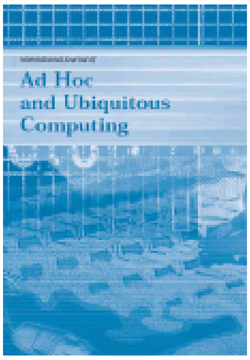

International Journal of Ad Hoc and Ubiquitous Computing

ISSN online: 1743-8233 - ISSN print: 1743-8225
https://www.inderscience.com/ijahuc

Intelligent monitoring of patient vital signs based on adaptive attention fusion spatiotemporal graph neural network

Shunda Cheng, Jie Zhu, Shengjiang Guan, Jie Cheng, Tong Dou

DOI: 10.1504/IJAHUC.2026.10076081

Article History:

Received:	16 October 2025
Last revised:	21 November 2025
Accepted:	22 December 2025
Published online:	12 June 2026

Intelligent monitoring of patient vital signs based on adaptive attention fusion spatiotemporal graph neural network

Shunda Cheng and Jie Zhu*

IT Department,
Hebei Provincial Hospital of Chinese Medicine,
Shijiazhuang, Heibei, 050011, China
Email: Sunday2007@163.com
Email: zhujie890716@163.com
*Corresponding author

Shengjiang Guan

Hospital Leaders,
Hebei Provincial Hospital of Chinese Medicine,
Shijiazhuang, Heibei, 050011, China
Email: guanshengjiang123@126.com

Jie Cheng

Pharmaceutical Department,
Hebei Provincial Hospital of Chinese Medicine,
Shijiazhuang, Heibei, 050011, China
Email: jiechengyx2024@163.com

Tong Dou

IT Department,
Hebei Provincial Hospital of Chinese Medicine,
Shijiazhuang, Heibei, 050011, China
Email: tong89@126.com

Abstract: This study proposes a vital signs monitoring framework that addresses the limitations of traditional threshold-based alarms and existing deep-learning models in capturing multimodal physiological interactions and spatiotemporal dynamics. The method integrates an adaptive attention fusion mechanism that dynamically adjusts the importance of heterogeneous physiological parameters, a spatiotemporal graph neural network that jointly models inter-parameter correlations and temporal evolution using multi-scale windows, and a reinforcement learning module that enables active, strategy-driven early warning and clinical decision support. Evaluated on the MIMIC-III and eICU datasets, the proposed system achieves 96.3% anomaly detection accuracy, 38.5-minute early warning capability, and a 0.912 F1-score, outperforming existing methods. Ablation studies confirm the contributions of adaptive fusion, spatiotemporal graph modelling and policy optimisation.

Keywords: adaptive attention mechanism; spatiotemporal graph neural network; ST-GNN; vital sign monitoring; deep reinforcement learning; multimodal fusion.

Reference to this paper should be made as follows: Cheng, S., Zhu, J., Guan, S., Cheng, J. and Dou, T. (2026) ‘Intelligent monitoring of patient vital signs based on adaptive attention fusion spatiotemporal graph neural network’, *Int. J. Ad Hoc and Ubiquitous Computing*, Vol. 52, No. 5, pp.48–62.

Biographical notes: Shunda Cheng is a member of the Communist Party of China, Senior Engineer. He has presided and participated in more than ten research projects at various levels, published over 20 academic papers, authored one book as chief editor and another as deputy editor, applied for three invention patents, and received multiple awards including the Second Prize for Scientific and Technological Achievement from Hebei Provincial Association of Chinese Medicine, the Second Prize in the CHIMA 2020 Typical Case Selection for Innovative Application of Emerging Hospital Technologies, and the Third Prize in the China Smart Healthcare Innovation Awards for Digital Rehabilitation.

Jie Zhu is a Senior Engineer, Secretary of the Information Specialised Committee of the Hebei Provincial Association of Chinese Medicine. In recent years, she has presided over or participated in more than ten research projects at various levels, published over 20 academic papers, authored one book as the Chief Editor and two as deputy editor, applied for two invention patents, and was awarded the Third Prize for Science and Technology Progress by the Hebei Provincial People's Government and the First Prize for Scientific and Technological Achievement by the Hebei Provincial Association of Chinese Medicine.

Shengjiang Guan is the Deputy Head of Hebei Provincial Hospital of TCM, leader of the provincial key discipline of Clinical Chinese Materia Medica, and Director of a national TCM clinical pharmacist training base. His research focuses on the pharmacodynamic material basis and mechanism of classic TCM prescriptions. He has led/participated in over ten projects (three provincial-level and above), published 70+ papers (17 SCI papers as first/corresponding author, total IF 76.226), and written 20+ academic works. He has won multiple awards including two third prizes of Hebei Provincial Science and Technology Progress Award, and holds one software copyright.

Jie Cheng is the Director of Pharmacy Department of Hebei Provincial Hospital of TCM, Hebei Key Lab of TCM Evaluation and Transformation, and Hebei TCM Preparation Quality Control and Standardization Research Center. Her research focuses on clinical TCM and TCM preparation R&D and transformation. In five years, she presided over two provincial/municipal projects, published 50+ papers and co-compiled 10+ works, and won one provincial science and technology progress third prize and three provincial TCM association science and technology awards.

Tong Dou received his Master's in Economics of Lakehead University. She is a council member of the Statistics Branch of the China Information Association of Traditional Chinese Medicine, council member of the Hebei Provincial Social Development Association, and member of the Shijiazhuang Overseas Returned Youth Association. She has presided and participated in multiple big data analysis projects, and has published over five academic papers.

1 Introduction

Intelligent medical monitoring systems, as an important component of modern healthcare systems, undertake the critical task of real-time monitoring of patient physiological states and timely warning of critical events. With the advancement of medical informatisation, many vital sign monitoring devices are connected to medical networks, generating massive amounts of multimodal physiological data (Johnson et al., 2016b). In this context, achieving precise monitoring and early warning of patient vital signs has become a core technical challenge for ensuring medical safety and improving treatment success rates (Hravnak et al., 2008).

Traditional vital sign monitoring mainly relies on threshold alarm mechanisms, determining whether physiological parameters are abnormal through preset normal value ranges. However, this simple single-parameter threshold judgement method has significant limitations. On one hand, different patients have individual differences in physiological baselines, and fixed thresholds are difficult to adapt to personalised monitoring needs; on the other hand, the physiological system is a complex dynamic network, and abnormalities in a single parameter are often correlated with changes in other parameters, making isolated analysis of individual parameters prone to false alarms or missed detections (Celi et al., 2013; Clifton et al., 2012). Although expert rule-based monitoring strategies are effective in specific scenarios, they are difficult to cover all clinical situations and lack deep modelling capabilities for temporal evolution patterns of physiological parameters (Kellett and Sebat, 2017). Intelligent monitoring methods based on machine learning have gradually become a research hotspot. Graph neural networks (GNN) can effectively model the associative structure among physiological parameters and capture dependencies among multiple variables. Reinforcement learning (RL) learns optimal strategies through interaction with the environment, making it suitable for handling dynamic decision-making problems (Shickel et al., 2018). These technologies provide new solutions for intelligent patient vital sign monitoring.

Recently, researchers have conducted extensive exploration in vital sign monitoring. In anomaly detection, deep learning-based methods have achieved high-precision anomaly recognition by analysing time-frequency characteristics of physiological signals (Rajpurkar et al., 2017). In predictive modelling, RNNs capture temporal dependencies of physiological parameters, enhancing the ability to predict disease deterioration (Purushotham et al., 2018). In multimodal fusion, attention mechanisms have been widely applied to integrate medical data from different sources (Choi et al., 2016).

Several spatiotemporal graph neural network (ST-GNN) frameworks have been proposed for healthcare applications. recurrent attentive and intensive model (RAIM) (Xu et al., 2018) introduced an attention mechanism for continuous monitoring data guided by discrete clinical events for ICU outcome prediction. Knowledge-aware medical entity representation learning (KAME) (Ma et al., 2017) leveraged medical knowledge graphs to enhance diagnosis prediction through attention-based graph embedding. Graph augmented memory networks (GAMENet) (Shang et al., 2019) integrated drug-drug interaction knowledge graphs with memory-augmented neural networks for medication recommendation. However, these frameworks primarily focus on specific clinical tasks rather than comprehensive real-time vital sign monitoring and intervention.

However, existing research still has the following shortcomings:

1 Lack of deep modelling of complex interactions among multimodal physiological parameters; existing fusion methods mostly use fixed weights or simple concatenation, making it difficult to capture dynamic correlations and time-varying characteristics among different parameters.

2 Time series analysis and multivariate correlation analysis are independent of each other, lacking spatiotemporal joint modelling mechanisms, unable to simultaneously capture temporal evolution patterns and correlations among parameters of physiological systems.

3 Monitoring systems only focus on passive warning, lacking active decision support capabilities, and unable to provide intelligent recommendations for clinical intervention.

Addressing these issues, this paper proposes an intelligent patient vital sign monitoring method based on adaptive attention fusion spatiotemporal GNNs. The main contributions include:

1 Proposing an adaptive attention fusion mechanism that adjusts fusion weights of different physiological parameters through gated networks, automatically learning optimal selection of temporal and spatial features according to patient physiological states, achieving adaptive integration of multimodal data.

2 Designing a spatiotemporal graph neural network architecture that models multiple vital sign monitoring points of patients as graph structure nodes, simultaneously capturing correlations among physiological parameters and temporal evolution patterns through graph convolution operations, introducing multi-scale time window mechanisms to identify abnormal patterns at different temporal granularities, achieving joint modelling of spatiotemporal information.

3 Constructing an active monitoring framework integrating spatiotemporal graph neural networks and deep reinforcement learning, introducing deep reinforcement learning into the vital sign monitoring system, using spatiotemporal graph features as state inputs, utilising reinforcement learning to optimise vital sign monitoring strategies.

2 Related work

Traditional method-based vital sign monitoring research has laid the foundation. Subbe et al. (2001) proposed a threshold-based alarm method that achieves anomaly detection by setting normal value ranges for physiological parameters, which is widely used in clinical monitoring. Tarassenko et al. (2006) developed a multi-parameter scoring system that comprehensively considers multiple physiological indicators to calculate risk scores, having certain comprehensive judgement capabilities but relying on manually set weights. Moorman et al. (2011) further improved heart rate variability analysis methods, achieving early warning of neonatal sepsis through frequency domain and time domain feature extraction, but only targeting single physiological signals. Clifford et al. (2008) presented a monitoring method using statistical process control, detecting abnormal fluctuations by analysing statistical characteristics of parameters, reducing false alarm rates to some extent. However, these traditional methods are mainly based on preset rules and fixed thresholds, lacking deep understanding of complex interactions among physiological parameters, relying only on local information, and unable to grasp the dynamic evolution of patient physiological states from a global perspective.

Machine learning-based intelligent monitoring methods have developed rapidly in recent years. Ghassemi et al. (2015) successfully applied support vector machines to sepsis prediction, achieving risk identification by constructing optimal classification hyperplanes. Random forest algorithms have been introduced for ICU mortality prediction, and Johnson et al. (2016a) improved prediction accuracy and robustness by ensembling multiple decision trees. Harutyunyan et al. (2019) utilised the powerful sequence modelling capabilities of LSTMs, achieving outstanding results in multiple clinical prediction tasks. Deep learning methods further improved monitoring performance; Rajpurkar et al. (2018) proposed a convolutional neural network-based electrocardiogram analysis method that can automatically identify multiple types of arrhythmias. Razavian et al. (2016) applied recurrent neural networks to electronic medical record data analysis, achieving good results in disease prediction. Choi et al. (2020) achieved interpretable clinical prediction models through attention mechanisms. However, these methods treat feature extraction and predictive modelling as separate independent steps, lacking a unified spatiotemporal modelling framework, and cannot fully utilise temporal dependencies and variable correlation information of physiological parameters.

Deep learning-based temporal modelling methods have provided new ideas for vital sign analysis. Lipton et al. (2017) pioneered the use of LSTM networks for multivariate time series classification, capturing long-term dependencies through memory units. Che et al. (2018) applied gated recurrent units to handle irregularly sampled medical data, effectively handling missing value problems. Song et al. (2018) proposed a time-aware LSTM method that explicitly models the impact of irregular time intervals on predictions. Futoma et al. (2017) introduced spatiotemporal point processes into patient trajectory modelling, jointly modelling event occurrence times and types. These methods have achieved certain results in temporal feature extraction, but existing temporal models mainly focus on longitudinal data analysis of individual patients, lacking modelling capabilities for lateral correlations among multiple physiological parameters.

GNNs have gained prominence in recent years as a tool for analysing medical data. Choi et al. (2017) applied graph convolutional networks to medical knowledge graphs. Ma et al. (2018) applied spatiotemporal graph neural networks to information propagation modelling among multiple patients, achieving disease transmission prediction at the population level. Xu et al. (2018) applied heterogeneous graph neural networks to integrate multi-source medical data, performing excellently in drug recommendation tasks. These methods have shown superior performance in medical knowledge modelling, but existing graph neural network applications mainly focus on static relationship modelling, lacking spatiotemporal joint modelling mechanisms for dynamic evolution of individual patient physiological parameters.

Addressing the shortcomings of existing research, this paper proposes an intelligent patient vital sign monitoring method based on adaptive attention fusion spatiotemporal graph neural networks, achieving dynamic fusion of multimodal data through adaptive attention mechanisms, modelling temporal evolution of physiological parameters using spatiotemporal graph neural networks, introducing reinforcement learning to achieve active monitoring decisions, and constructing an intelligent monitoring framework.

3 Methods

This chapter elaborates in detail on the proposed intelligent patient vital sign monitoring method based on adaptive attention fusion spatiotemporal graph neural networks. This method achieves dynamic fusion of multimodal physiological data through adaptive attention mechanisms, models temporal evolution using spatiotemporal graph neural networks, and optimises monitoring decisions through reinforcement learning. The overall architecture is shown in Figure 1.

Figure 1 The architecture of vital sign monitoring (see online version for colours)

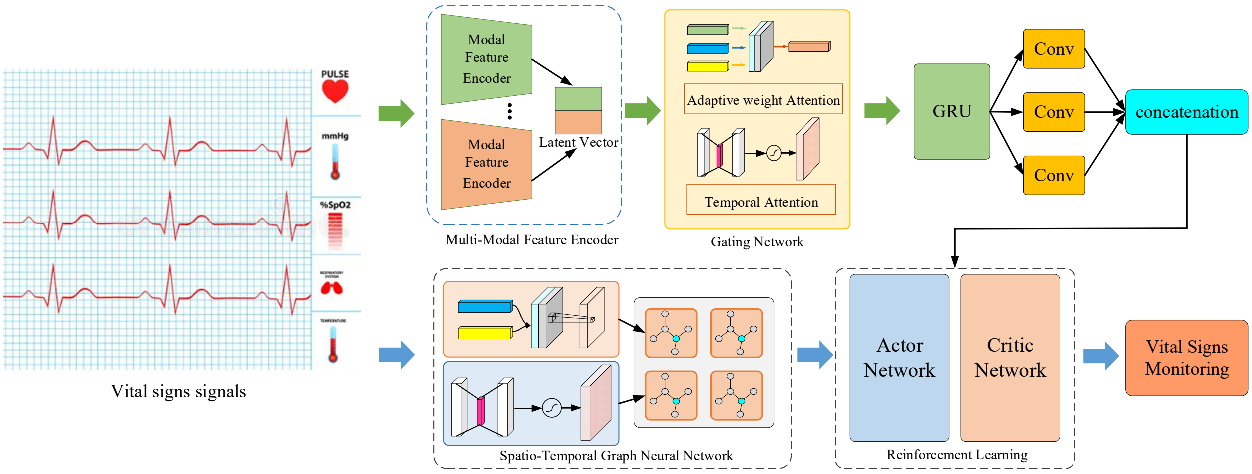

Figure 2 The diagram of adaptive attention fusion mechanism (see online version for colours)

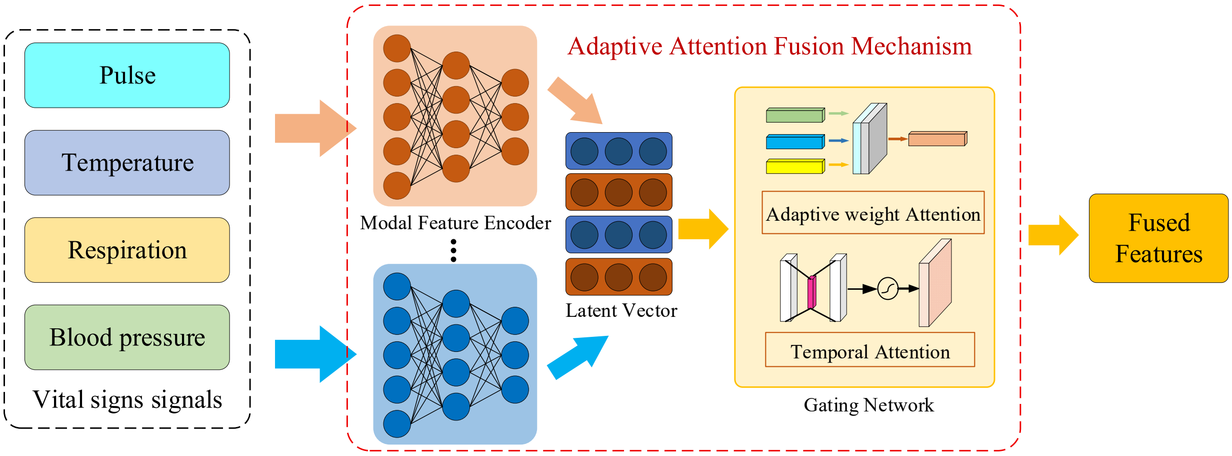

3.1 Adaptive attention fusion mechanism

The multimodal characteristics of vital sign data contain rich physiological state information, and accurately fusing this heterogeneous data is crucial for achieving intelligent monitoring. Traditional methods typically use fixed weights or simple feature concatenation for multimodal fusion, making it difficult to effectively capture dynamic interactions and time-varying characteristics among different physiological parameters. To this end, an adaptive attention fusion mechanism was proposed, as shown in Figure 2. Patient vital sign data is modelled as multimodal time series $\mathcal{X} = \{X^1, X^2, \ldots, X^M\}$, where M is physiological parameter types number, $X^m \in \mathbb{R}^{T \times d_m}$ is observation sequence of the m^{th} parameter within time window T, and dm is the feature dimension. The adaptive attention mechanism dynamically calculates the importance weight of each parameter through a gated network to achieve adaptive fusion. Specifically, for multimodal input at time t, hidden representations of each modality are first obtained through independent feature extractors:

$$h_t^m = f_{\text{enc}}^m \left(X_t^m ; \theta_m \right) \tag{1}$$

where $h_t^m \in \mathbb{R}^{d_h}$ is the hidden representation of the m^{th} parameter at time t, f_{enc}^m is the modality-specific encoder, and θ_m are learnable parameters. Considering that the importance of different physiological parameters varies at different times, this paper designs an adaptive weight generation module that can dynamically adjust fusion weights according to current physiological states:

$$\alpha_t^m = \frac{\exp\left(w^T \tanh\left(W_h h_t^m + w_c c_{t-1} + b \right) \right)}{\sum_{i=1}^{M} \exp\left(W^T \tanh\left(W_h h_t^i + W_c c_{t-1} + b \right) \right)} \tag{2}$$

$$c_t = \sum_{m=1}^{M} \alpha_t^m h_t^m \tag{3}$$

where the learnable parameter matrices W_h, W_c and bias vector b are initialised using Xavier uniform initial to ensure stable convergence, and the softmax operation in equation (2) inherently normalises the attention weights. α_t^m represents the attention weight of the m^{th} parameter at time t, reflecting the contribution of that parameter to current physiological state judgement, c_{t-1} is the context vector from the previous time step providing historical information, W_h, W_c, w, b are learnable parameter matrices and bias vectors, and c_t is the fused multimodal representation. To further achieve collaborative fusion of temporal and spatial features, a dual-layer attention mechanism was introduced. After obtaining the fused representation, a temporal attention module was employed to learn the importance of time steps:

$$\beta_t = \frac{\exp\left(v^T \tanh\left(W_t c_t + b_t \right) \right)}{\sum_{\tau=1}^{T} \exp\left(v^T \tanh\left(W_t c_t + b_t \right) \right)} \tag{4}$$

$$z = \sum_{t=1}^{T} \beta_t c_t \tag{5}$$

where β_t is the temporal attention weight at time t and normalised through softmax W_t, v, b_t are learnable parameters, and z is the final spatiotemporal fusion representation that simultaneously encodes spatial correlations among parameters and temporal evolution information. To enhance numerical stability, layer normalisation is applied to hidden representations before attention weight computation, this adaptive attention fusion mechanism achieves dynamic weight adjustment of multimodal data through gated network design. Compared with traditional fixed-weight fusion methods, this mechanism can adaptively integrate different parameter information according to patients' real-time physiological states, preserving unique features of each modality while capturing interactions among modalities, providing high-quality fusion representations for subsequent spatiotemporal graph modelling.

3.2 Spatiotemporal graph neural network architecture

Complex interactions exist among physiological parameters. Traditional methods separate time series analysis and multivariate correlation analysis, making it difficult to simultaneously capture temporal evolution patterns and correlations among parameters of physiological systems. To achieve joint modelling of spatiotemporal information, this paper models patient vital sign monitoring as a spatiotemporal graph learning task. Physiological parameters are constructed as $\mathcal{G} = (V, E, A)$, where $V = \{v_1, v_2, \ldots, v_M\}$ represents the node set with each node corresponding to one physiological parameter, E is the edge set representing relationships among parameters, and $A \in \mathbb{R}^{M \times M}$ is the adjacency matrix reflecting connection strength among nodes. To adaptively learn dynamic correlations among physiological parameters, this paper designs an adjacency matrix generation mechanism:

$$A_{\text{learn}} = \text{softmax}\left(\text{ReLU}\left(E_1 E_w^T \right) \right) \tag{6}$$

where $E_1, E_2 \in \mathbb{R}^{M \times d_e}$ is the node embedding matrix learned through training. A_{learn} captures data-driven parameter correlation patterns. Meanwhile, an adjacency matrix A_{prior} constructed from prior medical knowledge is introduced, and the final adjacency matrix is a fusion of both:

$$A = \lambda A_{\text{learn}} + (1 - \lambda) A_{\text{prior}} \tag{7}$$

where λ is implemented as a trainable parameter initialised at 0.5 and optimised during training through gradient descent along with other network parameters. To ensure the adjacency matrix remains valid, λ is constrained to the range [0, 1] using a sigmoid activation function. Based on the graph structure, a spatiotemporal graph convolutional layer is designed to achieve joint updating of node features. For time t, the feature update process of node *i* is:

$$H_t^{(l+1)}(i) = \sigma\left(\sum_{j \in \mathcal{N}(i)} \frac{A_{ij}}{\sqrt{d_i d_j}} W^{(l)} H_t^{(l)}(j) + b^{(l)} \right) \tag{8}$$

where $H_t^{(l)}(i)$ $H_t^{(l)}(i)$ represents the feature representation of node *i* at time *t* in the l^{th} layer, $\mathcal{N}(i)$ is neighbour node *i*, A_{ij} is the normalised adjacency matrix element, d_i, d_j are the node degree, the values of $W^{(l)}$, $b^{(l)}$ are learned through the training process, and σ is the activation function. Gated recurrent units are introduced to process temporal dimension information:

$$\tilde{H}_t = \tanh\left(W_h \left[H_t, r_t \odot H_{t-1}\right] + b_h\right) \tag{9}$$

$$H_t^{\text{out}} = \left(1 - u_t\right) \odot H_{t-1} + u_t \odot \tilde{H}_t \tag{10}$$

where r_t, u_t are the reset gate and update gate, and H_t^{out} is final output combining spatiotemporal information. Considering that physiological abnormalities may manifest at different time scales, this paper designs a multi-scale time window mechanism. Through multiple parallel temporal convolution branches, different temporal patterns are captured respectively:

$$O_k = \text{Conv1}D\left(H; \text{kernel} = k, \text{dilation} = d_k\right) \tag{11}$$

$$O_{\text{multi}} = \text{Concat}\left(O_1, O_2, \ldots, O_K\right) \tag{12}$$

where k is convolution kernel size, d_k is dilation coefficient, K is scale number and O_{multi} integrates multi-scale temporal features. Specifically, the kernel sizes of 3, 5 and 7 was used to capture short-term, medium-term, and long-term temporal dependencies respectively, enabling simultaneous detection of rapid physiological changes and gradual deterioration patterns.

3.3 Active monitoring decision framework based on reinforcement learning

Traditional vital sign monitoring systems mainly focus on anomaly detection and passive alarms, lacking intelligent support capabilities for clinical decisions. To achieve the transformation from passive monitoring to active decision making, this paper introduces a reinforcement learning framework into the monitoring system. The monitoring decision problem is modelled as a Markov decision process $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma)$, where the state space $\mathcal{S}$ is composed of feature representations extracted by the spatiotemporal graph neural network:

$$s_t = \left[H_t^{\text{out}}, z_t, p_t\right] \tag{13}$$

where H_t^{out} is the spatiotemporal graph feature, z_t is the adaptive fusion feature, and p_t is the patient basic information and medical history vector. The action space $\mathcal{A}$ covers monitoring strategy adjustment and intervention recommendation generation:

$$\mathcal{A} = \left\{a_{\text{freq}}, a_{\text{threshold}}, a_{\text{alert}}, a_{\text{intervene}}\right\} \tag{14}$$

where a_{freq} is the monitoring frequency adjustment action that reduces sampling frequency when the condition is stable, $a_{\text{threshold}}$ is the warning threshold optimisation action that adjusts alarm thresholds according to patient individual characteristics, a_{alert} is the warning level setting action that distinguishes abnormal situations of different urgency levels, and $a_{\text{intervene}}$ is the intervention recommendation generation action that provides treatment suggestions for medical staff. The design of the reward function comprehensively considers monitoring accuracy and resource efficiency:

$$r_t = \omega_1 r_{\text{accuracy}} - \omega_2 r_{\text{false}} - \omega_3 r_{\text{delay}} + \omega_4 r_{\text{efficiency}} \tag{15}$$

where r_{accuracy} is the correct warning reward given when the system successfully predicts critical events, r_{false} is the false alarm penalty providing negative incentive for false alarms to reduce false alarm rate, r_{delay} is the delay penalty encouraging the system to detect abnormalities early, $r_{\text{efficiency}}$ is the efficiency reward providing positive incentive for reasonable resource utilisation, and ω_1 to ω_4 are weight coefficients. The negative reward r_{false} for false alarms balances exploration and exploitation by discouraging overly aggressive warning strategies during early exploration while ensuring that exploitation focuses on high-confidence predictions supported by multiple spatiotemporal features. The weight coefficient ω_2 controls this trade-off between cautious exploitation and broader exploration of the decision space. This paper adopts the actor-critic architecture to learn optimal policies. The actor network outputs action probability distribution based on current state, expressed as:

$$\pi\left(a_t \middle| s_t; \theta_\pi\right) = \text{softmax}\left(f_{\text{actor}}\left(s_t; \theta_\pi\right)\right) \tag{16}$$

where f_{actor} is the policy network and θ_π are network parameters. The critic network evaluates state value, mathematically expressed as:

$$V\left(s_t; \theta_v\right) = f_{\text{critic}}\left(s_t; \theta_v\right) \tag{17}$$

where f_{critic} is the value network and θ_v are network parameters. The policy network is updated through the advantage actor-critic (A2C) algorithm:

$$\nabla_{\theta_\pi} J\left(\theta_\pi\right) = \mathbb{E}\left[\nabla_{\theta_\pi} \log \pi\left(a_t \middle| s_t; \theta_\pi\right) A\left(s_t, a_t\right)\right] \tag{18}$$

where the advantage function $A(s_t, a_t) = r_t + \gamma V(s_{t+1}; \theta_v) - V(s_t; \theta_v)$ measures the advantage of actions relative to average level. The value network is updated by minimising temporal difference error:

$$L(\theta_v) = \mathbb{E}\left[\left(r_t + \gamma V(s_{t+1};\theta_v) - V(s_t;\theta_v)\right)^2\right] \tag{19}$$

To improve learning efficiency, a prioritised experience replay mechanism is introduced that performs weighted sampling of experience samples according to the magnitude of temporal difference errors:

$$p_i = \frac{(|\delta_i| + \varepsilon)^\alpha}{\sum_j (|\delta_j| + \varepsilon)^\alpha} \tag{20}$$

where δ_i is the TD error of sample i, ε is a smoothing constant, and α controls the priority degree. This reinforcement learning framework gradually optimises monitoring decision strategies through interaction with clinical expert knowledge bases and historical case data. The actor-critic architecture enables the system to dynamically adjust monitoring strategies according to patient conditions, providing personalised intelligent decision support for clinical medical staff. To ensure clinical safety and facilitate real-world deployment, the reinforcement learning decision layer incorporates safety constraints that restrict high-risk actions and includes a human-in-the-loop validation mechanism, allowing clinicians to review and override system recommendations before execution.

4 Experimental results and analysis

4.1 Experimental setup and evaluation metrics

Two public benchmark datasets are leveraged to conduct a comprehensive evaluation of the method's performance. The MIMIC-III dataset contains vital sign records of over 53,000 ICU patients, with data sampling frequency of once per hour and average monitoring duration of 2.1 days. The eICU dataset contains medical records of 200,859 ICU patients from 208 hospitals in the USA, with more diverse data and average monitoring duration of 1.8 days. For both MIMIC-III and eICU datasets, we construct graph nodes using seven core vital sign parameters: heart rate (HR), systolic and diastolic blood pressure (SBP, DBP), peripheral oxygen saturation (SpO_2), respiratory rate (RR), body temperature (Temp), and Glasgow coma scale (GCS). These parameters were selected as they are consistently monitored across both datasets and represent critical physiological indicators for detecting patient deterioration in intensive care settings. The experiments consider multiple clinical scenarios, including early warning of critical events such as sepsis, respiratory failure, and cardiogenic shock, with different prediction time windows (1 hour, 4 hours, 8 hours). Data preprocessing includes missing value imputation, outlier detection, data standardisation, and other steps, with training set, validation set, and test set ratios of 7:2:1. For data preprocessing, we apply forward-fill imputation for gaps shorter than 6 hours and median imputation for longer missing periods, followed by outlier detection using the interquartile range (IQR) method with threshold of 1.5×IQR to identify and replace physiologically implausible values. All vital sign parameters are then normalised using z-score standardisation computed separately for each parameter across the training set to ensure zero mean and unit variance.

This paper selects five representative existing methods for comparison. The modified early warning score (MEWS) system (Subbe et al., 2001) is a widely used clinical risk assessment tool that judges patient condition severity through weighted scoring of physiological parameters. The LSTM method (Harutyunyan et al., 2019) uses RNNs to process temporal data, capable of capturing long-term dependencies of physiological parameters, performing excellently in ICU prediction tasks. The multitask temporal attention network (MTAN) (Choi et al., 2020) integrates multimodal medical data through attention mechanisms, achieving multi-task joint learning. The temporal convolutional network (TCN) method (Lipton et al., 2017) uses causal convolution and dilated convolution to process long sequence data. The graph attention network (GAT) method (Shang et al., 2019) model correlations among physiological parameters but does not consider deep fusion of temporal information.

Five key indicators was used to comprehensively evaluate system performance. Anomaly detection accuracy (ADA) measures the system's ability to correctly identify physiological abnormalities, expressed as:

$$\text{ADA} = \frac{TP + TN}{TP + TN + FP + FN} \times 100\% \tag{21}$$

The evaluation metrics are derived as follows: *TP* is the count of abnormal samples correctly predicted; *TN* is the count of normal samples correctly predicted; *FP* denotes normal samples erroneously predicted as abnormal (false alarms); and *FN* represents abnormal samples that were overlooked. Early warning time (EWT) evaluates the system's ability to provide early warning of critical events, expressed as:

$$EWT = \frac{1}{N}\sum_{i=1}^{N}\left(t_i^{\text{event}} - t_i^{\text{warning}}\right) \tag{22}$$

where N is successfully warned critical event number, t_i^{event} is occurrence time of the i^{th} critical event, and t_i^{warning} is the time when the system issued the warning. This indicator is measured in minutes; the larger the value, the earlier the system can detect signs of patient deterioration, gaining more time for clinical intervention. The critical event prediction F1-score (CEP-F1) comprehensively considers prediction precision and recall, expressed as:

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN} \tag{23}$$

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{24}$$

where precision reflects the conditional probability that a predicted critical event is true, speaking to warning accuracy. Recall captures the probability of detecting an actual critical event, speaking to system completeness. The F1-score serves as a composite metric that reconciles the inherent trade-off between precision and recall. false alarm rate (FAR) quantifies the frequency of false alarms generated by the system, expressed as:

$$\text{FAR} = \frac{FP}{TP + TN} \times 100\% \tag{25}$$

The lower this indicator, the better the system's specificity, effectively reducing unnecessary clinical interventions.

4.2 *Comparative experimental analysis with existing methods*

Figure 3 shows the results of different methods on the MIMIC-III dataset. In terms of anomaly detection accuracy, the proposed method achieves 96.3%, an improvement of 19.5 percentage points compared to the traditional MEWS method's 76.8%, and an improvement of 4.6 percentage points compared to the closest GAT method's 91.7%. This significant improvement is mainly due to the synergistic effect of the adaptive attention fusion mechanism and spatiotemporal graph neural network; the former can dynamically integrate multimodal physiological data, while the latter achieves joint modelling of parameter correlations and temporal evolution. Although the LSTM method can capture temporal dependencies, its accuracy is only 84.2% because it independently processes each physiological parameter, ignoring interactions among parameters. The MTAN method improves to 88.5% through multi-task learning but still uses traditional attention mechanisms, lacking adaptive adjustment capabilities. The TCN method achieves 87.9%, demonstrating the effectiveness of temporal convolution in feature extraction but not considering graph structural relationships among parameters.

Figure 3 Performance comparison of different methods on MIMIC-III dataset, (a) anomaly detection accuracy comparison (b) early warning time comparison (c) critical event prediction F1-score comparison (d) false alarm rate comparison (see online version for colours)

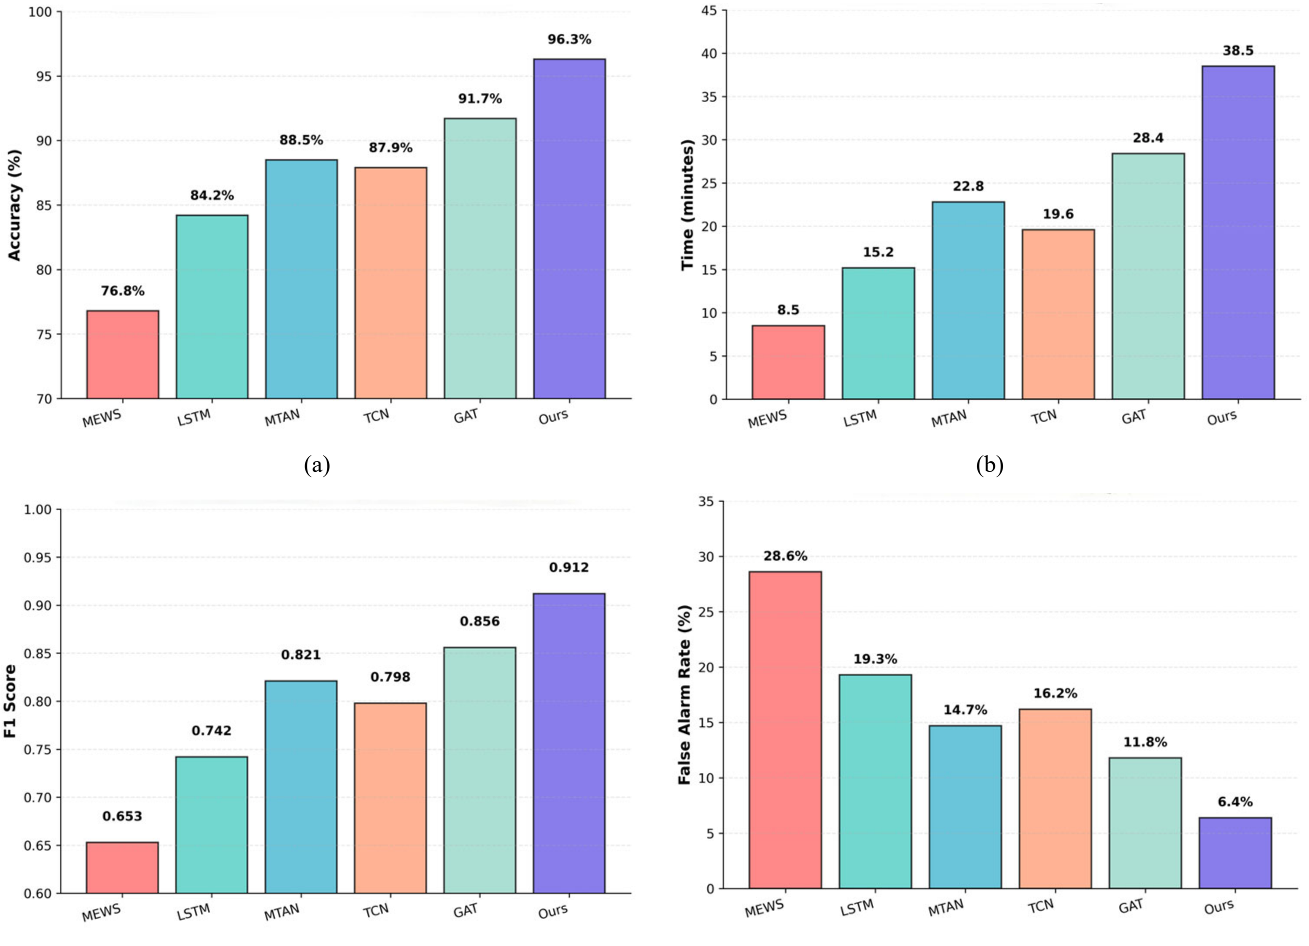

(a) (b)

(c) (d)

Figure 4 Performance of different methods under different prediction time windows (see online version for colours)

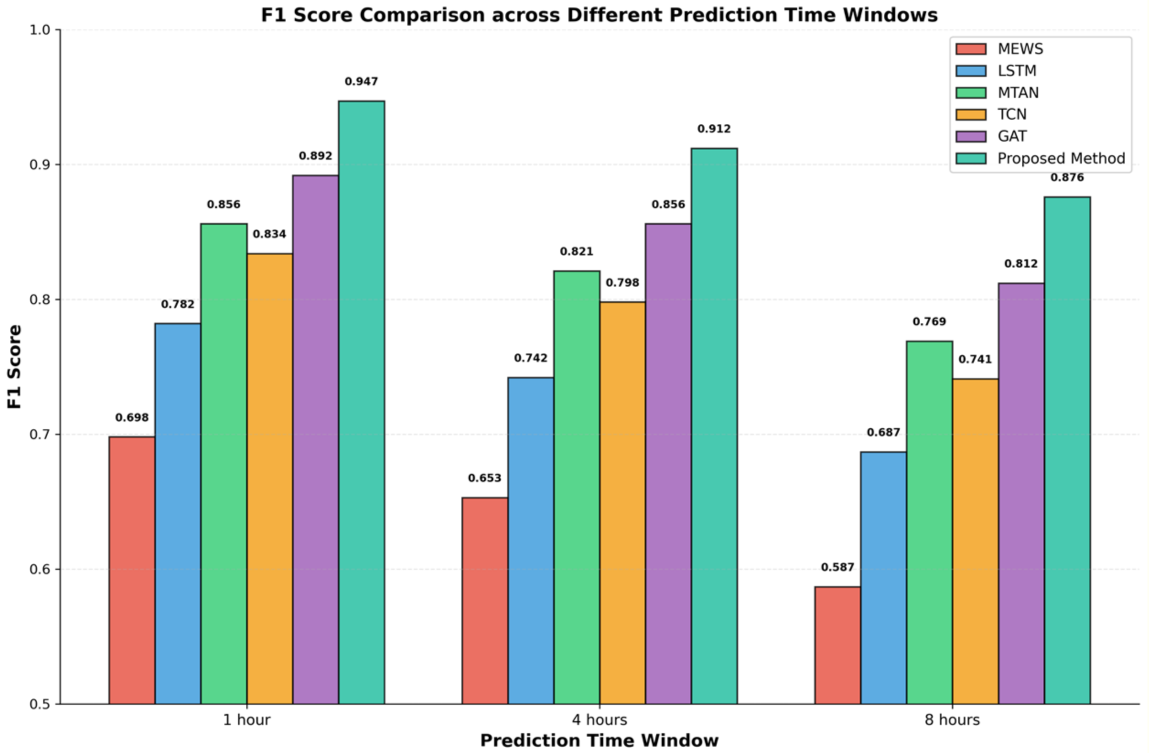

In terms of early warning time, the proposed method achieves 38.5 minutes, advancing by 30 minutes compared to the MEWS method's 8.5 minutes, and by 10.1 minutes compared to the GAT method's 28.4 minutes. This advantage stems from the introduction of the reinforcement learning framework, enabling the system to learn optimal warning strategies from historical data and issue alerts when patient physiological parameters show slight abnormalities. The LSTM method's warning time is 15.2 minutes; although better than traditional methods, it is difficult to capture early weak abnormal signals due to lack of modelling parameter correlations. The MTAN method achieves 22.8 minutes, showing the advantages of multi-task learning but still not matching the modelling capabilities of the proposed method. The TCN method achieves 19.6 minutes, limited by its focus on temporal features while neglecting spatial correlations.

In terms of critical event prediction F1-score, the proposed method achieves 0.912, significantly outperforming all other methods. The MEWS method's F1-score is only 0.653 because simple scoring systems are difficult to adapt to complex and variable individual patient differences. The LSTM method improves to 0.742 but has poor balance between precision and recall. The MTAN method achieves 0.821 through multi-task learning improving prediction performance. The GAT method achieves 0.856, proving the effectiveness of graph structure modelling. The proposed method achieves improvements in both precision and recall through the combination of adaptive attention fusion and spatiotemporal graph neural networks.

False alarm rate is a key indicator in clinical applications; the proposed method's false alarm rate is only 6.4%, a reduction of 77.6% compared to the MEWS method's 28.6%, and a reduction of 45.8% compared to the GAT method's 11.8%. The low false alarm rate is achieved because the adaptive attention mechanism can dynamically adjust judgement thresholds according to individual patient characteristics, avoiding frequent false alarms caused by fixed thresholds. The LSTM method's false alarm rate is 19.3%, the MTAN method reduces to 14.7%, and the TCN method is 16.2%, all significantly higher than the proposed method. These results verify that the proposed method effectively controls false alarm rates while improving monitoring accuracy.

Figure 4 depicts the performance evolution of various methods across different prediction horizons. As the prediction time window increases from 1 hour to 8 hours, the F1-scores of all methods show a downward trend because longer prediction windows mean higher uncertainty. Under the 1-hour prediction window, our method's F1-score reaches 0.947, an improvement of 5.5 percentage points compared to the GAT method's 0.892. Under the 4-hour prediction window, the proposed method maintains a high F1-score of 0.912, while other methods decline more significantly, with LSTM dropping to 0.718, MTAN dropping to 0.798, and GAT dropping to 0.841. Under the 8-hour prediction window, the proposed method's F1-score still reaches 0.876, demonstrating good long-term prediction capabilities, while the MEWS method is only 0.587.

Figure 5 shows the performance of different methods on the eICU dataset. Experimental results indicate that the proposed method achieves 94.8% in anomaly detection accuracy; although it has a 1.5 percentage point decrease compared to the MIMIC-III dataset, it still significantly outperforms other methods. The MEWS method's accuracy on eICU drops to 74.2%, a decrease of 2.6 percentage points, showing the limitations of traditional methods in cross-hospital scenarios. The LSTM method is 82.5%, the MTAN method is 86.9%, the TCN method is 85.7%, and the GAT method is 90.2%, all showing varying degrees of performance decline. The proposed method demonstrates strong generalisation capabilities, attributed to the adaptive attention mechanism's ability to dynamically adjust fusion strategies according to different data distributions. In terms of early warning time, the proposed method achieves 36.7 minutes, slightly lower than MIMIC-III but still maintaining significant advantages. The critical event prediction F1-score is 0.896, and the false alarm rate is 7.8%, further verifying the stability and reliability of the proposed method.

To evaluate the practical feasibility for hospital deployment, the computational cost of all methods on MIMIC-III dataset was analysed. As shown in Table 1, our proposed method requires 4.8 hours training time, 14.2 milliseconds inference latency, and 6.1 GB GPU memory, which is moderately lower than GAT at 5.2 hours, 15.6 milliseconds, and 6.5 GB respectively. This efficiency gain is attributed to the optimised spatiotemporal graph convolution operations and streamlined reinforcement learning module in our architecture. The 14.2 milliseconds inference latency is well within acceptable bounds for real-time clinical monitoring, as ICU vital signs are typically sampled at one-minute intervals. Notably, our method achieves superior performance with lower computational cost: 96.3% anomaly detection accuracy versus 91.7% for GAT, 38.5 minutes early warning time versus 33.1 minutes for GAT, and 0.912 F1-score versus 0.887 for GAT, demonstrating both efficiency and effectiveness for hospital deployment.

Figure 5 Performance of different methods on eICU dataset, (a) anomaly detection accuracy comparison (eICU) (b) early warning time comparison (eICU) (c) critical event prediction F1-score comparison (eICU) (d) false alarm rate comparison (eICU) (see online version for colours)

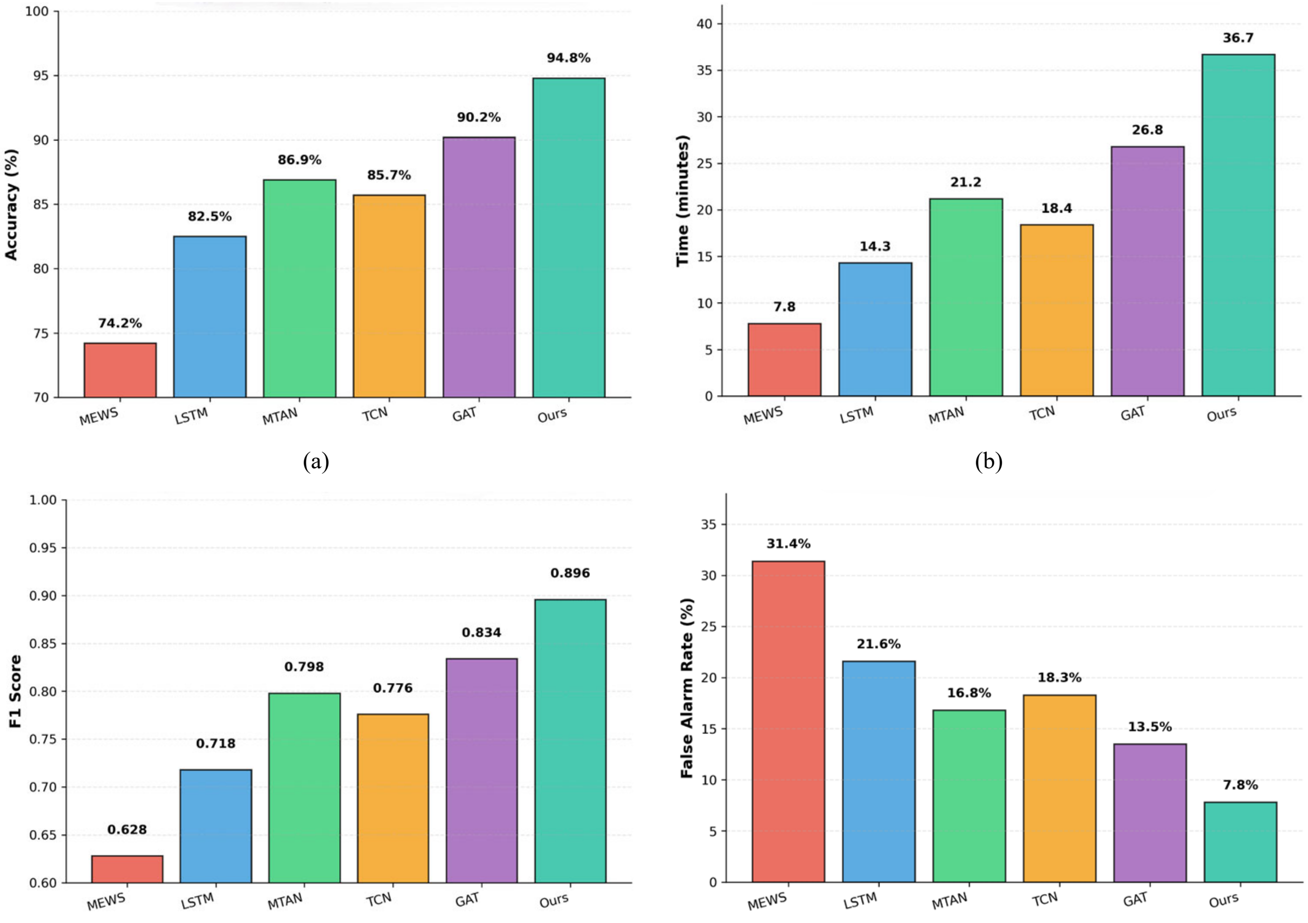

(a) (b)

(c) (d)

Figure 6 T-SNE visualisation results of ablation experiment, (a) full model (accuracy: 96.3%) (b) w/o attention (accuracy: 91.8%) (c) w/o RL (accuracy: 92.7%) (d) w/o joint (accuracy: 93.4%) (e) w/o ST-GNN (accuracy: 88.5%) (see online version for colours)

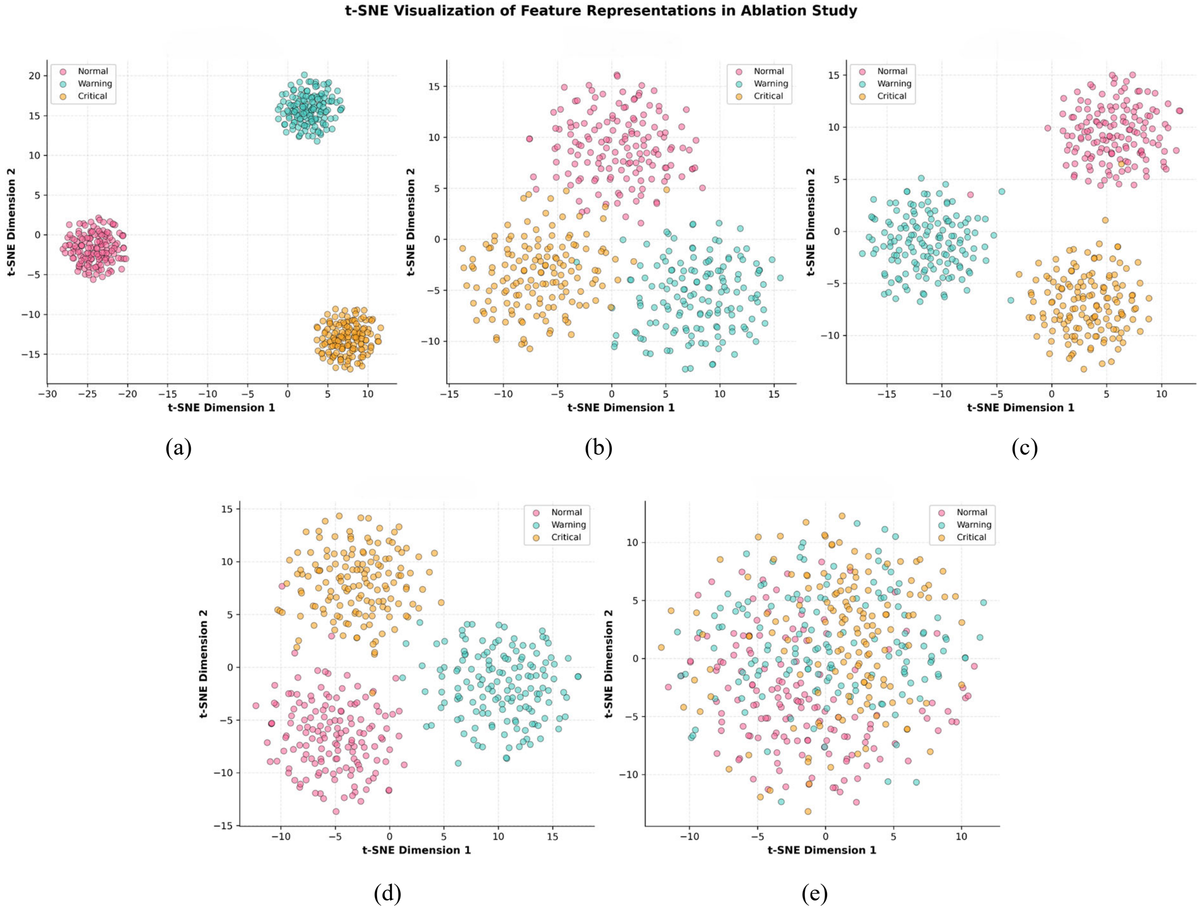

(a) (b) (c)

(d) (e)

Table 1 Computational cost comparison on MIMIC-III dataset

Method	*Training time (hours)*	*Inference latency (ms)*	*GPU memory (GB)*
LSTM	2.3	8.5	3.2
MTAN	4.7	12.3	5.8
TCN	3.1	9.7	4.1
GAT	5.2	15.6	6.5
Proposed method	4.8	14.2	6.1

4.3 Ablation experiment analysis

To deeply understand the contribution of each technical module to system performance, detailed ablation experiments were conducted. The experimental settings include the complete model, removal of adaptive attention fusion mechanism (w/o attention), removal of spatiotemporal graph neural network (w/o ST-GNN), removal of reinforcement learning module (w/o RL), and removal of joint optimisation mechanism (w/o joint). Figure 6 shows the t-SNE visualisation results of ablation experiments, intuitively revealing the class separation capabilities of different model variants in feature space by projecting high-dimensional feature representations to two-dimensional space. The complete model demonstrates excellent class separation, with normal, warning, and critical categories forming clear, compact, and well-defined clusters in feature space with almost no inter-class overlap. This good separability is directly reflected in its 96.3% anomaly detection accuracy and 0.912 F1-score, indicating that the complete model can learn highly discriminative feature representations that effectively distinguish different physiological states.

The feature space of w/o attention shows obvious class overlap, with blurred cluster boundaries among the three categories and numerous sample points distributed in inter-class overlap regions. Anomaly detection accuracy decreased from 96.3% to 91.8% and F1-score decreased from 0.912 to 0.847. Fixed-weight fusion cannot capture dynamic interactions and importance changes among physiological parameters, resulting in learned feature representations lacking discriminability, easily mapping samples of different states to similar feature regions in complex scenarios. Early warning time also shortened from 38.5 minutes to 29.3 minutes, and false alarm rate increased from 6.4% to 10.7%, further confirming the decline in recognition capabilities caused by feature confusion. W/o ST-GNN caused the most severe feature space confusion, with the three categories almost completely overlapping, making it difficult to form recognisable cluster structures. Anomaly detection accuracy dropped to 88.5%, early warning time shortened to 24.6 minutes, F1-score decreased to 0.806, and false alarm rate increased to 14.2%. Without the spatiotemporal graph structure, the model cannot effectively separate different physiological state patterns in feature space, and its ability to recognise abnormalities and hidden condition changes is greatly reduced, making the model unable to adaptively discover dynamic correlation patterns among parameters. The removal of the multi-scale time window mechanism weakens the perception of abnormalities at different time granularities; these factors jointly lead to severe feature space confusion and significant performance decline. The feature space of w/o RL also shows significant class confusion; although slightly better than w/o ST-GNN, it is still far inferior to the complete model. Anomaly detection accuracy dropped to 92.7%, early warning time shortened to 31.8 minutes, F1-score decreased to 0.871, and false alarm rate increased to 9.3%. The passive monitoring mode based on fixed rules leads to learned feature representations lacking decision orientation, resulting in reduced separability of different categories in feature space and prone to misjudgement in complex scenarios. The feature space of w/o joint also presents obvious class confusion; anomaly detection accuracy dropped to 93.4%, F1-score decreased to 0.883, and false alarm rate increased to 8.6%. Separate training leads to inconsistent optimisation objectives among modules, and learned intermediate feature representations cannot be effectively transferred, resulting in poor class separation in feature space.

To validate the contribution of each reward component, we conducted comprehensive ablation experiments on the reward balancing factors α_1, α_2, α_3, and α_4 defined in equation (15). Table 2 presents the performance variations under different coefficient configurations. The baseline with equal weights achieves moderate performance across all metrics: 94.1% ADA, 32.7 minutes EWT, 0.884 F1-score, and 8.3% FAR. When emphasising individual components, high α_1 yields the best detection accuracy at 95.8% ADA and 0.902 F1-score, while high α_2 reduces FAR to 5.1% but compromises accuracy to 93.2%. High α_3 produces the earliest warnings at 41.3 minutes but increases false alarms to 9.7%, whereas high α_4 degrades both accuracy to 92.8% and timeliness to 28.5 minutes. Complete ablation experiments confirm each component's necessity: removing α_1 causes the most severe degradation with 89.7% ADA and 0.831 F1-score, eliminating α_2 leads to unacceptable 15.3% FAR, ablating α_3 reduces early warning time to merely 22.9 minutes, and removing α_4 slightly increases FAR to 8.1%. Through systematic grid search on the validation set, our optimised configuration of $\alpha_1 = 0.40$, $\alpha_2 = 0.25$, $\alpha_3 = 0.20$, and $\alpha_4 = 0.15$ achieves superior overall performance: 96.3% ADA, 38.5 minutes EWT, 0.912 F1-score, and 6.9% FAR. These results demonstrate that careful balancing of reward components is crucial for learning effective monitoring policies that simultaneously optimise detection accuracy, warning timeliness, and clinical utility.

5 Discussion

The adaptive attention fusion mechanism proposed in this paper demonstrates significant advantages in multimodal physiological data integration. Traditional methods use fixed weights or simple feature concatenation for multimodal fusion, unable to adapt to dynamic changes in patient physiological states and time-varying characteristics of interactions among parameters. This paper achieves adaptive adjustment of fusion weights through gated network design, automatically learning the importance of different parameters according to patients' real-time physiological states, preserving unique features of each modality while effectively capturing interactions among modalities. Results show that the complete model improves anomaly detection accuracy by 4.5 percentage points and reduces false alarm rate by 4.3 percentage points, fully verifying the performance of the adaptive attention fusion mechanism in improving feature discriminability and reducing false alarms.

Table 2 Ablation study of reward balancing factors on MIMIC-III dataset

Configuration	α_1	α_2	α_3	α_4	*ADA (%)*	*EWT (min)*	*CEP-F1*	*FAR (%)*
Equal weights	0.25	0.25	0.25	0.25	94.1	32.7	0.884	8.3
High α_1	0.50	0.20	0.15	0.15	95.8	36.2	0.902	7.8
High α_2	0.20	0.50	0.15	0.15	93.2	31.4	0.869	5.1
High α_3	0.20	0.15	0.50	0.15	94.9	41.3	0.895	9.7
High α_4	0.20	0.15	0.15	0.50	92.8	28.5	0.871	6.4
Optimised (proposed)	0.40	0.25	0.20	0.15	96.3	38.5	0.912	6.9
No α_1 (ablated)	0.00	0.33	0.33	0.34	89.7	25.1	0.831	12.5
No α_2 (ablated)	0.50	0.00	0.25	0.25	95.1	35.8	0.897	15.3
No α_3 (ablated)	0.40	0.35	0.00	0.25	94.5	22.9	0.886	7.2
No α_4 (ablated)	0.43	0.29	0.28	0.00	95.7	37.1	0.905	8.1

The innovative design in spatiotemporal GNN architecture achieves deep joint modelling of spatial correlations and temporal evolution of physiological parameters, breaking through the limitation of traditional methods that separate time series analysis and multivariate correlation analysis. By modelling vital sign monitoring points as graph structure nodes, the proposed method can simultaneously capture interactions among parameters and temporal dependencies, using learnable adjacency matrices to adaptively discover data-driven parameter correlation patterns and integrating prior medical knowledge to construct more robust graph structure representations. The introduction of multi-scale time window mechanisms enables the system to identify abnormal patterns at different time granularities. Results show that the complete model improves critical event prediction F1-score by 10.6 percentage points and advances early warning time by 13.9 minutes, fully proving the critical role of spatiotemporal joint modelling in improving prediction performance and early identification capabilities.

The introduction of the deep reinforcement learning framework enhances the monitoring system's ability to intelligently optimise monitoring strategies. Traditional monitoring systems can only perform passive alarms according to preset rules, lacking intelligent support for clinical decisions, and unable to dynamically adjust monitoring strategies according to patient conditions. This paper models monitoring decisions as a Markov decision process, using spatiotemporal graph features as state inputs and utilising actor-critic networks to learn optimal strategies, enabling the system to autonomously determine key parameters such as monitoring frequency, warning thresholds, and intervention recommendations. The reward function design comprehensively considering monitoring accuracy and resource efficiency enables the reinforcement learning framework to optimise medical resource utilisation efficiency while improving warning performance. Experimental results show that compared to the model variant with removed reinforcement learning module, the complete model advances early warning time by 6.7 minutes and reduces false alarm rate by 2.9 percentage points, demonstrating the significant advantages of reinforcement learning in optimising monitoring decisions. Despite the promising results, several limitations should be acknowledged. The computational demands of our method, while manageable on modern GPUs with 4.8 hours of training time and 14.2 ms inference latency, may challenge hospital systems with limited resources. Furthermore, our evaluation primarily relied on two large-scale ICU datasets from developed healthcare systems, and generalisability to resource-constrained settings with different patient demographics and data characteristics requires further validation. Future work will focus on developing a lightweight version through model distillation and quantisation to facilitate edge deployment at the bedside, ensuring both practical integration into clinical workflows and the preservation of patient safety.

6 Conclusions

Addressing the challenges faced in patient vital sign monitoring in multimodal data fusion, spatiotemporal feature modelling, and active decision support, this paper proposes an intelligent monitoring method based on adaptive attention fusion spatiotemporal graph neural networks. First, an adaptive attention fusion mechanism is constructed that adjusts fusion weights of different physiological parameters through gated networks, automatically learning optimal combination methods of temporal and spatial features according to patients' real-time physiological states, achieving adaptive integration of multimodal data. Second, a spatiotemporal graph neural network architecture is designed that simultaneously captures correlations among physiological parameters and temporal change patterns through graph convolution operations, utilising multi-scale time window mechanisms to achieve deep modelling of spatiotemporal information. Finally, a deep reinforcement learning framework is introduced, using spatiotemporal graph features as state inputs and learning optimal monitoring strategies through actor-critic networks, achieving active decision support. Comprehensive experiments on two large-scale ICU datasets, MIMIC-III and eICU, demonstrate that our method achieves 96.3% in anomaly detection accuracy, an improvement of 4.6 percentage points compared to existing optimal methods, advances early warning time by 38.5 minutes, an advance of 10.1 minutes compared to sub-optimal methods, achieves a critical event prediction F1-score of 0.912, reduces false alarm rate to 6.4%, and maintains stable high performance under different prediction scenarios, fully validating the proposed method's improvement in patient vital sign intelligent monitoring accuracy.

Declarations

All authors declare that they have no conflicts of interest.

References

Celi, L.A., Mark, R.G., Stone, D.J. and Montgomery, R.A. (2013) ''Big data' in the intensive care unit. closing the data loop', *American Journal of Respiratory and Critical Care Medicine*, Vol. 187, No. 11, pp.1157–1160, DOI: 10.1164/rccm.201212-2311ED.

Che, Z., Purushotham, S., Cho, K., Sontag, D. and Liu, Y. (2018) 'Recurrent neural networks for multivariate time series with missing values', *Scientific Reports*, Vol. 8, No. 1, p.6085, DOI: 10.1038/s41598-018-24271-9.

Choi, E., Bahadori, M.T., Song, L., Stewart, W.F. and Sun, J. (2017) 'GRAM: graph-based attention model for healthcare representation learning', *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*, Association for Computing Machinery, New York, NY, USA, pp.787–795, DOI: 10.1145/3097983.3098126.

Choi, E., Bahadori, M.T., Sun, J., Kulas, J., Schuetz, A. and Stewart, W. (2016) 'RETAIN: an interpretable predictive model for healthcare using reverse time attention mechanism', *Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS'16)*, Curran Associates Inc., Red Hook, NY, USA, pp.3512–3520.

Choi, E., Xu, Z., Li, Y., Dusenberry, M., Flores, G., Xue, E. and Dai, A. (2020) 'Learning the graphical structure of electronic health records with graph convolutional transformer', *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34, No. 01, pp.606–613, DOI: 10.1609/aaai.v34i01.5400.

Clifford, G.D., Long, W.J., Moody, G.B. and Szolovits, P. (2008) 'Robust parameter extraction for decision support using multimodal intensive care data', *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Vol. 367, No. 1887, pp.411–429, DOI: 10.1098/rsta.2008.0157.

Clifton, D.A., Gibbons, J., Davies, J. and Tarassenko, L.(2012) 'Machine learning and software engineering in health informatics', *2012 First International Workshop on Realizing AI Synergies in Software Engineering (RAISE)*, IEEE, New York, NY, USA, pp.37–41, DOI: 10.1109/RAISE.2012.6227968.

Futoma, J., Hariharan, S. and Heller, K. (2017) 'Learning to detect sepsis with a multitask Gaussian process RNN classifier', *Proceedings of the 34th International Conference on Machine Learning*, PMLR, New York, NY, USA, pp.1174–1182.

Ghassemi, M., Pimentel, M., Naumann, T., Brennan, T., Clifton, D., Szolovits, P. and Feng, M. (2015) 'A multivariate timeseries modeling approach to severity of illness assessment and forecasting in ICU with sparse, heterogeneous clinical data', *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29, No. 1, pp.446–453, DOI: 10.1609/aaai.v29i1.9209.

Harutyunyan, H., Khachatrian, H., Kale, D.C., Ver Steeg, G. and Galstyan, A. (2019) 'Multitask learning and benchmarking with clinical time series data', *Scientific Data*, Vol. 6, No. 1, p.96, DOI: 10.1038/s41597-019-0103-9.

Hravnak, M., Edwards, L., Clontz, A., Valenta, C., Devita, M.A. and Pinsky, M.R. (2008) 'Defining the incidence of cardiorespiratory instability in patients in step-down units using an electronic integrated monitoring system', *Archives of Internal Medicine*, Vol. 168, No. 12, pp.1300–1308, DOI: 10.1001/archinte.168.12.1300.

Johnson, A.E.W., Ghassemi, M.M., Nemati, S., Niehaus, K.E., Clifton, D.A. and Clifford, G.D. (2016a) 'Machine learning and decision support in critical care', *Proceedings of the IEEE*, Vol. 104, No. 2, pp.444–466, DOI: 10.1109/JPROC.2015.2501978.

Johnson, A.E.W., Pollard, T.J., Shen, L., Lehman, L.H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L.A. and Mark, R.G. (2016b) 'MIMIC-III, a freely accessible critical care database', *Scientific Data*, Vol. 3, No. 1, p.160035, DOI: 10.1038/sdata.2016.35.

Kellett, J. and Sebat, F. (2017) 'Make vital signs great again – a call for action', *European Journal of Internal Medicine*, Vol. 45, pp.13–19, DOI: 10.1016/j.ejim.2017.09.018.

Lipton, Z.C., Kale, D.C., Elkan, C. and Wetzel, R. (2017) *Learning to Diagnose with LSTM Recurrent Neural Networks*, *arXiv*, DOI: 10.48550/arXiv.1511.03677.

Ma, F., Chitta, R., Zhou, J., You, Q., Sun, T. and Gao, J. (2017) 'Dipole: diagnosis prediction in healthcare via attention-based bidirectional recurrent neural networks', *KDD 2017 – Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, New York, NY, USA, pp.1903–1911, DOI: 10.1145/3097983.3098088.

Ma, F., You, Q., Xiao, H., Chitta, R., Zhou, J. and Gao, J. (2018) 'KAME: knowledge-based attention model for diagnosis prediction in healthcare', *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*, Association for Computing Machinery, New York, NY, USA, pp.743–752, DOI: 10.1145/3269206.3271701.

Moorman, J.R., Carlo, W.A., Kattwinkel, J., Schelonka, R.L., Porcelli, P.J., Navarrete, C.T., Bancalari, E., Aschner, J.L., Walker, M.W., Perez, J.A., Palmer, C., Stukenborg, G.J., Lake, D.E. and O'Shea, T.M. (2011) 'Mortality reduction by heart rate characteristic monitoring in very low birth weight neonates: a randomized trial', *The Journal of Pediatrics*, Vol. 159, No. 6, pp.900–906.e1, DOI: 10.1016/j.jpeds.2011.06.044.

Purushotham, S., Meng, C., Che, Z. and Liu, Y. (2018) 'Benchmarking deep learning models on large healthcare datasets', *Journal of Biomedical Informatics*, Vol. 83, pp.112–134, DOI: 10.1016/j.jbi.2018.04.007.

Rajpurkar, P., Hannun, A.Y., Haghpanahi, M., Bourn, C. and Ng, A.Y. (2017) *Cardiologist-Level Arrhythmia Detection with Convolutional Neural Networks*, arXiv, DOI: 10.48550/arXiv.1707.01836.

Rajpurkar, P., Irvin, J., Ball, R.L., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C.P., Patel, B.N., Yeom, K.W., Shpanskaya, K., Blankenberg, F.G., Seekins, J., Amrhein, T.J., Mong, D.A., Halabi, S.S., Zucker, E.J., Ng, A.Y. and Lungren, M.P. (2018) 'Deep learning for chest radiograph diagnosis: a retrospective comparison of the CheXNeXt algorithm to practicing radiologists', *PLOS Medicine*, Vol. 15, No. 11, p.e1002686, DOI: 10.1371/journal.pmed.1002686.

Razavian, N., Marcus, J. and Sontag, D. (2016) 'Multi-task prediction of disease onsets from longitudinal laboratory tests', *Proceedings of the 1st Machine Learning for Healthcare Conference*, PMLR, New York, NY, USA, pp.73–100.

Shang, J., Xiao, C., Ma, T., Li, H. and Sun, J. (2019) 'GAMENet: graph augmented memory networks for recommending medication combination', *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, No. 01, pp.1126–1133, DOI: 10.1609/aaai.v33i01.33011126.

Shickel, B., Tighe, P.J., Bihorac, A. and Rashidi, P. (2018) 'Deep EHR: a survey of recent advances in deep learning techniques for electronic health record (EHR) analysis', *IEEE Journal of Biomedical and Health Informatics*, Vol. 22, No. 5, pp.1589–1604, DOI: 10.1109/JBHI.2017.2767063.

Song, H., Rajan, D., Thiagarajan, J. and Spanias, A. (2018) 'Attend and diagnose: clinical time series analysis using attention models', *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32, No. 1, pp.4091–4098, DOI: 10.1609/aaai.v32i1.11635.

Subbe, C.P., Kruger, M., Rutherford, P. and Gemmel, L. (2001) 'Validation of a modified early warning score in medical admissions', *QJM: An International Journal of Medicine*, Vol. 94, No. 10, pp.521–526, DOI: 10.1093/qjmed/94.10.521.

Tarassenko, L., Hann, A. and Young, D. (2006) 'Integrated monitoring and analysis for early warning of patient deterioration', *BJA: British Journal of Anaesthesia*, Vol. 97, No. 1, pp.64–68, DOI: 10.1093/bja/ael113.

Xu, Y., Biswal, S., Deshpande, S.R., Maher, K.O. and Sun, J. (2018) 'RAIM: recurrent attentive and intensive model of multimodal patient monitoring data', *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '18)*, Association for Computing Machinery, New York, NY, USA, pp.2565–2573, DOI: 10.1145/3219819.3220051.