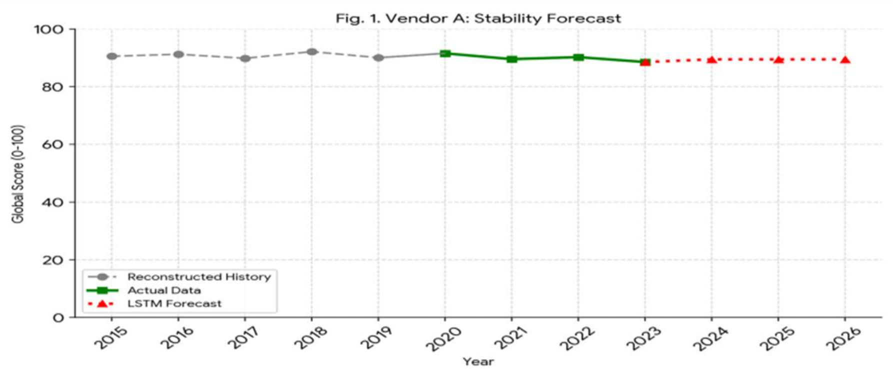

Fig. 1. Vendor A: Stability Forecast (High Performance).

Analysis and interpretation

A close examination of the trajectory for Vendor A reveals a narrow band of oscillation, with values remaining steadily between 88.0 and 90.0. The projected outlook, as indicated by the red dotted line, extends this stability into the future, exhibiting minimal deviation from the established average. No evidence emerges of underlying structural risk. The sustained absence of downward volatility suggests that this vendor has achieved a resilient operational balance. For MFG, such a visual pattern substantiates Vendor A's status as a dependable, low-maintenance partner, well-suited for long-term strategic collaboration.

5.3 Profile B: The Consistent Incumbent (Vendor B)

Figure 2 shows the so-called ‘Stagnant’ scenario. In contrast to the subtle yet healthy oscillations observed with Vendor A, the trajectory for Vendor B is strikingly unvaried, displaying an almost unbroken smoothness throughout. This pattern raises questions regarding dynamism and adaptability within the supplier’s operational profile.

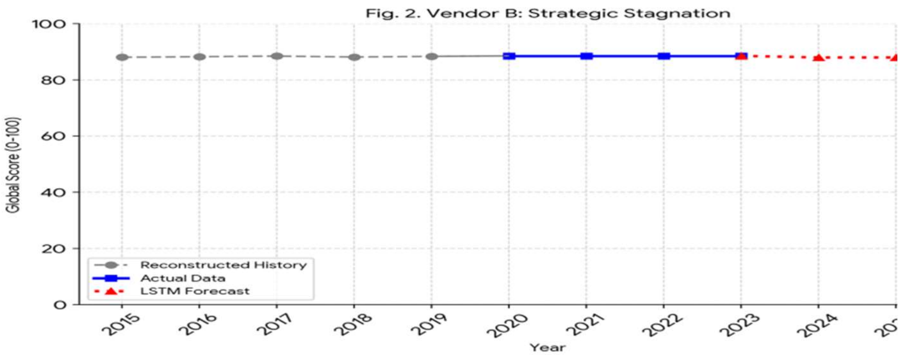

Fig. 2. Vendor B: Strategic Stagnation (Flat Trajectory).

Analysis and interpretation

The graph shows a distinguishing plateau effect. The shift from the solid blue line (actual data) to the red dotted forecast indicates a near-perfect horizontal continuation, sitting just below the 88.0 level. Whereas the performance is technically adequate, the visual absence of any upward gradient is critical. The LSTM anticipates operational inertia: it expects this vendor to continue meeting basic requirements but is unlikely to drive innovation. This visual diagnosis of Stagnation serves as a trigger for management to initiate vendor development programs to break the inertia.

5.4 Profile C: The Volatile Partner (Vendor C)

Figure 3 demonstrates the most critical analytical finding. It shows a vendor with a high-volatility “V-shaped” history, dictating a clear visual warning against recency bias.

Analysis and interpretation

The graph is distinguished by two key visual features: a Structural Failure in 2021/2022 (when performance collapsed below 40.0) and a Deceptive Recovery in 2023 (a sharp climb to 92.0). Significantly, the LSTM forecast does not follow the upward momentum of the 2023 recovery. Instead, it breaks downward immediately in 2024, returning to a much lower range. That is the model's "Risk Intelligence" in action. A simple linear regression would have followed the 2023 up-ward slope and predicted a perfect score. However, the LSTM sees the deep trough of 2021 and interprets the

2023 recovery as a temporary outlier rather than a sustainable fix. The graph visually supports the diagnosis of "Deterioration", warning MFG that the risk of another collapse remains high despite the recent good year.

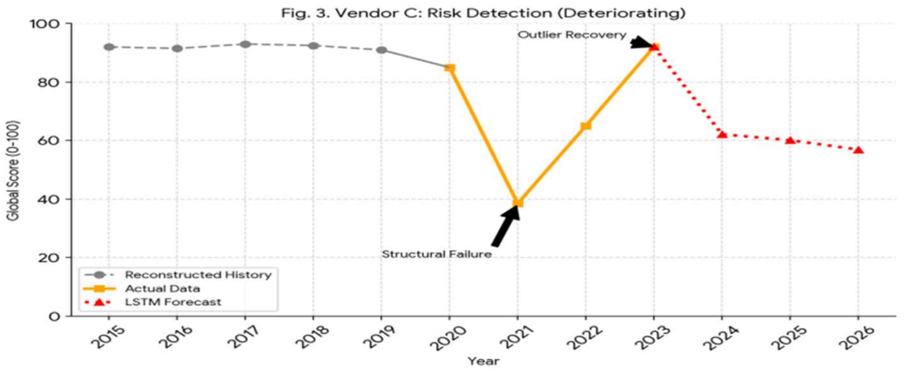

Fig. 3. Vendor C: Risk Detection (Deteriorating Profile).

6 Conclusion and Recommendations

This experiment validates a VAE-LSTM approach for forecasting vendor performance under data-scarce conditions. By rebuilding latent history, the adopted approach effectively qualifies the Deep Learning model where standard models fail. The adopted approach effectively revealed sustainable stability (Vendor A) and deceptive recovery (Vendor C), offering procurement managers a risk-aware shield.

MFG management should integrate this predictive scoring into quarterly reviews to transition from reactive troubleshooting to "proactive" sourcing. Future research should incorporate exogenous variables (e.g., market indices) to refine the model's sensitivity to external shocks.

Acknowledgements: The authors would like to thank the management and supply chain division of Mediterranean Float Glass (MFG) for granting access to the needed datasets that underpin this research. We also appreciate the constructive comments from the anonymous reviewers.

References

1. Mahamuni, A.: Internet of things, machine learning, and artificial intelligence in the modern supply chain. Def. Transp. J. 74(1), 14–17 (2018).
2. Kumar, P., et al.: Revolutionizing sourcing with AI: Harnessing technology for efficiency. World J. Adv. Res. Rev. 23(2), 925–938 (2024)
3. Pacella, M., Papadia, G.: Evaluation of deep learning with LSTM networks for time series forecasting. Procedia CIRP 99, 604–609 (2021)

4. Bubak, A., et al.: An LSTM network-based genetic algorithm for integrated procurement. Int. J. Prod. Res. 63(11), 4036–4065 (2025)
5. Gidiagba, J.O., et al.: A systematic review of machine learning applications in sustainable supplier selection. Decis. Anal. J. 14, 100547 (2025)
6. Mediterranean Float Glass : Company presentation: Le verre accède à une nouvelle ère. Cevital (2022)
7. Bouriche, N., Kishk, M., Wilson, G.: An Integrated Performance Prism-Based Framework for Supply Chain Performance Measurement. In: Proc. 9th Int. Conf. Ind. Eng. Appl. (ICIEA), pp. 917–923. WCSE (2022). https://doi.org/10.18178/wcse.2022.04.185
8. Ma, Q., Li, H.: A decision support system for supplier quality evaluation. Expert Syst. Appl. 242, 122746 (2024)
9. Steinberg, F., et al.: A novel machine learning model for predicting late vendor deliveries. Supply Chain Anal. 1, 100003 (2023)
10. Achatbi, I.: Leveraging machine learning to forecast vendor reliability. Trans. Eng. Comput. Sci. 12(6), 54–63 (2024)
11. Oluwatomisin, A.J., et al.: Machine learning models for forecasting vendor performance variability. J. Mach. Learn. Deep Learn. 1(1), 1–15 (2024)
12. Presciuttini, A., et al.: Integrating lean management and emerging technologies for enhanced sustainable performance. Data Sci. Manag. (2025)
13. Shidpour, H., et al.: A multi-phase analytics framework for supply chain vendor selection. Supply Chain Anal. 12, 100172 (2025)
14. Kingma, D.P., Welling, M.: Auto-encoding variational Bayes. In: Proc. International Conference on Learning Representations (ICLR) (2014)
15. Doersch, C.: Tutorial on variational autoencoders. arXiv preprint arXiv:1606.05908 (2016).
16. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and approximate inference in deep generative models. In: Proc. International Conference on Machine Learning (ICML), pp. 1278–1286 (2014)
17. Blei, D.M., Kucukelbir, A., McAuliffe, J.D.: Variational inference: A review for statisticians. J. Am. Stat. Assoc. 112(518), 859–877 (2017)
18. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Comput. 9(8), 1735–1780 (1997)
19. Graves, A.: Supervised Sequence Labelling with Recurrent Neural Networks. Springer, Heidelberg (2012)
20. Gers, F.A., Schmidhuber, J., Cummins, F.: Learning to forget: Continual prediction with LSTM. Neural Comput. 12(10), 2451–2471 (2000)
21. Gabellini, M., et al.: Investigating the potential of machine learning and deep learning models in probabilistic supply risk forecasting. IFAC-Papers On Line 59(10), 2957–2962 (2025)
22. Mohammed, I.A.: Machine learning algorithms for enhancing vendor risk assessment. In: Contemp. Approaches Multidiscip. Res. Theory Pract. Impact, pp. 131–149 (n.d.)
23. Kingma, D.P., Welling, M.: 3-encoding variational Bayes. In: Proc. ICLR (2014)