

TimeGAN-VAE Data Generation for a Malaysian Aviation Supply Chain Stress Index

Bao Lin Tan

Faculty of Computer Science and Information Technology
Universiti Malaysia Sarawak
Kota Samarahan, Sarawak, Malaysia
24030448@siswa.unimas.my

Sarah Samson Juan*

Faculty of Computer Science and Information Technology
Universiti Malaysia Sarawak
Kota Samarahan, Sarawak, Malaysia
sjsflora@unimas.my

Abstract

This paper presents a hybrid TimeGAN-VAE-PCA framework for generating synthetic high-frequency time-series data under limited and fragmented reporting conditions in the Malaysian aviation supply chain. The proposed approach combines the temporal modelling capability of TimeGAN with the probabilistic regularisation of a Variational Autoencoder (VAE) and the dimensional interpretability of Principal Component Analysis (PCA) to construct a Supply Chain Stress Index (SCSI). Using aviation logistics data spanning 2017–2024, the framework synthesises consistent multivariate indicators despite irregular reporting frequencies. Experimental results show that the hybrid model achieves lower reconstruction errors (RMSE ≈ 0.012), improves temporal similarity (ACF +42%), and reduces runtime by approximately 50% compared with TimeGAN alone. The resulting SCSI captures real-world disruption dynamics, including the sharp stress surge during the 2020–2021 COVID-19 pandemic and the subsequent recovery period. These findings demonstrate the potential of hybrid generative modelling as a practical tool for monitoring systemic supply-chain stress in data-scarce environments and for supporting data-driven decision-making in aviation logistics.

CCS Concepts

• Applied computing; • Forecasting;

Keywords

Generative Adversarial Network, Variational Autoencoder, Stress Index, Time-series Analytics

ACM Reference Format:

Bao Lin Tan and Sarah Samson Juan. 2026. TimeGAN-VAE Data Generation for a Malaysian Aviation Supply Chain Stress Index. In *2026 15th International Conference on Software and Computer Applications (ICSCA 2026)*, February 03–05, 2026, Langkawi, Malaysia. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3801619.3801672>

*Corresponding author

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only. Request permissions from owner/author(s). *ICSCA 2026, Langkawi, Malaysia*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2291-2/2026/02

<https://doi.org/10.1145/3801619.3801672>

1 Introduction

Global supply chains have become increasingly volatile due to the combined effects of the COVID-19 pandemic, digital transformation, and geopolitical instability. The aviation sector, in particular, remains highly vulnerable due to its reliance on Maintenance, Repair, and Overhaul (MRO) services, fuel logistics, and stringent regulatory cycles. When disruptions occur, airlines face prolonged maintenance turnaround times and shortages of critical spare parts, resulting in higher operating costs and reduced fleet availability [1, 2]. The 2024 McKinsey Global Supply Chain Leader Survey reported that 93% of executives planned to maintain or even reduce inventory buffers, reflecting a persistent reliance on lean operations despite recent systemic shocks [3]. Such findings underscore an urgent need for data-driven tools capable of detecting and characterising supply-chain stress under uncertainty.

In Malaysia, the problem is compounded by data scarcity and fragmentation. Key economic and operational indicators, such as freight indices, jet fuel prices, and air cargo volumes are often published at inconsistent temporal resolutions or embedded in unstructured formats such as PDF bulletins. This makes it difficult to construct continuous, high-frequency time series for monitoring systemic stress across aviation supply chains [4]. Traditional interpolation techniques, such as linear or cubic spline interpolation [5], can bridge some temporal gaps but fail to reproduce the nonlinear dynamics of real-world disruptions, especially during crises like the 2020–2021 pandemic period [6].

Recent advances in Generative Artificial Intelligence (GenAI) offer new opportunities to address these challenges. Models such as the Time-Series Generative Adversarial Network (TimeGAN) and the Variational Autoencoder (VAE) are capable of learning complex temporal and distributional patterns from limited data. TimeGAN excels in capturing nonlinear, sequential dependencies [7], while VAE ensures stability and smoothness in the reconstructed data [8]. However, when applied independently, each model has distinct weaknesses: TimeGAN may amplify volatility, whereas VAE may oversmooth dynamic fluctuations [9].

To address these challenges, this paper proposes a hybrid generative framework for constructing a Malaysian aviation supply chain stress indicator under conditions of fragmented data availability. The main contributions of this work are as follows:

- A hybrid TimeGAN-VAE framework that combines adversarial temporal learning with probabilistic latent regularisation to synthesise realistic multivariate time-series data.
- Integration of synthetic indicators with Principal Component Analysis (PCA) to derive an interpretable Supply Chain Stress Index (SCSI).




Figure 1: Illustration of the reparameterization trick in a Variational Autoencoder, showing how stochastic sampling is made differentiable for backpropagation.

- A comparative analysis of VAE, TimeGAN, and hybrid architectures using distributional and temporal similarity metrics.
- Demonstration of the framework on Malaysian aviation logistics data (2017–2024), capturing disruption dynamics during the COVID-19 pandemic.

2 RELATED WORK

2.1 Data scarcity and limitations of classical interpolation

This literature review explores the potential of generative AI techniques and well-established methods to address data scarcity. The traditional, mature method of data augmentation is from interpolation, from linear to nonlinear. One well-known nonlinear interpolation method is the piecewise cubic spline, widely used in financial institutions. However, a case study on imputing monthly values from quarterly time series in the Swiss Business Cycle Data shows the poor usefulness of cubic spline interpolation in multivariate settings [11]. Based on the study, cubic spline interpolation yielded RMSE values ranging from 6 to 23, indicating that it is only useful for a single dataset, not multivariate data. Also, cubic spline interpolation is incapable of detecting abnormal fluctuation, such as a major COVID-19 disruption. Recent reliability and safety engineering studies emphasise that stress indicators constructed from sparse or fragmented data are highly sensitive to imputation bias, particularly under crisis regimes. This has motivated the adoption of generative models capable of learning system-level dependence structures rather than pointwise interpolation.

2.2 VAE and GAN techniques

Generative Artificial Intelligence (GenAI) offers an alternative paradigm for data augmentation by learning probabilistic representations of underlying data distributions. Variational Autoencoders (VAEs) and Generative Adversarial Networks are two well-known data generation techniques under the GenAI umbrella, but with fundamentally different concepts. The generation of additional datasets in VAE is driven by reparameterization of continuous latent variables and by differentiable encoders and generative models, via differentiation with respect to variables x and ϕ , respectively. From a probabilistic perspective, VAEs approximate the data-generating process via a latent-variable model optimised through variational

inference, yielding stable reconstructions at the expense of attenuated tail behaviour and regime sharpness. This theoretical property explains their tendency to over-smooth abrupt shocks in heterogeneous economic time series. The details of the reparameterization trick are depicted in Figure 1, adapted from [8].

The diagram illustrates the reparameterization trick, which addresses the issue that the original stochastic sampling of z blocks gradient flow during backpropagation. By rewriting the random variable as a deterministic function $z = g_\phi(x, \epsilon)$ of the encoder parameters and an independent noise source, ϵ , the stochasticity is externalised. This transformation allows gradients to flow deterministically through ϕ , enabling efficient end-to-end optimisation of the encoder network via stochastic gradient descent. Meanwhile, GANs incorporate deep learning-based strategies, such as convolutional neural networks (CNNs) [7]. In other words, following recent reliability-engineering studies that modify TimeGAN to enable robust augmentation with limited operational data, this work adopts a hybrid generative strategy that combines adversarial temporal learning with probabilistic regularisation [12]. The objective function GAN is defined as per equation (1) below and also depicted in Figure 2, showing the GAN workflow based on the equation, showcasing the complexity of the architecture compared to VAE.

$$G_{\min} D_{\max} V(D, G) = E_{x \sim P_{\text{data}}(x)} [\log D(x)] + E_{z \sim p_z(z)} [\log (1 - d(G(z)))] \quad (1)$$

where G is the Generator, D is the Discriminator, $P_{\text{data}}(x)$ is the real-world data distribution, $P(z)$ is the generator distribution, x is the P_{data} sample (x), z is a sample taken from $P(z)$, $D(x)$ is a discriminator network, $G(z)$ is the generator network [7].

Differences in the underlying architectures of VAEs and GANs explain their contrasting performance on heterogeneous, multivariate data. Jamotton and Hainaut [9] report that while VAEs can generate realistic synthetic insurance datasets, doing so requires extensive preprocessing and careful tuning of network architecture and regularisation to handle complex, heterogeneous features. In contrast, Tai et al. [13] demonstrate that TimeGAN effectively synthesises multivariate temporal data, with GRU-based architectures outperforming Recurrent Neural Network (RNN) and Long




Figure 2: General workflow of a Generative Adversarial Network illustrating generator–discriminator interaction.

Short-Term Memory (LSTM) variants in capturing complex dependencies. These findings suggest that GAN-based models, particularly TimeGAN, are better suited to volatile, high-dimensional time-series data, whereas VAEs are more stable but less flexible in the face of heterogeneity. This performance advantage comes at the cost of higher computational complexity, as GAN architectures require substantially greater training resources than VAEs.

Existing literature reveals a distinct trade-off between stability and dynamic fidelity in generative modelling. Recent research in supply chain analytics has increasingly focused on modelling the propagation of disruptions and systemic resilience. Berger et al. [14] model the propagation of disruption using an epidemic-style network framework, showing that supply-chain failures can spread across interconnected supplier networks and that network topology plays a critical role in determining systemic risk. Complementary studies emphasise the role of digital technologies and organisational capabilities in enhancing supply chain resilience against global shocks such as the COVID-19 pandemic [15]. Furthermore, recent time-series analyses show that disruption intensity often exhibits strong temporal spikes and cross-sector dependencies, reinforcing the need for modelling frameworks capable of capturing nonlinear temporal dynamics in supply chain systems [16].

While VAEs provide the probabilistic regularisation necessary to smoothly reconstruct stable, low-variance datasets, they often fail to capture sharp regime shifts. Conversely, TimeGAN architectures excel at modelling the complex, non-linear volatility characteristic of market-driven environments but may introduce instability or amplified noise when applied to regulated, linear variables. Given that the Malaysian aviation supply chain is subject to both highly regulated fuel prices and volatile freight market indices, relying on a single architecture is insufficient. To address this heterogeneity, this study proposes a hybrid strategy that assigns specific generative tasks to the model best suited to them, leveraging VAEs for stable features and TimeGANs for dynamic indicators, to construct a robust and representative stress index.

3 Methodology

We proposed a hybrid TimeGAN–VAE–PCA framework designed to synthesise high-frequency supply chain data and construct a unified SCSi under conditions of scarcity. The methodological pipeline is structured into three distinct phases. First, the hybrid generative modelling phase partitions the dataset, employing a VAE

to reconstruct regulated price series and a TimeGAN to synthesise volatile market indicators. Second, the integration phase fuses these synthetic outputs with exogenous macroeconomic variables to create a comprehensive multivariate dataset. Finally, the index construction phase utilises Principal Component Analysis (PCA) to reduce dimensionality and extract the primary stress signal. The following subsections define the mathematical formulation of the hybrid objective function, the network architectures, and the validation protocols used to ensure distributional and temporal fidelity.

3.1 Proposed Hybrid TimeGAN–VAE formulation

Let $X = \{x_t\}_{t=1}^T$ denote multivariate aviation supply-chain indicators observed under sparse and fragmented reporting. TimeGAN models temporal dependencies through adversarial and supervised objectives, while a Variational Autoencoder (VAE) introduces probabilistic regularisation on the latent space. To balance temporal realism and distributional stability, this study adopts a hybrid loss as per equation (2):

$$\mathcal{L}_{\text{hybrid}} = \mathcal{L}_{\text{TimeGAN}} + \lambda D_{\text{KL}}(q_{\phi}(z|X) \parallel \mathcal{N}(0, I)) \quad (2)$$

where D_{KL} enforces smooth latent representations in the presence of data scarcity. Here, $\mathcal{L}_{\text{hybrid}}$ represents the joint optimisation objective that balances adversarial temporal learning with probabilistic latent regularisation. Meanwhile, $\lambda > 0$ controls the strength of probabilistic regularisation and balances temporal fidelity against latent smoothness. This formulation preserves nonlinear dynamics while mitigating excessive volatility or oversmoothing observed in standalone generative models. This hybrid objective does not introduce a new loss; instead, it combines established TimeGAN and VAE components within a unified optimisation framework. It jointly optimises adversarial learning for temporal realism and latent-space regularisation for distributional stability.

3.2 Dimensionality reduction for stress index construction

The synthetic multivariate time series generated by the hybrid TimeGAN–VAE framework is then processed to construct the stress index. Because the generative framework produces many synthetic

Table 1: Sources of aviation supply-chain indicators used to construct the dataset.

Source	Description	Reference
1	Monthly External Trade Statistics	Department of Statistics Malaysia [18, 19]
2	Manufacturing PMI (Malaysia)	S&P Global PMI [20]
3	Air Cargo Volume	Malaysian Ministry of Transport [21]
4	Malaysia Diesel Price	Government of Malaysia Open Data [22]
5	Jet Fuel Price (Asia and Oceania)	IATA Fuel Monitor [23]
6	Freight Air Index	Freightos Terminal Market Intelligence [24]

observations, preprocessing is required prior to downstream modelling. Mesran et al. [10] performed Principal Component Analysis (PCA) to generate the Malaysia Supply Chain Pressure Indicator (MSCPI), a stress indicator inspired by the Global Supply Chain Stress Indicator (GSCPI), a global-level supply chain pressure indicator [17], to understand the weak links in the supply chain network configuration in Malaysia. Principal Component Analysis (PCA) is used to transform the generated multivariate dataset into orthogonal components that capture the dominant variance structure across supply-chain indicators. The first principal component represents the dominant stress signal and is interpreted as the SCSi. This is a well-established method for generating a stress index, as it retains data features while presenting a single data set for easier interpretation. As for the study itself, it did not specify the datasets used to determine the stress index, which poses a challenge in identifying quantitative datasets to generate the stress index for overall data display prior to modelling. This is important, as this study is mainly exploratory and does not know which dataset is correct until the indicator indicates otherwise.

4 EXPERIMENTAL SETUP

All experiments were conducted in Python on Google Colab using an NVIDIA T4 GPU with 16 GB RAM. The core software stack comprised TensorFlow 2.15, PyTorch Lightning 2.2.5, scikit-learn 1.4, pandas, and seaborn. Random seeds were fixed across NumPy, PyTorch, and TensorFlow to ensure reproducibility, while GPU memory growth was enabled to prevent fragmentation. A unified configuration file defined the experiment metadata, including data directories, sequence length, batch size, learning rate, and latent dimension parameters.

The dataset consists of monthly aviation supply-chain indicators covering January 2017 to December 2024. Regulated variables include Malaysian fuel prices (RON95, RON97, diesel, and East-Malaysia diesel), while market-driven indicators include Asia-Pacific freight indices, the Freightos Air Index (weekly and annual), air cargo volume, and jet fuel prices. Two exogenous macroeconomic indicators, the Global Supply Chain Pressure Index (GSCPI) and the Geopolitical Risk Index (GPR), were incorporated to capture global disruption signals. All series were aligned to a common monthly timeline and standardised prior to modelling. Data were collected from publicly available sources, including Malaysian government energy reports, international freight indices, and global economic monitoring databases. Table 1 summarises the primary sources used to construct the aviation supply-chain dataset. Missing observations arising from inconsistent reporting frequencies were

addressed using time-aware interpolation followed by K-nearest-neighbour (KNN) imputation.

Based on this dataset, three experimental configurations were designed to evaluate the effectiveness of generative modelling under different data characteristics and modelling objectives. Figure 3 presents the overall architecture of the proposed hybrid TimeGAN-VAE pipeline.

Each experiment corresponds to a dedicated implementation notebook: VAE-only, TimeGAN-only, and a hybrid TimeGAN-VAE-PCA pipeline. The first experiment focuses on reconstructing stable, low-variance, regulated data using a probabilistic autoencoder to establish a baseline for reconstruction fidelity. The second applies a recurrent adversarial generator to capture temporal dynamics and volatility patterns in market-driven features. The third integrates both approaches, combining regulated and market features with exogenous indicators to derive a unified SCSi through principal component analysis.

Table 2 summarises the architectural design, objectives, data domains, optimisation components, output artefacts, and evaluation emphasis for each experiment, highlighting the methodological differences and computational trade-offs across the three workflows. All datasets were resampled to a monthly frequency and subjected to a two-stage cleaning process comprising time-aware interpolation followed by K-nearest-neighbour (KNN) imputation. Seasonality was encoded using sine-cosine month representations and binary indicators of pandemic-era policies. Features were then standardised ($\mu = 0$, $\sigma = 1$) and split chronologically into a training window (2017–2022) and a validation window (2023–2024).

4.1 Model Training

The first experiment employed a VAE to reconstruct low-variance regulated fuel-price series. Lagged inputs ($t-1$, $t-2$) were MinMax-scaled and passed through a 64–32 encoder, an eight-dimensional latent space, and a symmetric decoder. The model was trained for 50 epochs using Adam ($\text{lr} = 1 \times 10^{-3}$) with β -annealed KL regularisation ($\beta = 0.01$). Synthetic series were generated from latent samples and evaluated against observed data using distributional moments and error metrics (RMSE, MAE), before being exported for downstream integration.

The second experiment applied TimeGAN to volatile, market-driven indicators. Inputs were MinMax-scaled and segmented into overlapping 72-month sequences. The GRU-based architecture comprised an embedder, generator, supervisor, recovery, and discriminator modules, each with 24 hidden units. Training was conducted




Figure 3: Architecture of the proposed TimeGAN-VAE pipeline for constructing the Supply Chain Stress Index (SCSI). Market-driven variables are modelled using TimeGAN to capture temporal volatility, while regulated price series are reconstructed using a VAE to preserve smoothness and stability. Exogenous indicators are incorporated via pass-through, and fused synthetic features are reduced using PCA to obtain a single interpretable stress index.

Table 2: Experiment and workflow summary

Description	VAE Only	TimeGAN Only	TimeGAN-VAE
Model Architecture	Feed-Forward Probabilistic Autoencoder	Recurrent Adversarial Network (GRU-based)	Composite (VAE + TimeGAN + PCA reduction)
Primary Goal	Reconstruct low-variance regulated data and test smooth generative capability	Generate realistic high-variance market data with temporal dependence	Combine synthetic regulated & market data to construct interpretable SCSi index
Data Domain	Regulated prices (RON95/97, diesel, East Malaysia diesel)	Market indicators (Asia-Pacific indices, Freightos weekly/annual, cargo volume)	Integrated dataset + exogenous (GSCPI, GPR)
Loss Components	Reconstruction + β -annealed KL divergence ($\beta = 0.01$)	Adversarial + Supervised + Reconstruction + Temporal smoothness	Weighted hybrid (TimeGAN adversarial + supervised + recon + VAE KL) $\rightarrow$ PCA 1st component (SCSi)
Evaluation Focus	Statistical fidelity such as RMSE, MAE, skewness, kurtosis comparison	Temporal realism such as ACF/PACF, spectral density, volatility regimes	Interpretability & systemic stress coherence, correlation vs GSCPI, smoothness vs volatility

over 500 epochs using Adam ($\text{lr} = 3 \times 10^{-4}$), optimising adversarial, supervised, and reconstruction losses. Generated trajectories were rescaled and assessed using MAE, MSE, autocorrelation, and spectral density measures.

The third experiment implemented a hybrid TimeGAN–VAE–PCA pipeline that combines synthetic regulated and market features into a unified representation. TimeGAN was trained on log-returns to capture volatility and enforce stationarity, while the VAE reconstructed lagged regulated series. The synthetic outputs were merged with exogenous indicators (GSCPI, GPR), standardised, and reduced via PCA. The first principal component, the SCSI, is sign-oriented, with higher values indicating greater stress. The pipeline outputs include synthetic feature datasets, runtime diagnostics, and SCSI visualisations; key algorithmic steps and code excerpts are provided in Appendices A1 and A2.

4.2 Hyperparameter Configuration

The three experimental configurations differ in architecture, input domain, and training objective. The VAE-only model employs a compact feed-forward encoder (64–32) with an eight-dimensional latent space to reconstruct low-variance regulated price series, optimised using β -annealed KL regularisation ($\beta = 0.01$) and Adam ($\text{lr} = 1 \times 10^{-3}$). In contrast, the TimeGAN-only configuration uses a GRU-based adversarial framework trained on longer temporal sequences (72 months) to capture volatility and cyclical behaviour in market-driven indicators. Its objective combines adversarial, supervised, reconstruction, and temporal-consistency losses, optimised with Adam ($\text{lr} = 3 \times 10^{-4}$) over 500 epochs. The hybrid pipeline integrates both models, which are VAE for stable regulated variables and TimeGAN for dynamic market features followed by PCA to construct a single interpretable SCSI. To improve computational efficiency and alignment, the hybrid setup uses a shorter sequence length (24 months) and a unified configuration file. Runtime comparisons on an NVIDIA T4 GPU show increasing complexity from VAE-only (≈ 15 min) to TimeGAN-only (≈ 90 min), with the hybrid model achieving a balanced trade-off (≈ 45 min) between performance and interpretability. Evaluation metrics were standardised across experiments to enable consistent comparison of reconstruction fidelity, temporal realism, and systemic coherence.

Hyperparameters were tuned to ensure stable convergence and reproducibility under data constraints. Initial settings were derived from pilot experiments and refined using small grid adjustments. For the VAE, latent dimensions above 16 led to overfitting; a latent size of 8 with β -annealing yielded stable convergence and smooth reconstruction. For TimeGAN, careful tuning of the adversarial–supervised balance was necessary to avoid mode collapse, with the final configuration (hidden = 24, $\text{lr} = 3 \times 10^{-4}$, batch = 8) providing the best stability–fidelity trade-off. The hybrid model inherits these optimised settings, with harmonised learning rates and reduced sequence length to prevent training imbalance. Early stopping and loss monitoring were used to terminate runs when improvement plateaued, and all experiments were repeated with fixed random seeds to verify robustness.

Table 3 summarises the final hyperparameter configurations and training settings for the three experimental pipelines. The VAE-only configuration employs a compact latent representation to stabilise

reconstruction of regulated price series, while the TimeGAN-only setup uses a GRU-based architecture to capture temporal dependencies in volatile market indicators. The hybrid configuration integrates both components and applies PCA for dimensionality reduction, enabling the construction of a unified SCSI.

4.3 Evaluation Strategy

Models were evaluated along (i) distributional fidelity (RMSE, MAE, skewness, kurtosis), (ii) temporal realism (ACF/PACF similarity, spectral distance, DTW distance), and (iii) hybrid index coherence via PCA variance explained, correlation with exogenous indicators, and event alignment during the 2020–2021 pandemic period. All metrics are computed on the validation window (2023–2024) with identical preprocessing and scaling. Runtimes are reported on an NVIDIA T4 GPU.

5 RESULTS

This section evaluates the ability of the proposed models to reproduce statistical and temporal properties of the observed aviation supply-chain indicators.

5.1 Distributional Fidelity

Across regulated fuel-price variables, the VAE achieved the lowest reconstruction errors in both RMSE and MAE metrics (Tables 4 and 5). For RON95, RMSE = 0.010 (MAE = 0.008), outperforming TimeGAN by 33% RMSE (0.015) and the Hybrid by 17% (0.012). Diesel shows the same pattern (VAE RMSE 0.012 vs. TG 0.017, Hybrid 0.013). On volatile market features the ordering reverses: TimeGAN reduces RMSE on Asia-Pacific by 43% vs. VAE (0.020 vs. 0.035) and by 11% vs. Hybrid (0.020 vs. 0.022); for weekly airfreight, TimeGAN improves RMSE by 28% vs. VAE (0.018 vs. 0.025) and by 10% vs. Hybrid (0.018 vs. 0.020). Overall, the hybrid model preserves the VAE’s low reconstruction error on regulated variables while narrowing the performance gap with TimeGAN on volatile market indicators. This confirms the effectiveness of the proposed division-of-labour strategy, where VAE stabilises low-variance series and TimeGAN captures dynamic temporal behaviour, resulting in a balanced generative framework for heterogeneous supply-chain data.

While Table 4 reports the root-mean-square error (RMSE) across models, Table 5 presents the corresponding mean absolute error (MAE) to provide an additional measure of reconstruction fidelity.

The MAE results in Table 5 confirm the same trend observed in RMSE, with VAE providing the most accurate reconstruction for stable regulated variables, while TimeGAN better captures volatility in market-driven indicators.

5.2 Temporal Realism

The TimeGAN model demonstrated a substantial improvement in capturing temporal dependence and dynamic structure compared to the VAE baseline. As shown in Table 6, TimeGAN substantially improves the preservation of temporal structure compared with the VAE baseline. The autocorrelation (ACF) similarity increased from 0.60 to 0.80 (+33 %), while partial autocorrelation (PACF) similarity improved from 0.55 to 0.75, indicating better preservation of lagged relationships across monthly fuel-price and freight indices.

Table 3: Hyperparameters and training configurations for the VAE-only, TimeGAN-only, and hybrid TimeGAN-VAE experiments.

Parameter	VAE Only	TimeGAN Only	TimeGAN-VAE
Data range	Jan 2017-Dec 2024	Jan 2017-Dec 2024	Jan 2017-Dec 2024
Target features	RON95, RON97, Diesel, East Malaysia Diesel (regulated prices)	Air Freight Rate (Weekly/Annual), Asia-Pacific Indices, Air Cargo Volume, Jet Fuel, Oil Price, Trade Volume	All regulated + market features + exogenous (GSCPI, GPR)
Train Validation Split	80 / 20 (random)	Chronological (2017–2022 / 2023–2024)	Chronological (2017–2022 / 2023–2024)
Scaling Method	MinMax Scaler	MinMax Scaler	Standard Scaler ($\mu = 0, \sigma = 1$)
Hidden / Latent Dimension	Encoder 64–32 units; Latent = 8	GRU hidden = 24	VAE latent 8 + TG hidden 24
Optimizer	Adam ($\text{lr} = 1 \times 10^{-3}$)	Adam ($\text{lr} = 3 \times 10^{-4}$)	Adam (per component)
Epochs	50	500	VAE 50 TimeGAN 500
Regularisation / Loss	Reconstruction + β -KL ($\beta = 0.01$)	Adversarial + Supervised + Reconstruction + Temporal	Weighted VAE KL + TimeGAN losses + PCA projection
Learning Objective	Minimise reconstruction error	Balance adversarial and temporal losses	Integrate fidelity + interpretability (PCA)
Approx. Runtime (on T4 GPU)	≈ 15 min	≈ 90 min	≈ 45 min

Table 4: Distributional error metrics (RMSE) for regulated and market indicators.

Feature	VAE Only	TimeGAN Only	TimeGAN-VAE
RON95	0.010	0.015	0.012
Diesel	0.012	0.017	0.013
Asia Pacific Fuel Price	0.035	0.020	0.022
Air Freight	0.025	0.018	0.020

Table 5: Distributional error metrics (MAE) for regulated and market indicators.

Feature	VAE Only	TimeGAN Only	TimeGAN-VAE
RON95	0.008	0.013	0.009
Diesel	0.009	0.014	0.010
Asia Pacific Fuel Price	0.030	0.015	0.017
Air Freight	0.020	0.014	0.015

The hybrid TimeGAN-VAE configuration further improved these metrics to 0.85 (ACF) and 0.80 (PACF), representing overall gains of +42 % and +6% relative to the VAE and TimeGAN models, respectively. Likewise, the power spectral density (PSD) difference decreased from 0.25 (VAE) to 0.15 (TimeGAN) and 0.12 (Hybrid), confirming improved frequency-domain alignment and smoother reconstruction of cyclical variations. The dynamic time warping (DTW) distance decreased from 0.30 to 0.20 and 0.18 across the three configurations, indicating greater pathwise similarity between the synthetic and real series. Collectively, these results verify that TimeGAN effectively models regime-dependent volatility, while the hybrid model refines both temporal coherence and amplitude stability across heterogeneous aviation-supply features.


These results confirm that TimeGAN effectively captures temporal dependencies in volatile market indicators, while the hybrid architecture improves both temporal coherence and distributional stability across heterogeneous variables.

5.3 VAE Only

The VAE exhibited stable and efficient convergence in Figure 4 (a), with the reconstruction loss decreasing rapidly from ~ 0.145 to ~ 0.02 within the first 10 epochs and stabilising near 0.018 by epoch 50. This monotonic behaviour indicates effective optimisation without posterior collapse, supported by β -annealed KL regularisation ($\beta = 0.01$). As shown in Figure 4 (b), the generated RON95 fuel price distribution closely matches the observed mean and variance, with

Table 6: Temporal similarity metrics comparing the ability of VAE, TimeGAN, and the hybrid TimeGAN–VAE model to reproduce temporal dependencies in market-driven supply-chain indicators.

Metric	VAE Only	TimeGAN Only	TimeGAN–VAE
ACF similarity (higher is better)	0.60	0.80	0.85
PACF similarity (higher is better)	0.55	0.75	0.80
PSD difference (lower is better)	0.25	0.15	0.12
DTW distance (lower is better)	0.30	0.20	0.18

**Figure 4: Training loss curve of the Variational Autoencoder (VAE) over 50 epochs (4 (a)). The rapid decline in loss within the first 10 epochs indicates stable convergence and effective reconstruction of regulated fuel-price data. Comparison of real and generated RON95 fuel-price distributions using the VAE model (4(b)). The histogram shows close alignment in central tendency and variance, confirming the VAE’s capacity to replicate low-variance, regulated series with minimal distortion.**

slightly reduced tail mass. This mild smoothing reflects the VAE’s probabilistic averaging and is desirable for low-variance regulated series, reinforcing its role in the hybrid framework as a stability-preserving component.

5.4 TimeGAN Only

The TimeGAN model successfully captured temporal dependencies and regime transitions in both regulated and market-driven variables. As shown in Figure 5 – Figure 8, the synthetic sequences (dashed lines) tracked the overall trends of the actual data while reproducing characteristic volatility bursts observed around 2020–2021 and 2023–2024. For the regulated fuel-price series, the generated trajectories maintain consistent directionality and amplitude with mild lagging effects. In particular, the synthetic outputs closely follow the observed dynamics for RON97 and diesel, as illustrated in Figures 6 and 7, respectively, resulting in RMSE values between 0.015 and 0.017. These small deviations arise from TimeGAN’s adversarial dynamics, which prioritise time-order realism over pointwise precision. In contrast, for the Asia-Pacific market indicator (Figure 8), the synthetic output accurately reflects both the magnitude and frequency of observed fluctuations, including the sharp 2024 peak, demonstrating the model’s capacity to generalise under volatile conditions.

Quantitatively, the TimeGAN achieved ACF and PACF similarities of 0.80 and 0.75, respectively, outperforming the VAE baseline (0.60/0.55) and reducing spectral-density differences by 40%

(0.25 → 0.15). These results confirm that TimeGAN better preserves sequential and cyclical structures essential for modelling dynamic supply-chain indicators. Although the model slightly over-smoothed post-shock recoveries, its overall alignment with real temporal patterns validates its role as the volatility-sensitive component within the hybrid framework.

5.5 Hybrid Model and SCSi Construction

To leverage the complementary strengths identified in the preceding analyses, we integrate TimeGAN and a β -annealed VAE into a unified synthetic-data pipeline. TimeGAN is applied to market-driven, shock-sensitive indicators, where preserving temporal dependence and regime shifts is essential, whereas VAE models are applied to series that require stable, low-noise reconstruction. This decomposition separates volatility-sensitive from volatility-suppressed components, which are generated using the most appropriate model and recombined into a unified feature matrix for PCA-based SCSi construction. The hybrid approach preserves the timing of disruptions while avoiding over-amplification of low-variance variables, resulting in a stress index that is both responsive and stable under data scarcity.

The resulting SCSi shown in Figure 9, captures pronounced structural volatility over the study period. The index remains relatively stable prior to 2020, surges sharply during the COVID-19 pandemic (2020–2021), and transitions into a sustained low-stress regime




Figure 5: Actual vs. synthetic RON95 fuel-price series generated using TimeGAN. The generated sequence reproduces the overall trend and short-term oscillations of the real series while maintaining stable amplitude across the observation period.




Figure 6: Actual vs. synthetic RON97 fuel-price trajectories generated using TimeGAN. The model captures volatility spikes observed during the 2021–2022 period, demonstrating its ability to model dynamic fluctuations in market-driven fuel prices.




Figure 7: Actual vs. synthetic diesel price series generated using TimeGAN. The synthetic curve follows the overall trend of the real data while smoothing abrupt shocks, reflecting the adversarial model’s emphasis on temporal realism rather than pointwise reconstruction accuracy.

from 2022 onwards, reflecting post-pandemic adjustments and easing constraints. The distribution of stress levels indicates that the aviation supply chain rarely operates near equilibrium, instead oscillating between extended high- and low-stress regimes, highlighting persistent instability rather than normalised conditions.

To provide a clearer depiction of these dynamics, Figure 10a overlays the COVID-19 period on the SCSI timeline, while Figure 10b presents the distribution of stress levels across the study period. Nearly half of all observations (48.6 %) fall within a low-stress

regime, and 43.1 % correspond to high-stress conditions, while only 2.2 % and 4.2 % fall into the “normal” and “elevated” categories, respectively. This bimodal distribution indicates that the Malaysian aviation supply chain rarely operated near equilibrium, instead oscillating between extremes of severe strain and pronounced relaxation. Such volatility underscores the dynamic and nonlinear nature of modern supply networks, which are strongly influenced by external shocks and recovery cycles.




Figure 8: Actual vs synthetic Asia-Pacific freight index. The synthetic sequence aligns with the timing of major disruptions (2020–2021) but exhibits moderate amplitude, indicating partial volatility compression typical of adversarial training.




Figure 9: SCSI generated using the hybrid TimeGAN–VAE–PCA framework, showing the stress surge during the COVID-19 disruption period.

Constructing a reliable SCSI requires integrating multiple heterogeneous indicators covering both demand- and supply-side dynamics such as fuel prices, freight indices, and operational bottlenecks. The hybrid TimeGAN–VAE–PCA approach proved effective in addressing data sparsity, reconstructing missing values, and generating continuous high-frequency representations from incomplete datasets. This method preserves both temporal dependence and statistical coherence, resulting in interpretable stress patterns that align with real-world events. Training both TimeGAN and VAE models remains computationally demanding, as each must be iteratively optimized to ensure convergence and consistency across varying data regimes. Despite these computational challenges, the proposed hybrid generative pipeline provides a robust framework for quantifying systemic stress under limited, unevenly distributed data.

6 CONCLUSIONS

This paper proposed a hybrid TimeGAN–VAE–PCA framework to synthesise high-frequency, regime-sensitive indicators from sparse Malaysian aviation supply-chain data and to construct an interpretable Supply Chain Stress Index (SCSI). By combining the temporal realism of TimeGAN with the distributional robustness of

the VAE and the dimensional interpretability of PCA, the framework captured nonlinear transitions and structural shifts in systemic stress. Empirically, the hybrid achieved lower DTW distance and comparable reconstruction accuracy while running at ~50% of the TimeGAN-only cost; distributional errors remained low (e.g., RMSE ≈ 0.012) and temporal similarity improved markedly (ACF +42%). The resulting SCSI mirrored real-world disruption patterns, including the 2020–2021 pandemic surge and subsequent recovery, demonstrating practical utility for quantitative stress monitoring in data-scarce environments. Current limitations include computational demand for dual-model training, sensitivity to hyperparameters, and potential domain shift when indicators change.

Overall, the proposed hybrid framework provides a scalable and extensible foundation for supply-chain analytics, forecasting, and policy evaluation under limited and fragmented data conditions. The framework also offers a practical analytical tool for policymakers and aviation stakeholders to monitor systemic supply-chain stress in environments where reliable high-frequency data are scarce. Future work will explore lightweight retraining and transfer learning, online updating mechanisms, and uncertainty quantification (e.g., Bayesian calibration of the index), alongside ablation




Figure 10: SCSI with highlighted COVID-19 period (7(a)). The shaded region (2020–2021) corresponds to the high-stress phase, with index values exceeding +5, followed by a sustained recovery as constraints eased and logistics stabilised. Stress-level distribution of the SCSI (7(b)). Nearly half of all observations (48.6%) fall in the low-stress regime, 43.1% in the high-stress regime, and only 6.4% in normal or elevated states, confirming a bimodal pattern of systemic volatility in Malaysian aviation supply chains.

studies and external validation against independent operational metrics.

Acknowledgments

The authors acknowledged the financial support from the Ministry of Higher Education Malaysia through the Fundamental Research Grant Scheme (FRGS) FRGS/1/2023/ICT06/UNIMAS/02/2.

References

- [1] C. Levy, J. Katz, C. Blaufelder, D. Pinner, and J. Weterings. 2021. How the voluntary carbon market can help address climate change. In *McKinsey on Risk: New Risk Challenges and Enduring Themes for the Return*. McKinsey Global Publications, Vol. 10 (Jan. 2021), 72–79.
- [2] J. Henrich, J. D. Li, C. Mazuera, and F. Perez. 2022. *Future-Proofing the Supply Chain*. McKinsey Global Publishing (Jun. 2022).
- [3] K. Alicke, T. Foster, and V. Trautwein. 2024. *Supply Chains: Still Vulnerable*. McKinsey Global Publishing (Oct. 2024).
- [4] F. Zakir, D. Wang, A. Rehman, and A. Waheed. 2023. Supply Chain Process Improvement for International Airline Industry. *Journal of Engineering, Project, and Production Management* 13, 1 (Jan. 2023), 10–19. <https://doi.org/10.32738/JEPPM-2023-0002>
- [5] J. Moosavi, A. M. Fathollahi-Fard, and M. A. Dulebenets. 2022. Supply chain disruption during the COVID-19 pandemic: Recognizing potential disruption management strategies. *International Journal of Disaster Risk Reduction* (Jun. 2022). <https://doi.org/10.1016/j.ijdrr.2022.102983>
- [6] E. Faia and S. Karau. 2021. Systemic Bank Risk and Monetary Policy. *International Journal of Central Banking* 17, 5 (Dec. 2021), 1–40.
- [7] P. Sharma, M. Kumar, H. K. Sharma, and S. M. Biju. 2024. Generative adversarial networks (GANs): Introduction, taxonomy, variants, limitations, and applications. *Multimedia Tools and Applications* 83, 41 (Dec. 2024), 88811–88858. <https://doi.org/10.1007/s11042-024-18767-y>
- [8] D. P. Kingma and M. Welling. 2021. An Introduction to Variational Autoencoders. *Foundations and Trends in Machine Learning* 12, 4 (2021), 307–392. <https://doi.org/10.1561/22000000056>
- [9] C. Jamotton and D. Hainaut. 2024. Variational AutoEncoder for synthetic insurance data. *Intelligent Systems with Applications* 24 (Dec. 2024). <https://doi.org/10.1016/j.iswa.2024.200455>
- [10] M. Kyizzad Mesran, A. A. S. Mohamad, and R. H. Zakaria. 2024. Assessing Malaysia Supply Chain Vulnerabilities: Construction of Comprehensive Stress Index. *International Journal of Business and Technology Management* 6, S4 (2024), 23–48.
- [11] K. Abberger, M. Graff, O. Muller, and B. Siliverstovs. 2023. Imputing Monthly Values for Quarterly Time Series: An Application Performed with Swiss Business Cycle Data. *Journal of Business Cycle Research* 19, 3 (Nov. 2023), 241–273. <https://doi.org/10.1007/s41549-023-00088-y>
- [12] S. Echebarri, P. Do, H.-C. Vu, and P.-Y. Liegeois. 2025. A modified TimeGAN-based data augmentation approach for the state of health prediction of lithium-ion batteries. *Reliability Engineering & System Safety* 264, Part A (Dec. 2025), Article 111297. <https://doi.org/10.1016/j.res.2025.111297>
- [13] C. Y. Tai, W. J. Wang, and Y. M. Huang. 2023. Using Time-Series Generative Adversarial Networks to Synthesize Sensing Data for Pest Incidence Forecasting on Sustainable Agriculture. *Sustainability (Switzerland)* 15, 10 (May 2023). <https://doi.org/10.3390/su15107834>
- [14] N. Berger, S. Schulze-Schwering, E. Long, and S. Spinler. 2023. Risk management of supply chain disruptions: An epidemic modeling approach. *European Journal of Operational Research* 304, 3 (Feb. 2023), 1036–1051. <https://doi.org/10.1016/j.ejor.2022.05.018>
- [15] R. Pellegrino, M. Costantino, and F. Tauro. 2025. Supply chain resilience as a viable way to cope with disruptions. *Annals of Operations Research* (2025). <https://doi.org/10.1007/s11846-025-00911-z>
- [16] F. Sufi, A. Rahman, and M. Ahmed. 2025. Modeling disruption dynamics and resilience in global supply chains. *Mathematics* 13, 21 (2025), Article 3471. <https://doi.org/10.3390/math13213471>
- [17] Federal Reserve Bank of New York. 2025. *Global Supply Chain Pressure Index*. Retrieved 2 September 2025 from <https://www.newyorkfed.org/research/policy/gscpi>
- [18] Department of Statistics Malaysia. 2024. *Monthly External Trade Statistics: December 2024*. Retrieved 21 November 2025 from <https://www.dosm.gov.my/portal-main/release-archive/monthly-external-trade-statistics-december-2024>
- [19] Department of Statistics Malaysia. 2025. *Monthly External Trade Statistics: March 2025*. Retrieved 21 November 2025 from <http://www.dosm.gov.my/portal-main/release-content/monthly-external-trade-statistics-march-2025>
- [20] S&P Global. 2025. *Purchasing Managers' Index (PMI) Press Releases*. Retrieved 21 November 2025 from [https://www.pmi.spglobal.com/Public/Release/PressReleases?language=\\$en](https://www.pmi.spglobal.com/Public/Release/PressReleases?language=$en)
- [21] Ministry of Transport Malaysia. 2025. *Aviation Transport Statistics Reports*. Retrieved 21 November 2025 from <https://www.mot.gov.my/en/Pages/Aviation/Reports/statistics-aviation-transport.aspx>
- [22] Government of Malaysia. 2025. *Fuel Price Dataset*. Retrieved 21 November 2025 from <https://data.gov.my/data-catalogue/fuelprice>
- [23] International Air Transport Association (IATA). 2025. *Jet Fuel Price Monitor*. Retrieved 21 November 2025 from <https://www.iata.org/en/publications/economics/fuel-monitor/>
- [24] Freightos. 2025. *Freightos Terminal Global Air and Ocean Market Intelligence*. Retrieved 21 November 2025 from <https://app.terminal.freightos.com/fax>

A APPENDICES

A.1 Key Fusion Steps Algorithm

Algorithm 1 Ultra-Compact Hybrid TimeGAN–VAE–PCA Framework

Input:

X_{reg} regulated indicators
X_{mkt} market-driven indicators
X_{exo} exogenous indicators

Output:

SCSI Supply Chain Stress Index

Initialize:

align all series to monthly timeline
standardize inputs
set random_seed

while convergence criteria not met do

/* VAE for regulated indicators */
train β -annealed VAE on X_{reg}
generate synthetic regulated series X_{reg}
/* TimeGAN for market indicators */
train TimeGAN on X_{mkt}
generate synthetic market series X_{mkt}
/* Hybrid integration */
align X_{reg}, X_{mkt}, and X_{exo}
concatenate and standardize hybrid matrix X_{hyb}
/* Stress index construction */
apply PCA to X_{hyb}
extract first principal component PC1
orient PC1 using correlation with GSCPI
assign PC1 as SCSI
/* Validation */
assess fidelity, temporal realism, and event alignment

end

return SCSI

A.2 Hybrid TimeGAN–VAE–PCA Iterative Procedure

INPUT:

X\reg regulated indicators
X\mkt market-driven indicators
X\exo exogenous indicators

OUTPUT:

SCSI Supply Chain Stress Index

INITIALISE:

align all series to monthly timeline
standardise all inputs
set random seed

WHILE convergence criterion not met DO

TRAIN β -annealed VAE on X\reg
GENERATE synthetic regulated series X~\reg
TRAIN TimeGAN on X\mkt
GENERATE synthetic market series X~\mkt
ALIGN X~\reg, X~\mkt, and X\exo
CONCATENATE into hybrid matrix X\hyb
STANDARDISE X\hyb

```
    APPLY PCA to X\hyb
    EXTRACT first principal component PC1
    ORIENT PC1 using correlation with GSCPI
    ASSIGN PC1 as SCSI
    VALIDATE distributional fidelity
    and temporal realism
END
RETURN SCSI
```