Optimizing Emergency Department Patient Flow Forecasting: A Hybrid VAE-GRU Model

Amel ZIDI[1], Rayen JMILI[2], Issam NOUAOURI[3], Ines BEN JAAFAR[4]

***Abstract*— Emergency departments (EDs) face increasing patient demand, leading to overcrowding and resource strain. Accurate forecasting of ED visits is critical for optimizing hospital operations and ensuring efficient resource allocation. This paper proposes a hybrid model combining Variational Autoencoder (VAE) and Gated Recurrent Unit (GRU) to enhance patient flow predictions. The VAE extracts meaningful latent features while handling missing data, whereas the GRU captures complex temporal dependencies, improving forecasting accuracy. Compared to traditional models such as LSTM, GRU, and 1D CNN, our hybrid VAE-GRU model demonstrates superior predictive performance. Experimental results, based on real-world hospital data, highlight the model's effectiveness in reducing prediction errors and improving decision-making in dynamic ED environments. Additionally, we compare the proposed model with ARIMA-ML, emphasizing the trade-offs between computational efficiency and prediction accuracy. The findings suggest that hybrid deep learning approaches can significantly enhance healthcare resource management, reducing patient waiting times and improving overall hospital efficiency.**

I. INTRODUCTION

The growing influx of patients in hospital emergency departments (EDs) presents significant challenges for healthcare systems, particularly in managing limited resources, maintaining the quality of care, and ensuring operational efficiency. Accurately predicting patient flow is essential to support timely decision-making and deliver appropriate medical interventions, especially during crises such as epidemic outbreaks or seasonal surges [1][2]. These rising demands call for smarter, data-driven strategies to effectively manage scarce medical staff and hospital beds.

Traditional time-series models like ARIMA and other autoregressive approaches [3][4] can perform reasonably well in stable environments, but they often fall short when it comes to capturing the nonlinear and highly dynamic nature of patient arrivals. In contrast, artificial intelligence (AI) and deep learning (DL) methods [9][10][11] provide more flexible and powerful modeling capabilities. They are better equipped to uncover complex patterns and adapt to the rapidly changing conditions that healthcare institutions frequently face [2][3].

*This work was not supported by any organization

[1]Amel Zidi is with The Higher Institute of Management of Tunis, Tunisia amel.zidi@esc.u-manouba.tn

[2]Rayen Jmili is with Higher School of Commerce of Tunis - University of Manouba, Tunisia jrayen926@gmail.com

[3]Issam Nouaouri is with Artois University, UR 3926, LGI2A Artois University 62400 Bethune, France issam.nouaouri@univ-artois.fr

[4]Ines Ben Jaafar is with Higher School of Commerce of Tunis - University of Manouba, Tunisia ines.benjaafar@esct.uma.tn

Despite their promise, many AI models still face key limitations such as handling missing data, capturing long-term dependencies, and providing interpretable results. This study seeks to address these gaps by introducing a hybrid forecasting model that combines Variational Autoencoder (VAE) and Gated Recurrent Unit (GRU) architectures. The VAE component supports robust learning of latent representations while effectively dealing with incomplete data, and the GRU module excels at modeling sequential dependencies over time.

This research is guided by the following core questions:

Can combining VAE and GRU architectures significantly enhance the accuracy of ED patient flow forecasting compared to traditional and other deep learning models?

What is the added value of using VAE in healthcare time-series forecasting, despite its higher computational complexity?

How well does the proposed model perform under real-world hospital conditions, particularly during periods of high patient influx?

To explore these questions, the model is trained and evaluated using real ED datasets and benchmarked against established forecasting approaches, including LSTM, CNN, and ARIMA-ML. Results show that the proposed model improves predictive accuracy and reduces error margins, offering valuable operational insights for hospital administrators, particularly in staff scheduling and emergency preparedness.

This work contributes to the ongoing advancement of AI-enabled healthcare solutions by offering a novel, empirically validated predictive model for resource management in emergency care settings.

The study is organized into four sections. Section I introduces the challenges in emergency departments and current forecasting issues. Section II explains the methodology, including the VAE component and GRU model. Section III covers data analysis and results. Section IV discusses advantages, limitations, and future implications.

II. REVIEW OF FORECASTING MODELS FOR HOSPITAL EMERGENCY DEPARTMENT FLOWS IN CRISIS SITUATIONS

The prediction of patient flows in hospital systems, particularly emergency department (ED) visits, is a critical research area. Forecasting methods are broadly categorized into classical statistical models and AI-based approaches, including machine learning (ML) and deep learning.

Traditional time series models such as ARMA, ARIMA, and VARMA are widely used for short-term forecasting due to their statistical rigor. For example, an ARMA model predicted annual ED visits at Troyes Hospital with 91.2% accuracy, notably during epidemics due to low residual correlation with epidemiological trends [3]. At Lille University Hospital, ARIMA models achieved a mean error of three patients per day in pediatric ED admissions [1], while Poisson regression has been applied for rare event modeling [9].

However, these classical models face limitations in capturing nonlinear and high-dimensional interactions. AI methods, such as random forests and gradient boosting, have demonstrated improved performance. For instance, random forests outperformed GBM and ARIMA in predicting daily ED visits at St. Joseph Mercy Ann Arbor [16], particularly on large heterogeneous datasets.

Deep learning models, especially recurrent neural networks (RNNs) like LSTM and GRU, are well suited for temporal ED data due to their ability to model sequential dependencies. A GRU-based model provided strong hourly forecasts in Lille's pediatric ED [16], while stacked LSTM-GRU networks enhanced accuracy in high-variability settings such as Singapore General Hospital [10].

Despite the power of AI models, classical approaches like ARIMA/SARIMA remain relevant for linear trends and seasonality [5], motivating hybrid methods combining statistical and AI techniques to leverage their complementary strengths.

Variational Autoencoders (VAEs) represent a recent innovation capable of handling missing data, reducing dimensionality, and capturing latent structures in noisy medical datasets, thus improving robustness in irregular time series forecasting.

The integration of VAEs with GRUs forms a hybrid model that combines representational learning and temporal modeling, enhancing adaptability and forecasting accuracy for managing complex and uncertain hospital environments.

III. VARIATIONAL AUTOENCODER : VAE

The Variational Autoencoder (VAE), introduced by Shangsong Liang [13] and further developed by Cinelli [14], is a probabilistic generative model grounded in variational Bayesian inference [20]. It belongs to the class of latent variable models and approximates complex data distributions through a learned lower-dimensional latent space. An encoder maps input x to a latent vector z, while a decoder reconstructs x from samples drawn from the learned distribution, as illustrated in Figure 1. This framework captures highly nonlinear dependencies in data, making it particularly effective in domains characterized by high-dimensional complexity.

Despite increased computational demands, VAEs outperform simpler models in dynamic contexts such as healthcare, where ED visit patterns are shaped by nonlinear factors including epidemics, mobility, staffing fluctuations, and public health policies [8][12]. Unlike traditional methods, VAEs can adapt to distributional shifts and capture latent interactions, yielding more accurate and resilient forecasts.

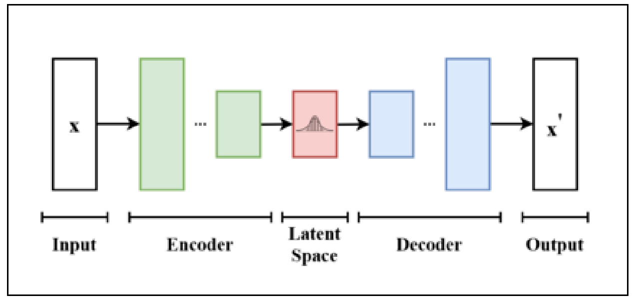

Fig. 1. Basic scheme of a VAE [7]

A. *Probabilistic Modeling in VAE*

From a probabilistic perspective, the objective is to maximize the likelihood of the data x under a chosen distribution $p_\theta(x) = p(x|\theta)$, often assumed Gaussian $N(x|\mu, \sigma)$, where μ and σ are learnable parameters. While this is tractable for simple distributions, introducing a prior over latent variables z renders the exact computation of the likelihood intractable. To address this, the marginal likelihood is expressed as shown in Equation (1).

Using the chain rule of probability, the joint distribution is decomposed as shown in Equation (2).

In the context of the VAE, the latent variable $z \in R^d$ is typically modeled as a continuous vector, and the conditional distribution $p_\theta(x \mid z)$ is assumed Gaussian, making $p_\theta(x)$ a mixture of Gaussians. The generative process is defined by a prior, a likelihood, and a posterior.

However, the posterior $p_\theta(z \mid x)$ is generally intractable due to its complex form. Therefore, it is approximated by a variational distribution, as shown in Equation (3).

$$p_\theta(x) = \int_z p_\theta(x, z)\, dz1 \tag{1}$$

$$p_\theta(x) = \int_z p_\theta(x \mid z)\, p_\theta(z)\, dz2 \tag{2}$$

$$q_\phi(z \mid x) \approx p_\theta(z \mid x)3 \tag{3}$$

IV. GATED RECURRENT UNIT : GRU

The Gated Recurrent Unit (GRU), introduced by Cho et al. [15], is a streamlined variant of the Long Short-Term Memory (LSTM) model [19], featuring fewer parameters and a simpler architecture. By merging the forget and input gates of LSTM into a single "update gate," the GRU achieves comparable performance in sequential data tasks with improved computational efficiency and easier training. Figure 2 depicts its basic structure.

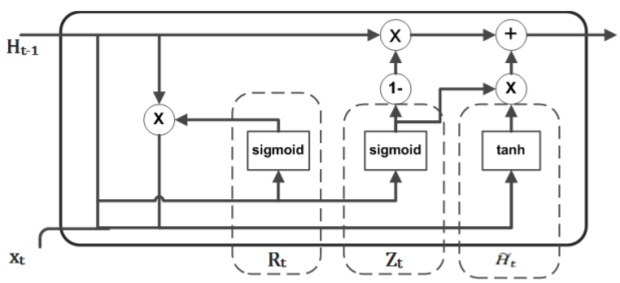

Fig. 2. Basic scheme of a GRU [6]

The GRU employs two gates: the reset gate R_t and the update gate Z_t, which control hidden state updates based on the current input X_t and the previous hidden state H_{t-1}. Their computations are shown in Equations (4) and (5).

The hidden state H_t is then updated using the gating mechanism described in Equation (6).

$$R_t = \sigma(W_r[H_{t-1}, X_t])4 \tag{4}$$

$$Z_t = \sigma(W_z[H_{t-1}, X_t])5 \tag{5}$$

$$H_t = (1-Z_t)\odot H_{t-1} + Z_t \odot \tanh(W[R_t \odot H_{t-1}, X_t])6 \tag{6}$$

where W_r, W_z are weight matrices, σ denotes the sigmoid function, and $\odot$ the element-wise product.

This gating mechanism enables selective retention and update of information, enhancing learning efficiency while mitigating vanishing gradient issues more simply than LSTMs, with similar predictive power.

V. HYBRIDIZATION OF THE VAE AND GRU MODELS

A. Implementation Steps

1) **Sequence Preparation**: Define the function `prepare_sequences()` which applies a sliding window of size `window_size` to extract input-output sequence pairs from the training dataset.
2) **VAE Encoder**: Construct the VAE encoder composed of two GRU layers, each with 64 units, to encode input sequences shaped as $(\texttt{window\_size}, 1)$. The output corresponds to the latent vector representation.
3) **VAE Decoder**: Design the VAE decoder accepting an input of dimension `latent_dim`. The latent vector is repeated `window_size` times, followed by two GRU layers that produce a one-dimensional output sequence.
4) **GRU Sequential Model**: Define a sequential GRU model, `mode_gru`, starting with a GRU layer of 64 units, with input shape $(\texttt{window\_size}, 1)$.
5) **Hybridization of VAE and GRU**: Enhance the GRU model by integrating VAE features. This can be achieved by inserting a dense layer between the VAE encoder and the GRU model or substituting the first GRU layer with the VAE encoder.

B. Benefits of VAE-GRU Hybridization

The hybrid VAE-GRU model presents several advantages by combining the strengths of both architectures:

- **Choice of VAE for Feature Extraction**: The VAE learns compact, informative latent representations of the input data, effectively reducing dimensionality while preserving salient features, which improves the efficiency of subsequent sequence modeling by the GRU.
- **Modeling Temporal Dynamics with GRU**: The GRU is designed to capture temporal dependencies in sequential data. It is computationally lighter than LSTM and exhibits similar performance, especially for short-term temporal patterns.
- **Improved Generalization**: By coupling the VAE's capacity for feature extraction with the GRU's temporal modeling, the hybrid model enhances generalization capabilities, resulting in improved predictive performance on sequence tasks.

The data flow in the hybrid VAE-GRU model is illustrated in Figure 3, which depicts how input sequences pass through the VAE encoder for feature extraction, followed by GRU layers to capture temporal dependencies, ultimately generating predictions at the output.

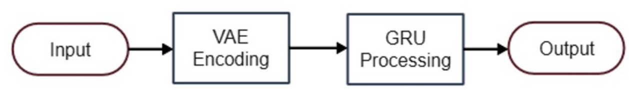

Fig. 3. Data Flow in the VAE-GRU Hybrid Model

C. Disadvantages Overcome by Hybridization

The hybrid approach addresses limitations inherent to each individual model:

- **VAE's Limitations for Sequential Data**: Conventional VAEs do not effectively capture temporal dependencies inherent in sequential data, a gap the GRU component compensates for.
- **GRU's Limitations for Learning Latent Representations**: GRU networks typically lack the ability to learn structured, compact latent representations, which are efficiently provided by the VAE.

In summary, the VAE is employed to extract latent features from time series inputs and reduce noise, while the GRU efficiently models sequential dependencies with fewer parameters than LSTM, thereby mitigating overfitting risks.

VI. DATA ANALYSIS

Hospital emergency services are experiencing growing demand, making resource optimization and accurate anticipation of patient flow essential for operational efficiency. Precise estimation of patient arrivals enables better scheduling of medical staff, optimized bed allocation, and reduction of patient waiting times.

A study was conducted within the emergency department of La Rabta Hospital [17], based on a detailed observation of the hospital process from patient arrival to bed allocation which allowed the collection of granular data on admission volumes over specific time intervals.

Patient arrivals were recorded continuously over a 24-hour period, divided into 24 one-hour segments, as illustrated in Figure 4.

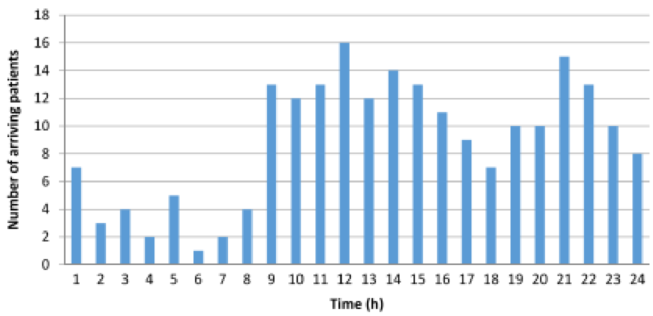

Fig. 4. Number of patient arrivals per hour

The arrival data was modeled using the Poisson distribution [17], justified by the random and independent nature of arrivals, making it suitable for predicting the frequency of events over fixed intervals. To improve prediction accuracy, contextual variables were incorporated into the model, including vacation periods, special events, and seasonal variations, which significantly influence patient flow. For instance, school holidays or major events can increase emergency department overcrowding, while specific weather conditions or seasonal activities may reduce admissions. Integrating these contextual factors allows capturing temporal and environmental fluctuations in patient arrivals more precisely. The Poisson distribution used to model these arrivals is defined by the following equation (7):

$$P(X = k) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{7}$$

where:

- $P(X = k)$ represents the probability of observing k patient arrivals in a given period,
- λ is the average arrival rate per interval,
- e is the base of the natural logarithm,
- $k!$ is the factorial of k.

Our database contains timestamped records of patient arrivals aggregated over 10-minute intervals, spanning from 2009 to 2016. Each entry records the number of arrivals within that time window, providing high-resolution data for detailed flow analysis.

The preprocessing phase followed the CRISP-DM methodology [18], beginning with data understanding, which included identifying missing values and detecting anomalies. Arrival counts exhibited expected variability, reflecting the typical dynamics of hospital activity.

Data preparation involved several key transformations. First, the Date Time column was converted to a datetime format to enable time series analysis. Second, the Number of Arrivals was normalized to maintain data consistency and suitability for statistical and machine learning methods. Finally, cleaning procedures were applied to handle missing or outlier values, ensuring the integrity of the dataset for modeling.

These preprocessing steps facilitated subsequent CRISP-DM stages, including modeling and evaluation of patient flow. The time series exhibited an overall stable trend, with occasional peaks reaching up to 450 arrivals per day, but typically fluctuating between 200 and 300 patients daily. This relative stability supports the feasibility of reliable demand forecasting. Figure 5 presents the daily arrivals time series over the study period.

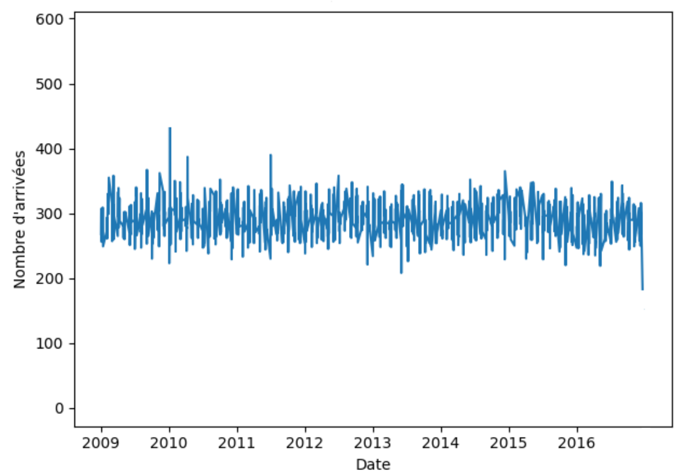

Fig. 5. Time series of daily patient arrivals

For model development, the dataset was split into 80% for training and 20% for validation. To ensure a robust evaluation on unseen data, the last 20% of instances were withheld for testing, respecting the temporal order inherent in time series data.

The model training employed the Adam optimizer with a learning rate of 0.001, balancing efficient convergence and stability during weight updates. The VAE model architecture consisted of two GRU layers (64 units each) in both encoder and decoder, with a latent vector dimension set to `latent_dim` = 10. The standalone GRU model comprised a single GRU layer of 64 units, followed by a dense layer with 8 neurons and a final single-neuron output layer.

Training used a batch size of 64 and 20 epochs, which provided sufficient iterations for learning while mitigating overfitting. A dropout rate of 0.2 was applied to each GRU layer to further prevent overfitting and improve generalization.

Table 1 summarizes the hyperparameters used during training.

VII. DISCUSSION

After tuning and training our method as well as the Deep Learning (DL) models using the training data, we evaluated their predictive performance using the test data. To precisely

Hyperparameter	Value
Learning rate	0.001
Number of GRU layers	2 (encoder), 2 (decoder)
Neurons per GRU layer	64
Latent vector dimension	10
Standalone GRU model	1 GRU layer, 64 neurons
Batch size	64
Number of epochs	20
Dropout rate	0.2
Optimizer	Adam
Loss function	Mean Squared Error (MSE)

TABLE I

SUMMARY OF HYPERPARAMETERS

quantify the performance of the models, we used several evaluation metrics, including Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the Coefficient of Determination (R^2). The explained variance (EV), which indicates the proportion of data variability explained by the model, was also considered. These metrics allow us to analyze the accuracy and efficiency of the predictions made by our models. The detailed results of these evaluations are presented in Table II.

Model	MAE	R^2	RMSE	EV (%)
VAE-GRU	0.32	0.67	0.52	68 %
ARIMA-ML	0.16	0.95	0.20	95 %
GRU	0.35	0.66	0.53	47 %
LSTM	0.36	0.66	0.53	47 %
1D CNN	0.36	0.65	0.54	46 %

TABLE II

EVALUATION METRICS RESULTS OF DIFFERENT PREDICTIVE MODELS

A. Computational Complexity and Training Time

Training the VAE-GRU model on an Intel Core i5 (Ideapad 3) required 44 minutes for 20 epochs, versus only 2 minutes for LSTM. This computational overhead, introduced by the VAE, is offset by improved latent feature extraction and forecasting accuracy.

B. Performance Comparison

The VAE-GRU model outperforms standalone deep models across multiple metrics, though ARIMA-ML remains superior overall:

- **MAE:** ARIMA-ML yields the lowest MAE (0.16), while VAE-GRU achieves 0.32, ahead of GRU (0.35), LSTM (0.36), and 1D CNN (0.36).
- R^2**:** VAE-GRU reaches 0.67, versus 0.66 (GRU/LSTM) and 0.65 (1D CNN).
- **RMSE:** VAE-GRU (0.52) outperforms GRU/LSTM (0.53) and 1D CNN (0.54).
- **Explained Variance:** VAE-GRU explains 68% of the variance, while GRU/LSTM explain 47%, and 1D CNN, 46%.

These results confirm the benefit of integrating VAE with GRU, enhancing temporal modeling through latent representation learning.

C. Comparison of Model Accuracies

Figure 6 and Figure 7 visualize predicted vs. actual values for VAE-GRU, GRU, LSTM, and 1D CNN. All models track general trends, but VAE-GRU shows the closest alignment with observed data, confirming its superior predictive capacity.

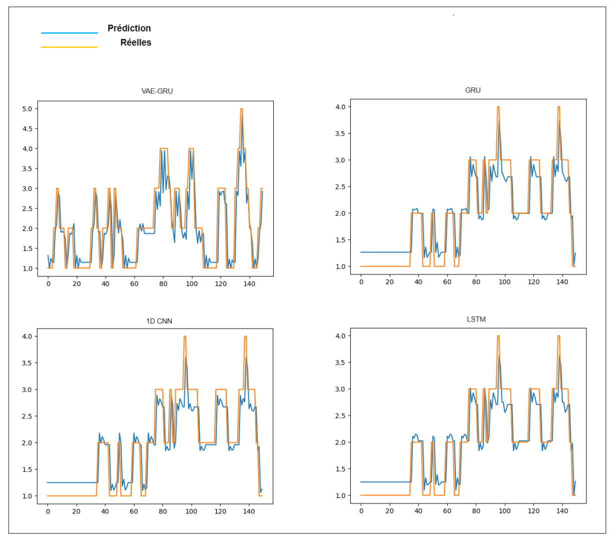

Fig. 6. Predicted vs. actual flows across models (VAE-GRU, GRU, 1D CNN, LSTM)

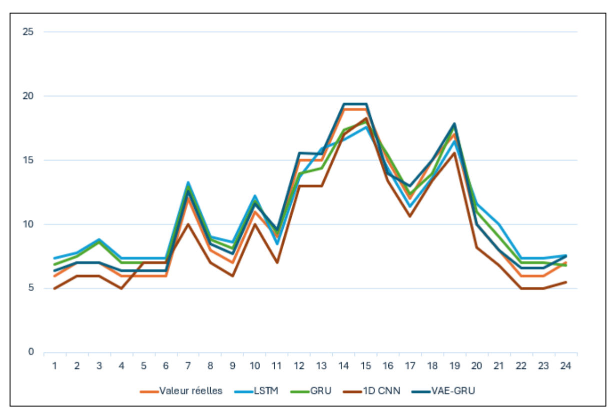

Fig. 7. Model comparison over 24 hourly intervals

D. Comparison with ARIMA-ML

Although the VAE-GRU model captures complex nonlinear temporal dependencies, it underperforms compared to ARIMA-ML in this context. ARIMA-ML achieves superior predictive accuracy with a lower MAE (0.16), higher R^2

(0.95), and reduced RMSE (0.20), reflecting better data variability representation. While VAE-GRU explains 68% of variance, it remains substantially below ARIMA-ML's 95%. Moreover, ARIMA-ML is computationally lighter and faster, favoring real-time deployment.

However, VAE-GRU excels in modeling multifactorial, nonlinear dynamics typical of emergency departments, where patient flow is influenced by complex and evolving factors. Despite higher computational costs and lower interpretability, it offers greater adaptability for long-term forecasting or highly dynamic environments. Thus, model selection should be guided by the application's priorities: short-term accuracy versus robustness to complex, changing conditions.

The next section discusses use cases, limitations, and trade-offs regarding computational demands and real-time feasibility.

E. Limitations and Perspectives

The VAE-GRU model, despite its strong performance, faces several limitations: its high computational complexity restricts use in resource-limited settings; sensitivity to noisy or unstructured data may reduce accuracy; reliance on representative training data limits generalization; and scalability to large-scale hospital environments remains challenging.

Nonetheless, promising results encourage further work on improving computational efficiency and adapting to diverse datasets. The model effectively captures complex dynamics and adapts to temporal variations, though its deployment requires balancing accuracy with resource constraints. Techniques like model compression, architectural refinement, and cloud solutions may enhance scalability and feasibility for practical hospital applications.

VIII. CONCLUSIONS

This paper introduces a novel hybrid model that combines Variational Autoencoders (VAE) and Gated Recurrent Units (GRU), harnessing the strengths of both techniques. VAEs effectively model underlying distributions and manage missing data, while GRUs excel at capturing complex temporal dependencies in time series. Accurate prediction of patient arrivals in emergency departments is essential for efficient resource allocation, reducing waiting times, and minimizing patient length of stay. Experimental results demonstrate that our hybrid model outperforms LSTM, GRU, and 1D CNN-based approaches in accuracy for forecasting emergency visits. These findings open up new avenues for future research and applications, not only in healthcare but also in any domain requiring effective data flow management.

REFERENCES

[1] M. Berquedich, A. Chebak, O. Kamach, O. Laayati, and M. Masmoudi, "An Artificial Immune System for the Management of the Emergency Divisions," in *EAI/Springer Innovations in Communication and Computing*, 2022.

[2] R. Mostafa and K. El-Atawi, "Strategies to measure and improve emergency department performance: a review," *Cureus*, vol. 16, no. 1, 2024.

[3] P. Yang, P. Cheng, N. Zhang, D. Luo, B. Xu, and H. Zhang, "Statistical machine learning models for prediction of China's maritime emergency patients in dynamic: ARIMA model, SARIMA model, and dynamic Bayesian network model," *Frontiers in Public Health*, vol. 12, p. 1401161, 2024.

[4] S. Aguirre, J. Kerexeta, and M. D. Espejo-Mambié, "Time series models for predicting the number of patients attending the emergency department in a local hospital," in *International Work-Conference on Bioinformatics and Biomedical Engineering*, Cham: Springer Nature Switzerland, 2024, pp. 344–355.

[5] D. R. Khan, A. B. Patankar, and A. Khan, "An experimental comparison of classic statistical techniques on univariate time series forecasting," *Procedia Computer Science*, vol. 235, pp. 2730–2740, 2024.

[6] K. Cho, B. van Merrienboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation," *arXiv preprint arXiv:1406.1078*, 2014.

[7] S. Cheng, C. A. Casas, S. Ouala, A. Farchi, C. Liu, P. Tandeo, R. Fablet, D. Lucor, B. Iooss, J. Brajard, D. Xiao, T. Janjić, W. Ding, Y. Guo, A. Carrassi, M. Bocquet, and R. Arcucci, "Machine Learning With Data Assimilation and Uncertainty Quantification for Dynamical Systems: A Review," *IEEE/CAA Journal of Automatica Sinica*, vol. 10, pp. 1361–1387, 2023.

[8] S. Jiang, Q. Liu, and B. Ding, "A systematic review of the modelling of patient arrivals in emergency departments," *Quantitative Imaging in Medicine and Surgery*, vol. 13, no. 3, p. 1957, 2022.

[9] X. Zhao, J. W. Lai, A. F. W. Ho, N. Liu, M. E. H. Ong, and K. H. Cheong, "Predicting hospital emergency department visits with deep learning approaches," *Biocybernetics and Biomedical Engineering*, vol. 42, no. 3, pp. 1051–1065, 2022.

[10] C. Peláez-Rodríguez, R. Torres-López, J. Pérez-Aracil, N. López-Laguna, S. Sánchez-Rodríguez, and S. Salcedo-Sanz, "An explainable machine learning approach for hospital emergency department visits forecasting using continuous training and multi-model regression," *Computer Methods and Programs in Biomedicine*, vol. 245, p. 108033, 2024.

[11] G. Gafni-Pappas and M. Khan, "Predicting daily emergency department visits using machine learning could increase accuracy," *The American Journal of Emergency Medicine*, vol. 65, pp. 5–11, 2023.

[12] H. Hwang, J. H. Jang, E. Lee, H. S. Park, and J. Y. Lee, "Prediction of the number of asthma patients using environmental factors based on deep learning algorithms," *Respiratory Research*, vol. 24, no. 1, p. 302, 2023.

[13] S. Liang, Z. Pan, W. Liu, J. Yin, and M. de Rijke, "A survey on variational autoencoders in recommender systems," *ACM Computing Surveys*, vol. 56, pp. 1–40, 2024.

[14] L. P. Cinelli, M. A. Marins, E. A. B. Da Silva, and S. L. Netto, *Variational methods for machine learning with applications to deep networks*, vol. 15, Springer, 2021.

[15] A. R. Sharafat and M. Bayati, "PatientFlowNet: A Deep Learning Approach to Patient Flow Prediction in Emergency Departments," *IEEE Access*, vol. 9, pp. 45552–45561, 2021, doi: 10.1109/ACCESS.2021.3066164.

[16] F. Xie, J. Zhou, J. W. Lee, M. Tan, S. Li, L. S. O. Rajnthern, *et al.*, "Benchmarking emergency department prediction models with machine learning and public electronic health records," *Scientific Data*, vol. 9, no. 1, p. 658, 2022.

[17] D. Daldoul, I. Nouaouri, H. Bouchriha, and H. Allaoui, "A stochastic model to minimize patient waiting time in an emergency department," *Operations Research for Health Care*, vol. 18, pp. 16–25, 2018.

[18] D. Peixoto, M. Faria, R. Macedo, H. Peixoto, J. Lopes, A. Barbosa, *et al.*, "Determining internal medicine length of stay by means of predictive analytics," in *EPIA Conference on Artificial Intelligence*, Cham: Springer International Publishing, Aug. 2022, pp. 171–182.

[19] J. Chen, Y. Li, C. Zhang, Y. Tian, and Z. Guo, "Urban flooding prediction method based on the combination of LSTM neural network and numerical model," *International Journal of Environmental Research and Public Health*, vol. 20, no. 2, p. 1043, 2023.

[20] S. Xu, J. Dimasaka, D. J. Wald, and H. Y. Noh, "Seismic multi-hazard and impact estimation via causal inference from satellite imagery," *Nature Communications*, vol. 13, no. 1, p. 7793, 2022.