AIMS Mathematics, 11(5): 13174–13195.
DOI: 10.3934/math.2026543
Received: 06 February 2026
Revised: 29 March 2026
Accepted: 07 April 2026
Published: 12 May 2026

https://www.aimspress.com/journal/Math

Research article

Deep generative neural network dynamics modeling and employee satisfaction prediction for human resource management

Jie Wang*

School of Accounting, Hangzhou College of Commerce, Zhejiang Gongshang University, Hangzhou 310000, China

* **Correspondence:** Email: wangjie198207@163.com.

Abstract: Traditional human resource management has limited ability to capture the temporal evolution of employee satisfaction and to learn informative representations from high-dimensional, multi-source data, thereby constraining the accuracy and interpretability of dynamic prediction. To address this issue, I developed a joint VAE-Neural ODE framework that integrates a variational autoencoder (VAE) with a neural ordinary differential equation (neural ODE). Specifically, the VAE encoder employs the reparameterization trick to learn the latent probability distribution of heterogeneous employee records and generate compact yet informative latent features. On this basis, the neural ODE models the continuous-time evolution of latent states and propagates them through a differentiable ODE solver for temporal forecasting. The decoder then maps the evolved latent state to a continuous employee satisfaction score normalized to the interval [0, 1]. Accordingly, the prediction task was formulated primarily as a regression problem, and model performance was evaluated using mean absolute error (MAE) and root mean squared error (RMSE). For interpretive comparison, the continuous outputs were further discretized using a fixed threshold to derive a binary accuracy metric. Experimental results showed that the proposed method achieves an accuracy of 94.7%, an MAE of 0.031, and an RMSE of 0.052 on the test set, while maintaining 87.2% accuracy in cross-quarter forecasting. These findings indicate that the proposed framework can effectively capture nonlinear satisfaction dynamics and provide quantitative support for data-driven human resource analysis. Moreover, by formulating employee satisfaction prediction as a latent continuous-time dynamical process and evolving hidden states through an ODE flow rather than purely discrete recursive updates, the framework produces smoother latent trajectories and reduces error accumulation over longer forecasting horizons and under irregular observation intervals, which are common in non-uniformly sampled HR records.

Keywords: human resource management; employee satisfaction prediction; variational autoencoder; neural ordinary differential equation; dynamic modeling
Mathematics Subject Classification: 34L99, 68T07

1. Introduction

Human resource management, as a critical component of modern organizational governance, is closely associated with employee performance, organizational efficiency, and the long-term development of enterprises. With the advancement of data-driven management and intelligent analytics, organizations have accumulated large volumes of multi-source employee data, including information on behavior, performance, attendance, and training, thereby providing an important foundation for understanding employee satisfaction and its dynamic evolution. Studies have shown that corporate social responsibility and internal corporate policy are closely related to employee satisfaction and organizational outcomes [1,2], while sustainability and ethical practices provide complementary evidence for satisfaction-oriented management [3]. Employee satisfaction not only reflects employees' identification with the organization and their degree of work engagement, but is also strongly linked to turnover intention, work efficiency, and innovation capability. Text-mining, workplace-health, and ESG-based studies further indicate that employee satisfaction is associated with corporate performance, productivity, costs, and sustainable organizational behavior [4–6].

Despite its importance, research on employee satisfaction faces several major challenges. Employee satisfaction exhibits complex and nonlinear dynamic characteristics and is influenced by multiple factors, including the organizational environment, managerial practices, and employees' individual psychological states. Empirical studies have revealed workplace factors and managerial conditions as important determinants of job satisfaction and employee performance [7,8]. In addition, employee-related data are typically high-dimensional, heterogeneous, and partially incomplete, while data collected from different sources often differ in temporal scale and feature space, which further increases modeling complexity. Machine-learning and HR data-modeling studies have shown that such heterogeneous feature structures make robust satisfaction prediction difficult [9,10]. Moreover, the time-series nature of employee satisfaction requires models capable of capturing continuous-time evolution. Deep generative, deep neural, and optimized neural-network approaches have been applied to job-satisfaction prediction, but their ability to model temporal dependencies under irregular sampling and asynchronous events remains limited [11–13].

Researchers have attempted to address these challenges using a variety of approaches. Traditional regression and factor-analysis methods examine employee satisfaction through linear or nonlinear associations, yet they are not well suited to high-dimensional, multi-source data or complex temporal evolution. BiLSTM-ANN and Bi-LSTM approaches have shown promise in employee satisfaction analysis and employee churn prediction, but their capacity to generate informative latent representations remains limited, and their robustness to missing or noisy data is often insufficient [14,15]. Graph neural networks and dynamic context-weighted embeddings have demonstrated advantages in modeling relational or contextual structures, yet they face limitations in continuous-time dynamic prediction [16,17]. Recurrent-network-based strategic HR decision models and generative-AI HR applications further highlight the need for methods that jointly support high-dimensional feature generation, continuous-time dynamic modeling, and robust learning under

incomplete data conditions [18,19].

To address the limited ability of approaches to model continuous-time satisfaction dynamics and to learn informative latent representations from high-dimensional, multi-source HR data, I propose a joint VAE-neural ODE framework for employee satisfaction prediction. Specifically, heterogeneous employee records are first projected into a latent space through a variational encoder, where the reparameterization trick is applied to learn the latent-variable distribution and derive compact yet expressive representations. On this basis, a neural ODE is introduced to characterize the continuous-time evolution of latent states, while a differentiable ODE solver is employed to propagate latent trajectories for forecasting. In this way, employee satisfaction prediction is reformulated as a continuous-time latent dynamical process driven by multi-source HR observations, moving beyond conventional fixed-step discrete sequence modeling. By coupling distributional representation learning with ODE-based state evolution, the proposed framework unifies latent feature extraction, continuous-time dynamical modeling, reconstruction, and prediction within an end-to-end architecture, thereby improving robustness to heterogeneous high-dimensional inputs and partial missingness. Furthermore, through empirical analyses of trajectory smoothness, step-size sensitivity, and long-horizon generalization, this study demonstrates that the continuous-time formulation can mitigate error accumulation in cross-quarter forecasting and provides a principled technical pathway for data-driven prediction and organizational human resource analytics.

2. Model construction and implementation

2.1. *Data feature representation and preprocessing*

The dynamic prediction of employee satisfaction relies on the effective representation of multi-source heterogeneous data, which covers multi-dimensional inputs such as performance indicators, job information, attendance logs, behavioral sequences, and questionnaire responses. The raw data contain long time-series records, high dimensionality, and substantial acquisition noise. Deep nonlinear state-space forecasting, non-stationary time-series modeling, and kernel-filter/time-attention forecasting studies show that missing-value treatment, normalization, and temporal alignment are necessary to stabilize dynamic learning [20–22]. Before modeling begins, missing information must be corrected, standardized, and timestamp aligned at the input stage to ensure consistent representation of data from different sources in the latent space.

Let the raw data matrix be:

$$X=\{x_{i,t}|i=1,\ldots,N;\ t=1,\ldots,T\}. \tag{1}$$

In formula (1), $x_{i,t}$ represents the multidimensional feature vector of the i-th employee at the time step t, N is the total number of employees, and T is the time span. Since the dimensions of different dimensional features are inconsistent, normalization operation is required, and the calculation method is:

$$\hat{x}_{i,t}^{(d)}=\frac{x_{i,t}^{(d)}-\mu^{(d)}}{\sigma^{(d)}}. \tag{2}$$

In formula (2), $\hat{x}_{i,t}^{(d)}$ represents the normalized eigenvalue of the d-th dimension, and $\mu^{(d)}$and

$\sigma^{(d)}$ represent the mean and standard deviation of the dimension feature over all samples and time, respectively. This process ensures that different dimensions maintain a uniform distribution scale when entering the latent space.

Missing features are imputed using a combination of time-series interpolation and local mean backfilling to ensure smoothness in the continuous-time domain. Multi-source data sampled at different clocks, denoted by $\{t_1,t_2,\ldots,t_T\}$, are resampled and aligned with a unified timestamp set to form a consistent input tensor $\widehat{X}$. This process ensures that the input data forms a stable trajectory in the latent space, thereby maintaining the computational accuracy of the neural ODE during continuous-time evolution.

To visually demonstrate the scale and complexity of the data, I conduct a comprehensive statistical and quantitative analysis of the raw features. Table 1 presents the key characteristics of the multi-source employee data used in the experiment, including sample size, time span, feature dimension, data density, and missing rate. These metrics reflect the quality and challenge of the data before input.

Table 1. Statistics of original employee multi-source data.

Data category	Sample size	Time span	Feature dimensions	Average sequence length	Missing rate	Data density
Core Employee Performance	512	12 months	18	320	2.9%	94.6%
Position and Compensation	512	12 months	10	320	1.5%	98.3%
Attendance and Work Hours	498	12 months	15	305	4.1%	91.7%
Behavioral Log Sequence	472	12 months	20	298	6.8%	87.9%
Survey Satisfaction Feedback	460	12 months	5	280	7.5%	85.4%

Table 1 shows the scale and quality characteristics of different data sources. Core performance and job-related data have the lowest missing rates and the highest data density, indicating the most stable structured attributes. Behavioral logs and questionnaire responses have relatively high missing rates, reflecting some discontinuities in the dynamic collection process. After preprocessing, the overall data integrity is significantly improved, providing a reliable input foundation for subsequent VAE latent distribution learning and neural ODE dynamic modeling.

After normalizing and temporally aligning the input features, the overall preprocessing process is shown in Figure 1. The framework demonstrates the complete path from raw data input to latent distribution generation, continuous dynamics modeling, and satisfaction prediction.

Figure 1 illustrates the process by which data enters the latent space via the encoder, is coupled with the continuous-time evolution of neural ODEs through probability distribution generation, and finally outputs predictions by the decoder. The diagram highlights the key connections between latent distribution generation and temporal dynamics modeling, emphasizing the model's closed-loop prediction capability in the continuous-time domain.

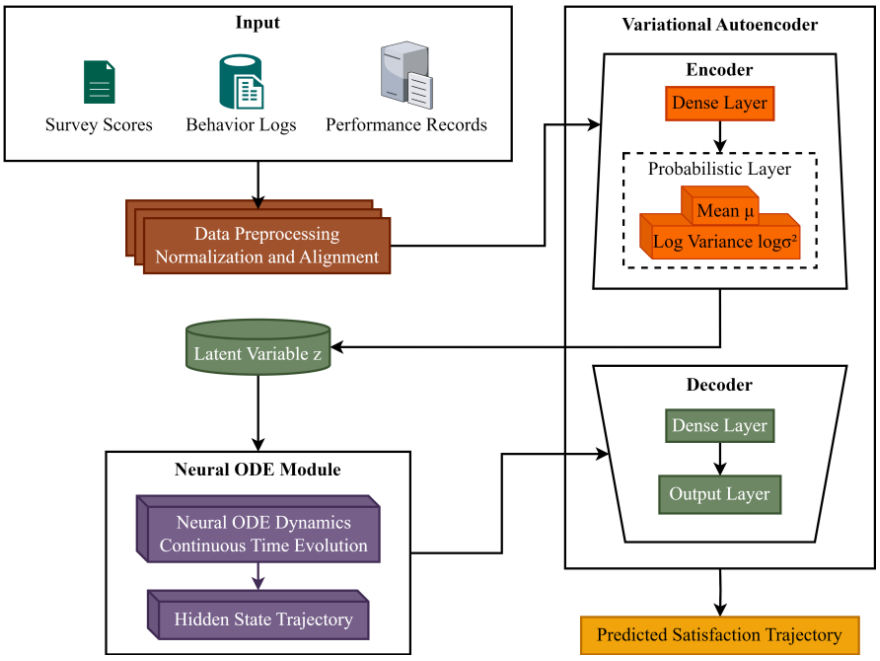

Figure 1. Framework diagram of the VAE-Neural ODE method.

2.2. Variational autoencoder-based latent distribution learning

For standardized and time-aligned data tensors, the model constructs a latent distribution to represent the probabilistic evolution of employee satisfaction. A factor-bridging job-satisfaction prediction method and an online-review-based analysis of employee personality, satisfaction, and turnover indicate that employee status should be represented through multi-factor and probabilistic information [23,24]. Because employee status is influenced by heterogeneous multi-source information, its latent characteristics exhibit high-dimensional complexity and uncertainty. VAE-based recommendation, psychological-data modeling, and heterogeneous graph recommendation studies demonstrate that variational autoencoders can learn nonlinear latent distributions from complex data [25–27]. Modeling the latent distribution with a VAE therefore unifies data generation and representation learning.

Encoder $\widehat{X}$ maps the input features to z, the conditional distribution of the latent variables. Through the neural network structure, the encoder outputs the latent mean vector μ_z and the logarithmic variance vector $\log \sigma_z^2$ to form a latent Gaussian distribution:

$$q_\phi(z|\widehat{X})=N(z;\mu_z,\mathrm{diag}(\sigma_z^2)). \tag{3}$$

In formula (3), ϕ represents the set of encoder parameters and σ_z is the standard deviation of the latent vector. This distribution describes the probability representation of the input features in the latent space.

Since gradient propagation cannot be performed by sampling directly from the distribution, the reparameterization technique is introduced to transform the sampling process into

$$z=\mu_z+\sigma_z\odot\epsilon,\ \epsilon\sim N(0,I). \tag{4}$$

In formula (4), $\odot$ is an element-by-element multiplication operation and ϵ is a noise vector that follows a standard normal distribution. This approach makes latent-variable sampling a differentiable operation, thereby ensuring end-to-end training.

Decoder z performs an inverse mapping on the latent variables and reconstructs the input features $\widehat{X}$. Its conditional probability distribution is defined as:

$$p_\theta(\widehat{X}|z)=\prod_{i=1}^{N}\prod_{t=1}^{T}\prod_{d=1}^{D} N(\hat{x}_{i,t}^{(d)};f_\theta^{(d)}(z_{i,t}),\sigma_\epsilon^2). \tag{5}$$

In formula (5), $z_{i,t}$ is the corresponding latent variable, $f_\theta^{(d)}(\cdot)$ represents the decoder's function mapping the latent variable to the d-th feature dimension, σ_ϵ^2 is the assumed observation noise variance, N is the total number of employees, T is the total number of time steps, and D is the feature dimension. This formula embodies the probabilistic model by which the decoder generates observation features from latent variables, clarifies the mapping mechanism from latent space to observation space, and provides a foundation for ELBO (Evidence Lower Bound) optimization.

This study uses ELBO to optimize the objective, which ensures that the latent distribution can approximate the true data distribution and remain continuously differentiable:

$$L_{VAE}(\theta,\phi;\widehat{X})=E_{q_\phi(z|\widehat{X})}[\log p_\theta(\widehat{X}|z)]-D_{KL}(q_\phi(z|\widehat{X})\|p(z)). \tag{6}$$

In formula (6), the first term is the reconstructed log-likelihood expectation, which ensures that the data generated by the decoder is as close to the original input as possible; the second term is the Kullback-Leibler divergence, which constrains the difference between the posterior distribution $q_\phi(z|\widehat{X})$ and the prior distribution $p(z)=N(0,I)$, thereby maintaining the regularization and structural stability of the latent space.

Table 2 summarizes the network structure and key training parameters of the VAE, making the modeling process transparent.

Table 2. VAE network structure and training parameters.

Module	Layer structure	Activation function	Output dimension	Optimizer	Learning rate	Batch size
Encoder	Fully Connected Layer → Fully Connected Layer → Output mean & variance	ReLU	128→64→(μ,σ)	Adam	1e-3	64
Latent Space	Reparameterization Sampling	–	16	–	–	–
Decoder	Fully Connected Layer → Fully Connected Layer → Reconstructed Output	ReLU	64→128→68	Adam	1e-3	64

Table 2 shows the encoder and decoder hierarchical structure and training configuration. The

encoder compresses the input high-dimensional features through layer-by-layer dimensionality reduction and outputs the latent mean and variance. The latent space dimension is set to 16 to ensure that characteristics related to employee satisfaction are retained while compressing information. The decoder performs the inverse mapping from the latent variables to the original feature space, providing a stable latent trajectory input for subsequent neural ODE continuous dynamics modeling.

2.3. Neural ODE continuous-time dynamics modeling

After the latent distribution representation is obtained, it is introduced into the continuous-time domain for state evolution modeling to characterize the dynamic changes in employee satisfaction. Traditional discrete recursive models assume a fixed time step, whereas the evolution of employee satisfaction often exhibits continuous, nonlinear characteristics. Comparative research on neural ODEs and neural operators, time-aware neural ODEs for incomplete time series, and latent ODEs with graph RNNs for multi-channel forecasting has shown that ODE-based latent dynamics can capture continuous temporal evolution and irregular observations [28–30]. Therefore, the neural ODE module is introduced to enable the latent state to evolve naturally over time in the continuous time domain.

The latent state vector is denoted as $z(t)$, and its dynamic process is described by the following differential equation:

$$\frac{dz(t)}{dt}=f_\theta(z(t),t). \tag{7}$$

In formula (7), f_θ is a nonlinear function defined by the neural network parameters θ that is used to characterize the change pattern of the latent state in the time dimension.

This function is implemented through a multi-layer perceptron, making the state evolution highly expressive.

From a dynamical systems perspective, this formulation learns a time-dependent vector field that governs the trajectory of the latent state. The continuity of the ODE solver enforces a smooth latent transition, acting as a strong regularizer that prevents abrupt jumps in the latent state, which are often observed in discrete-time models like RNNs. This continuous-time prior is particularly beneficial for forecasting over long or irregularly sampled time horizons as it decouples the model's internal dynamics from the specific times at which observations are made, thereby improving generalization to unseen temporal patterns.

For a given initial latent variable $z(t_0)$, the change in state over continuous time intervals $[t_0,t_1]$ is computed by an ordinary differential equation solver:

$$z(t_1)=z(t_0)+\int_{t_0}^{t_1} f_\theta\,(z(t),t)\,dt. \tag{8}$$

In formula (8), the integral part is implemented using a differentiable numerical solver, enabling end-to-end training. This mechanism enables the latent state to evolve smoothly in continuous time, consistent with the actual dynamic trend in employee satisfaction.

The advantage of neural ODEs is that they are not restricted by a fixed step size and can capture latent states at any point in time $z(t)$. This supports satisfaction prediction across quarters or non-uniform time intervals. Furthermore, continuous-time modeling ensures the stability of latent

trajectories over long time spans, avoiding the problem of accumulated errors associated with discrete recursion.

After introducing the neural ODE, the dynamic trajectory of the latent state is passed to the decoder, yielding continuous predictions of employee satisfaction. In this paper, I compare the actual satisfaction curve with the model-predicted curve, as shown in Figure 2.

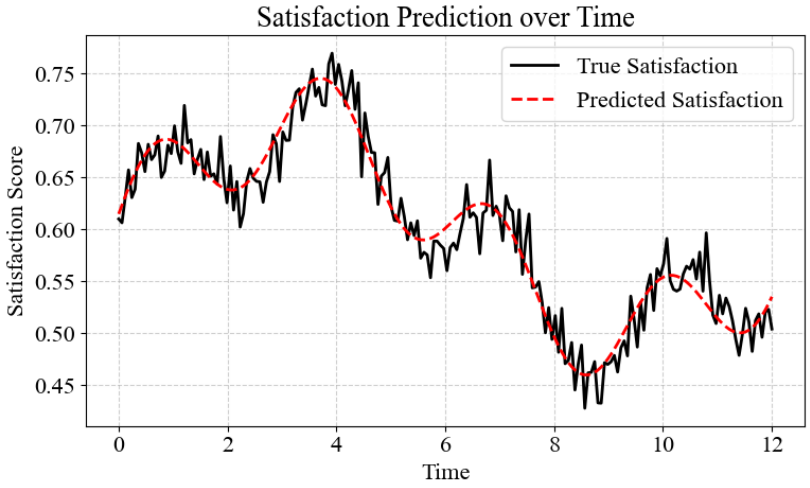

Figure 2. Schematic diagram of employee satisfaction prediction time series fitting.

Figure 2 shows a dynamic comparison of the true satisfaction curve and the model-predicted curve in the continuous time domain. The true curve shows a steady upward trend with periodic fluctuations, reflecting the changing characteristics of employee satisfaction influenced by long- and short-term factors. The predicted curve closely matches the true trajectory in the time series, showing a consistent trend and smoothness in local details, demonstrating the stability and generalization ability of the neural ODE in capturing nonlinear dynamic evolution. This result demonstrates that the model suppresses noise fluctuations while maintaining overall trend consistency, thereby ensuring the reliability of satisfaction prediction at the time series evolution level.

2.4. Joint optimization and end-to-end training

After constructing the VAE latent distribution learning module and the neural ODE continuous-time dynamics modeling module, end-to-end training is achieved through a joint optimization mechanism to ensure the coordinated convergence of the latent distribution generation capabilities and the temporal dynamics evolution capabilities. The key to joint optimization lies in designing a suitable loss function that optimizes latent space reconstruction and continuous-time prediction within the same framework.

The overall loss function consists of three parts: reconstruction error, KL (Kullback-Leibler) divergence regularization term, and temporal prediction error. It is defined as:

$$L_{total}=E_{q_\phi(z|\widehat{X})}[\|\widehat{X}-\widetilde{X}\|_2^2]+\beta D_{KL}(q_\phi(z|\widehat{X})\|p(z))+\lambda E_t[\|y_t-\hat{y}_t\|_2^2]. \quad (9)$$

In formula (9), the first term is the mean square error between the input features and the decoder reconstruction results $\tilde{X}$, which ensures that VAE can accurately restore the multi-source features of employees from the latent space; the second term is the KL divergence between the latent posterior distribution and the standard normal prior, and coefficient β is used to adjust the regularization strength of the latent space; and the third term is the continuous time prediction error, y_t, which represents the true satisfaction value, and $\hat{y}_t$ is the predicted value obtained by evolving and decoding through neural ODE.

Coefficient λ is used to balance the optimization goals of feature reconstruction and time prediction. The overall objective function comprises three terms: reconstruction loss, KL divergence, and trajectory prediction loss. To balance feature reconstruction, latent-space regularization, and continuous-time prediction, the corresponding weights are fixed at 2:0.2:1, respectively, in all experiments.

During joint training, the parameters θ, ϕ of the VAE and the parameters θ_{ODE} of neural ODE are updated synchronously via end-to-end backpropagation. Since the temporal evolution of the neural ODE is implemented using a differentiable numerical solver, the gradient is computed using the adjoint sensitivity method during backpropagation, ensuring stable parameter updates in the continuous time domain. This process ensures the coordinated optimization of latent distribution learning and temporal dynamics modeling, preventing the situation where one module converges while the other diverges.

The optimization strategy employs the Adam optimizer together with a cosine annealing learning-rate schedule, which promotes rapid convergence in the early training stage while preserving a smooth and stable decline in the later stage. This end-to-end training framework not only ensures high-quality representation in the latent space, but also supports closed-loop modeling of temporal evolution trajectories and dynamic prediction of employee satisfaction.

The joint optimization results demonstrate improved convergence speed and prediction stability, with the loss function showing a consistent downward trend over multiple iterations, thereby confirming the effective synergy between the VAE and neural ODE modules.

This end-to-end framework enables the model to simultaneously learn informative feature representations and capture temporal dynamics, providing a solid basis for subsequent experimental validation. Moreover, the multi-term objective admits a clear dynamical interpretation: The reconstruction term ensures that the latent variables faithfully encode heterogeneous HR signals, while the trajectory prediction term constrains the latent state to evolve along a coherent continuous-time flow that is consistent with observed satisfaction dynamics. The KL divergence term further prevents posterior collapse and stabilizes the geometry of the latent space, which is particularly important because the ODE module extrapolates along latent directions that may not be densely sampled. In our experiments, balancing these objective terms yields a latent representation that is descriptive for reconstruction and predictive for temporal evolution, thereby enhancing robustness in long-horizon forecasting.

3. Experimental design and implementation

3.1. Dataset composition and sampling scheme

The data used in the experiment comes from a real-world internal human resources information system, encompassing multi-dimensional, continuous records generated by employees throughout

their tenure. The dataset contains records of 512 employees after filtering incomplete profiles, and each employee is associated with a multivariate time series composed of 68 features derived from five categories: Performance metrics, attendance logs, behavioral records, and job information, with the satisfaction label collected from quarterly surveys. The data spans a full three-year period and includes information on basic employee attributes, attendance behavior, performance appraisals, training participation, and satisfaction scales collected through questionnaires. These records are accurately timestamped, providing continuous, dynamic data for constructing a time-series prediction model. The satisfaction score is obtained from standardized internal surveys where each employee reports their perceived job satisfaction on a normalized scale ranging from zero to one, ensuring consistency across time periods.

The 68 input features are grouped as follows: Performance metrics include quarterly sales volume, project completion rate, and supervisor evaluation score; attendance logs include monthly average overtime hours, number of absences, and tardiness frequency; behavioral records include peer feedback score, training participation count, and internal communication activity; and job information includes position level, years of tenure, and department code. The satisfaction score is a continuous variable ranging from 0 to 1, derived from the average of five Likert-scale survey questions on job contentment, work environment, and career development prospects. The distribution of satisfaction scores across the dataset is unimodal, with a mean of 0.72 and a standard deviation of 0.18, indicating slight positive skew. To quantify label noise, I compute the Cronbach's alpha for the survey questions at each time point, which consistently remains above 0.82, indicating high internal consistency.

The raw data maintains complete longitudinal tracking at the individual employee level. The lengths of each employee's data series vary, requiring alignment via timestamps during preprocessing. To ensure fairness and predictive stability, all employee data is segmented into quarterly time series. This maintains temporal evolution while mitigating bias caused by individual data imbalances. The sampling scheme extracts observation points at fixed time intervals, aligning discrete satisfaction scores with the multi-source features corresponding to the same time period, thereby establishing a one-to-one correspondence between input features and satisfaction labels. During the data partitioning process, the training set, validation set, and test set are strictly distinguished in the time dimension to avoid information leakage and examine the cross-time series generalization performance of the model so that the subsequent VAE-Neural ODE joint modeling can achieve stable prediction of the dynamic evolution process of employee satisfaction based on continuous time input.

3.2. Experimental environment and parameter settings

The experiments are conducted in a standard commercial computing environment. Table 3 presents the complete configuration of the experimental environment and key hyperparameters, facilitating replication and deployment on common hardware.

The hardware platform is a workstation equipped with an Intel Core i7-12700 processor at 2.1 GHz, 32 GB of RAM, and an NVIDIA GeForce RTX 3060 12 GB graphics processing unit (GPU). This meets the video memory and parallel computing requirements for deep generative network training. The operating system used is Windows 11 64-bit; the deep learning framework is PyTorch 2.2.0; the Python version is 3.10; and the CUDA driver version is 11.8. This environment facilitates the replication of experimental results on a standard high-performance personal computer. The random number seed is set to 42 to ensure reproducibility of the results.

The joint optimization of VAE and neural ODE focuses on stable convergence. The VAE learning rate is set to 1×10^{-3}, the batch size is 64, the latent space dimension is 16, and the reparameterization noise variance is 0.01. The optimizer uses Adam with a weight decay coefficient of 1×10^{-5}. The neural ODE module has a time step of 0.05, an adaptive Runge–Kutta 4th/5th order differentiable solver is used, the hidden layer dimension is 64, and the activation function uses Swish to enhance the nonlinear modeling of continuous-time dynamics.

Table 3. Experimental environment and hyperparameter configuration table.

Category	Specification and Version	Parameter	Value
Hardware	Intel Core i7-12700, 32 GB RAM, NVIDIA GeForce RTX 3060 12 GB	Learning rate	1×10^{-3}
Software	Windows 11 64-bit, Python 3.10, PyTorch 2.2.0, CUDA 11.8	Batch size	64
VAE module	—	Latent dimension	16
Neural ODE	—	Time step	0.05
Optimization	Adam, weight decay 1×10^{-5}	Epochs	200
Regularization	L2	Coefficient	0.001

The solver uses an adaptive error control mechanism with relative tolerance set to 1e-5 and absolute tolerance set to 1e-6, and dynamically adjusts the integration interval based on local truncation error to ensure numerical stability during latent state evolution. The joint loss function comprises reconstruction loss, KL divergence, and trajectory prediction loss, with weights set to 2:0.2:1, respectively. In addition, L2 regularization with a coefficient of 0.001 is applied to mitigate overfitting. Training is performed for 200 epochs, with early stopping implemented on the validation set to ensure generalization performance. The gradients of the neural ODE module are computed using the adjoint sensitivity method to reduce memory consumption, and gradient clipping with a threshold of 1.0 is applied to prevent instability during backpropagation.

3.3. Training process and verification strategy

For the experiment, I use real-world time series data from enterprise employees. During the data partitioning phase, the time series is split into training, validation, and test sets at a 7:2:1 ratio, ensuring independent validation across time periods. The training set is used to update model parameters, the validation set is used to monitor changes in the loss function and adjust hyperparameters, and the test set is used to evaluate predictive performance. Before data partitioning, missing value interpolation, normalization, and timestamp alignment are performed to ensure consistency and comparability across batches of input sequences.

The training process uses the Adam optimizer to perform gradient descent on the joint loss of the VAE and neural ODE. The loss function is composed of the reconstruction error, the Kullback-Leibler divergence term, and the neural ODE trajectory-fitting error, with the weights set at 2:0.2:1 to balance reconstruction accuracy and dynamic stability. Each training iteration uses a batch size of 64, and the learning rate is gradually decayed as training progresses using a cosine annealing strategy, starting from an initial value of 1×10^{-3}, to improve the stability of latent convergence. The time evolution of the neural ODE is performed using an adaptive Runge–Kutta 4th/5th order solver, which computes

the latent state and returns the gradient at each time step to ensure continuous differentiability during end-to-end backpropagation.

The validation strategy is centered around an early stopping mechanism. Training is terminated when the MAE on the validation set fails to improve for 15 consecutive iterations to prevent overfitting and conserve computing resources. During training, the loss curve, accuracy, MAE, and R^2 metrics are recorded, and TensorBoard is used to monitor the model's convergence and changes in the latent space distribution.

In this study, employee satisfaction prediction is defined as a continuous regression task where the model outputs a real-valued satisfaction score. The reported accuracy metric is obtained by transforming predicted scores and ground-truth scores into binary labels using a fixed threshold of 0.5. Satisfaction scores above 0.5 are classified as 'high satisfaction' and those below 0.5 as 'low satisfaction'. Accuracy is then calculated as the proportion of predictions whose binary classification matches the ground-truth label. This approach ensures alignment between regression outputs and classification-based evaluation for a unified metric. Finally, the prediction accuracy, MAE, and RMSE are output on the test set and compared with the best validation results from the training phase to confirm the model's time series generalization performance and stability.

4. Experimental results and analysis

4.1. Prediction accuracy evaluation

In an employee satisfaction prediction experiment, five methods are designed and implemented for multi-source time-series data, enabling a systematic comparison. The Gated Recurrent Unit (GRU) employs a gating mechanism to reduce parameter count and mitigate vanishing gradients; the LSTM captures long-term dependencies through its memory cell structure; the Temporal Convolutional Network (CNN) uses one-dimensional dilated convolutions to handle parallel computation over long sequences; the VAE focuses on generative learning of latent distributions but lacks a description of continuous dynamics; and the proposed VAE-neural ODE method combines VAE with neural ODE to construct a continuous time evolution equation within the latent space. All models are trained and validated using the same sampling frequency and data preprocessing and evaluated using accuracy, mean absolute error (MAE), and root mean square error (RMSE). The accuracy metric is derived from a threshold-based discretization of continuous satisfaction scores, while MAE and RMSE directly measure regression error, forming a unified evaluation framework that captures numerical precision and categorical consistency.

The probability distribution of the sample-by-sample prediction error is extracted from the test set output of the best model, yielding the results shown in Figure 3.

Figure 3 shows clear differences in accuracy among the methods. VAE-neural ODE achieves a mean accuracy of 94.7%, surpassing GRU's 88.5%, LSTM's 89.7%, Temporal CNN's 90.3%, and VAE-only's 91.2%. This difference stems from VAE-neural ODE's generation of high-dimensional features in the latent space and its introduction of continuous-time dynamics modeling, which enables the model to more fully capture nonlinear temporal evolution. Regarding error metrics, VAE-neural ODE achieves an MAE of 0.031 and an RMSE of 0.052, which are lower than those of other methods. This is primarily due to the continuous dynamics module's ability to smooth latent state evolution and reduce the impact of short-term fluctuations on predictions. The prediction error distribution is

concentrated between -0.005 and 0.005, with a maximum deviation of 0.034. This indicates that the model maintains stability for most samples, with minor deviations primarily due to random fluctuations in employee satisfaction. Therefore, the experimental results demonstrate that VAE-neural ODE exhibits strong performance in comprehensive prediction accuracy and error control.

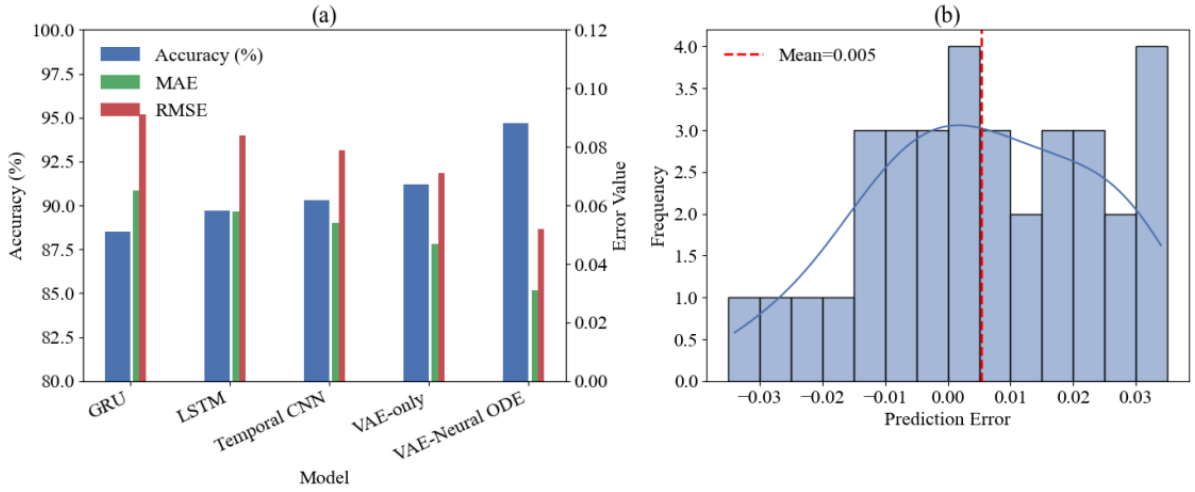

Figure 3. Comprehensive evaluation of prediction accuracy. (a) Comparison of prediction indicators across methods; and (b) VAE-neural ODE prediction error distribution.

To assess the statistical significance and robustness of these improvements, I perform five independent training runs for each model with different random seeds. The standard deviation of the VAE-neural ODE's accuracy across these runs is 0.6, indicating stable performance. I then conduct paired t-tests comparing the VAE-Neural ODE's performance against each baseline model across the five runs. The results are summarized in Table 4, which reports the mean difference, t-statistic, and p-value for accuracy, MAE, and RMSE metrics. All p-values are below 0.01, confirming that the observed improvements are statistically significant.

Table 4. Results of paired t-tests comparing the proposed VAE-neural ODE model against baseline models across evaluation metrics.

Comparison	Metric	Mean difference	t-statistic	p-value
VAE-Neural ODE vs. GRU	Accuracy	6.2	8.45	<0.001
VAE-Neural ODE vs. GRU	MAE	-0.018	-7.92	<0.001
VAE-Neural ODE vs. GRU	RMSE	-0.031	-8.11	<0.001
VAE-Neural ODE vs. LSTM	Accuracy	5	6.87	<0.001
VAE-Neural ODE vs. LSTM	MAE	-0.015	-6.54	<0.001
VAE-Neural ODE vs. LSTM	RMSE	-0.027	-7.03	<0.001
VAE-Neural ODE vs. Temporal CNN	Accuracy	4.4	5.92	0.002
VAE-Neural ODE vs. Temporal CNN	MAE	-0.012	-5.28	0.003
VAE-Neural ODE vs. Temporal CNN	RMSE	-0.021	-5.67	0.002
VAE-Neural ODE vs. VAE-only	Accuracy	3.5	4.76	0.005
VAE-Neural ODE vs. VAE-only	MAE	-0.014	-4.89	0.004
VAE-Neural ODE vs. VAE-only	RMSE	-0.015	-4.52	0.006

To further isolate the contribution of each component, a systematic ablation study is conducted by comparing the full model against four additional variants. First, the neural ODE module is replaced with a single-layer GRU operating on the latent space, forming a VAE+GRU model. Second, the VAE is removed, and the preprocessed features are directly fed into a neural ODE model for prediction, serving as a neural ODE without a VAE baseline. Third, the neural ODE module is removed and replaced with a direct latent regression layer, which corresponds to a VAE-only configuration. Fourth, the VAE encoder is replaced with a deterministic encoder without probabilistic sampling, forming a neural ODE with a deterministic encoder baseline. The full model is also evaluated with latent dimensions of 8, 16, and 32 to assess the impact of latent space capacity. The results are shown in Table 5. The VAE+GRU achieves an accuracy of 91.3%, lower than the full model, indicating the benefit of continuous-time dynamics over discrete recurrence in the latent space. The neural ODE without VAE achieves an accuracy of 90.7%, highlighting the importance of latent representation learning. The VAE-only configuration achieves an accuracy of 91.2%, while the deterministic encoder variant achieves 89.8%, confirming that the probabilistic latent representation provides additional expressiveness. Furthermore, the model achieves its highest accuracy at a latent dimension of 16, with performance dropping at 8 and 32, suggesting that overly constrained or overly flexible latent spaces hinder optimal learning of satisfaction dynamics.

Table 5. Results of the ablation study and latent dimension analysis.

Model Variant	Accuracy (%)	MAE	RMSE
VAE-Neural ODE (Full Model)	94.7	0.031	0.052
VAE+GRU	91.3	0.043	0.064
Neural ODE without VAE	90.7	0.049	0.071
VAE only	91.2	0.045	0.067
Neural ODE with deterministic encoder	89.8	0.052	0.074
VAE-LSTM	90.6	0.048	0.069
Full model (latent dim=8)	92.1	0.04	0.061
Full model (latent dim=32)	92.8	0.037	0.059

4.2. Continuous-time dynamic fitting results

For this experiment, I investigate the performance of fitting the continuous-time dynamics of employee satisfaction by constructing a three-part visualization analysis. Initially, a VAE-neural ODE model is used in the latent space to generate continuous evolutionary trajectories that reflect the temporal characteristics of the hidden state through dynamic changes in the latent variables. Thereafter, prediction residuals for different time periods are statistically analyzed to test the model's stability at different time scales. Furthermore, the neural ordinary differential equation time step is adjusted during the solution to assess its impact on fitting accuracy.

The results are shown in Figure 4. Figure 4(a) shows the boxplot distribution of the prediction residuals for six weekly time windows, reflecting the concentration of errors at different time spans. Figure 4(b) presents a continuous comparison curve of the true and model trajectories in the two-dimensional latent space, illustrating the dynamic evolution of the latent state. Figure 4(c) shows the relationship between the ODE solution step size and three accuracy metrics: MAE, RMSE, and R^2, quantifying the corresponding changes in numerical solution resolution and prediction accuracy.

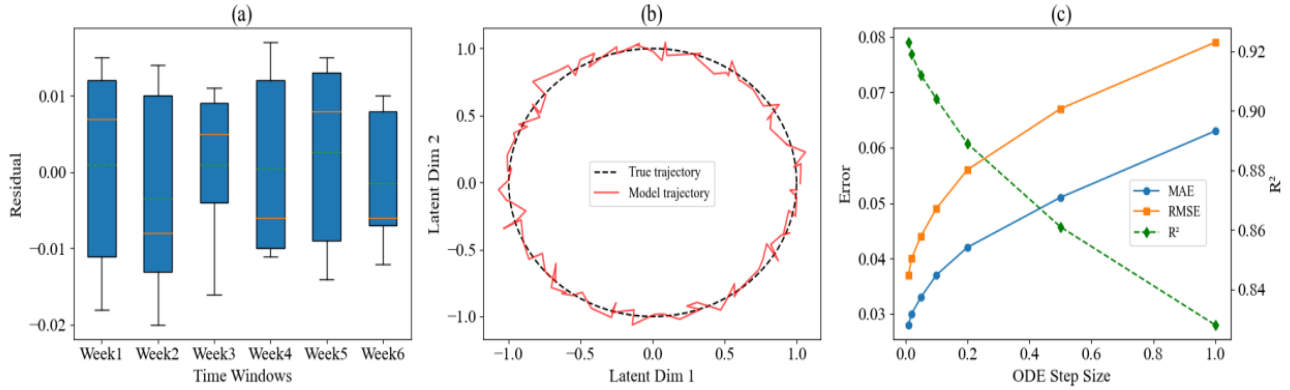

Figure 4. Combined diagram of the continuous-time dynamic fitting effect. (a) Distribution of prediction residuals in different time windows; (b) comparison of real and predicted trajectories in latent space; and (c) relationship between ODE step size and fitting accuracy.

Figure 4(a) shows that the residuals for the six time windows are highly concentrated, with means ranging from -0.003 to 0.003. This indicates that the forecast bias remains stable across weeks and does not accumulate over time. This is because the latent distribution fully captures the high-dimensional characteristics of employee states, thereby suppressing systematic drift. Figure 4(b) shows that the overall trends of the true and model trajectories are consistent, and the predicted curve exhibits a cyclical pattern similar to the true curve in the latent space, reflecting the dynamic equation's ability to track trends in the time domain. Figure 4(c) shows that as the step size increases from 0.01 to 1.00, the MAE and RMSE increase from 0.028 and 0.037 to 0.063 and 0.079, respectively, while the R^2 decreases from 0.923 to 0.828. This change is due to the larger step size, which reduces the accuracy of the numerical integration process in capturing subtle state changes, leading to a gradual accumulation of numerical integration error. Overall, this experiment reveals that the model can more stably reproduce the continuous evolution of the latent state when the step size is kept small.

From the numerical ODE viewpoint, increasing the step size reduces the solver's resolution and increases the local truncation error, which propagates over time as an integration bias. This effect becomes more pronounced when the learned vector field exhibits sharper curvature (i.e., locally faster dynamics), where coarse discretization may effectively under-sample the latent flow and produce an overly smoothed trajectory.

To comprehensively evaluate the model's stability in continuous-time modeling, I further examined the impact of solver type, regularization strength, and missing data patterns on prediction performance. Table 6 compares model predictions across configurations.

Table 6 shows that, in the comparison of solver types, the DOPRI5 method with adaptive step size achieves lower MAE and RMSE and higher accuracy compared to the Euler method with fixed step size, indicating that the adaptive solver can more accurately capture the continuous evolution of the latent state. In the regularization strength experiment, the model performs best when the original L2 coefficient is set to 0.001. Moreover, increasing the coefficient to 0.01 leads to an increase in prediction error, while decreasing the coefficient to 0.0001 also causes a slight decrease in performance, indicating that an appropriate regularization strength is crucial for balancing fitting accuracy and

generalization ability. Under the perturbation of missing data patterns, after randomly removing 20% of the input feature sequences and using linear interpolation for imputation, the model's MAE increases to 0.045, and the accuracy decreases to 91.2%, indicating that data integrity has a certain impact on model prediction, but the model maintains basic predictive ability.

Table 6. Comparison of model prediction performance under different solver types, regularization strengths, and missing data modes.

Configuration	Setting	MAE	RMSE	Accuracy (%)
Solver type	DOPRI5 (adaptive step size)	0.03	0.049	95.2
	Euler (fixed step size)	0.048	0.071	90.8
Regularization strength	L2 coefficient = 0.001 (original setting)	0.031	0.052	94.7
	L2 coefficient = 0.01	0.04	0.064	92.3
	L2 coefficient = 0.0001	0.035	0.058	93.6
Missing data pattern	20% feature sequence missing, linear interpolation imputation	0.045	0.069	91.2

4.3. Separability of latent-space representations

In the latent-space representation experiment, dimensionality reduction and distribution comparison are performed on multi-source employee features to verify the encoder's mapping effectiveness for different group characteristics. The experiment divides samples into three groups: high-, medium-, and low-satisfaction employees. Initially, the original feature data are dimensionally reduced using PCA to obtain a two-dimensional representation, revealing the interweaving of distributions across groups. Subsequently, the trained variational autoencoder is used to map the same group of samples into the latent space. A probability distribution generation mechanism is used to produce a continuous representation, which is then visually compared via dimensionality-reduction projection. This process yields the distribution results shown in Figure 5.

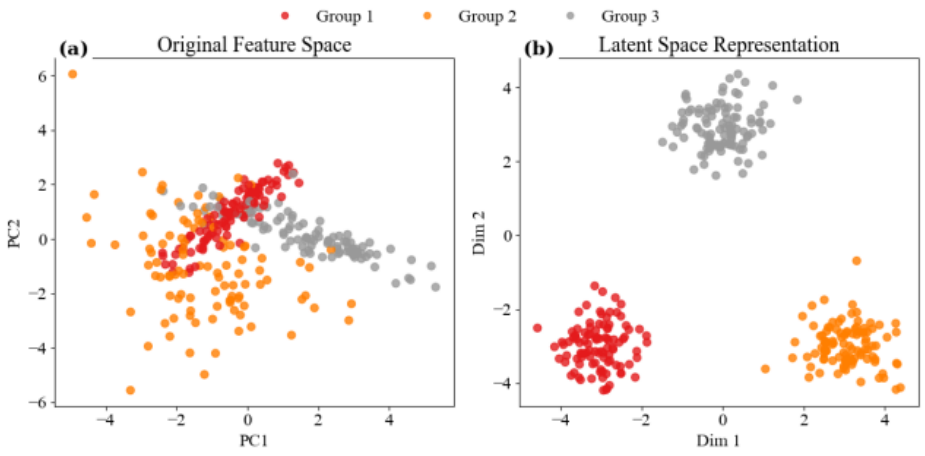

Figure 5. Latent space representation visualization. (a) Original feature space; and (b) Latent space representation.

Figure 5 shows significant overlap between groups in the original feature space projection results. This is because the original input dimensions contain many noisy and redundant features, and the dimensionality reduction process fails to effectively preserve discriminative structure during compression, resulting in the projected distributions of different groups becoming intertwined. In the latent-space mapping results, group clustering is significantly improved. This is primarily due to the encoder's reparameterized sampling and distribution constraints during training, which map high-dimensional inputs to the latent distribution. This reduces redundant dimensions and strengthens the representation of features related to group differences, resulting in a compact intra-class and distinct inter-class structure after projection. These experimental results demonstrate that the latent space can effectively reveal group differences.

To quantitatively assess the separability of the latent space, silhouette coefficients and ratios of inter-class to intra-class distances are calculated for the high-, medium-, and low-satisfaction groups. Table 7 compares the original feature space and the VAE latent space for the indicators above.

Table 7. Comparison of cluster separability indices between the original feature space and the VAE latent space.

Evaluation metric	Original feature space	VAE Latent space
Silhouette Score	0.12	0.54
Inter-Class Distance / Intra-Class Distance	0.21	0.48

Table 7 shows that the silhouette coefficients of the three groups in the original feature space are 0.12, indicating that the sample clustering structure is fuzzy and the boundaries between groups are unclear. In the latent space after VAE encoding, the silhouette coefficient increases to 0.54, indicating clearer cluster boundaries. The ratio of inter-class to intra-class distance is 0.21 in the original feature space, increasing to 0.48 in the latent space, indicating that the discriminability between satisfaction groups is significantly enhanced. These quantitative indicators confirm the conclusion visualized in Figure 5, namely, VAE encoding can effectively improve the feature separability of employees with different satisfaction levels in the latent space.

4.4. Temporal generalization capabilities

In the time series generalization experiment, multi-window forecasting tasks are designed to verify the model's stability and robustness across forecast horizons, with spans ranging from one week to two quarters. Six models, including LSTM, GRU, Transformer, VAE-LSTM, Latent ODE, and VAE-Neural ODE, are used as baselines for comparison. During the experiment, each model performs forecasts from the short term to the long term using the same training and validation set partitioning. Accuracy, mean absolute error (MAE), and root mean square error (RMSE) are calculated to assess performance degradation with increasing time horizons. This experimental approach directly demonstrates the performance differences between the models across forecast windows, as shown in Figure 6.

Figure 6 shows that as the forecast window gradually expands from one week to two quarters, the accuracy of all models decreases, while the error metrics increase. The LSTM achieves only 72.8% accuracy in the second-quarter forecast, with a MAE of 0.096 and an RMSE of 0.124. This is primarily due to its inability to capture long-term dependencies, leading to the accumulation of long-term

forecast bias. The GRU performs slightly better than the LSTM, but experiences error accumulation when forecasting is at the quarterly level, with its MAE increasing to 0.093. The Transformer maintains a relatively high fit for cross-month and cross-quarter forecasts, achieving 79.2% accuracy in the second quarter. This is primarily due to its self-attention mechanism, which preserves global dependencies over longer time series. The VAE-LSTM outperforms the three aforementioned models across window types, achieving 80.4% accuracy and a MAE of 0.075 in the second quarter, demonstrating that latent feature generation effectively improves representational capabilities. Latent ODE, with its continuous-time state-space modeling capabilities, demonstrates superior stability compared to VAE-LSTM in cross-quarter forecasting, achieving a forecast accuracy of 83.5% and a MAE of 0.065 for the second quarter. VAE-Neural ODE demonstrates greater stability and accuracy across all time windows, maintaining a Q2 forecast accuracy of 87.2% and an MAE of 0.050. This is primarily due to its introduction of continuous-time dynamic equations in the latent space, which makes state evolution smoother and more consistent with the distribution of actual time-series features. In summary, this method demonstrates greater stability and robustness in generalized forecasting across time windows.

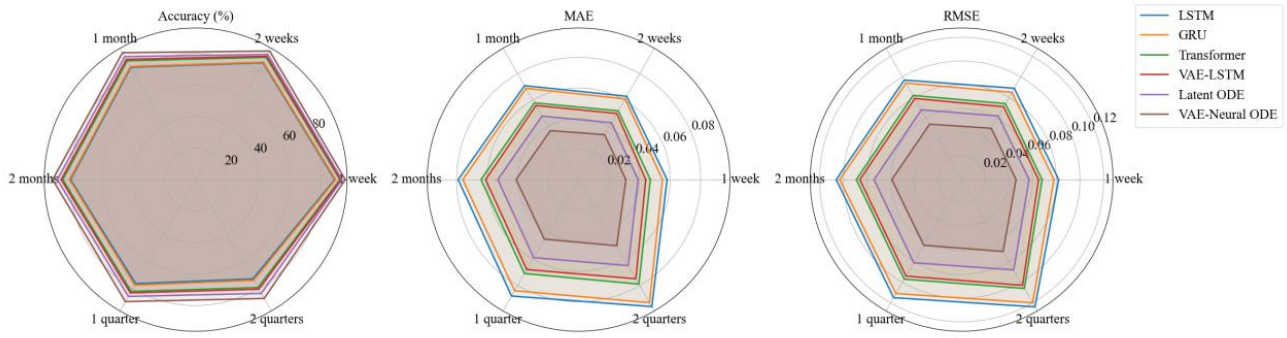

Figure 6. Comparison of prediction results in different time windows.

4.5. Model complexity and training convergence analysis

In an experiment to model dynamic, complex employee states, this study systematically compares differences in model size and training processes across methods. Five models are constructed: GRU, LSTM, Temporal CNN, VAE, and VAE-Neural ODE. Using the same dataset and training conditions, the number of parameters, single-epoch training time, and training and validation loss curves are recorded to evaluate each method's performance in terms of computational complexity and convergence efficiency. This experimental design makes it possible to directly assess the impact of model structural characteristics on computational cost and convergence trends. The results are shown in Figure 7.

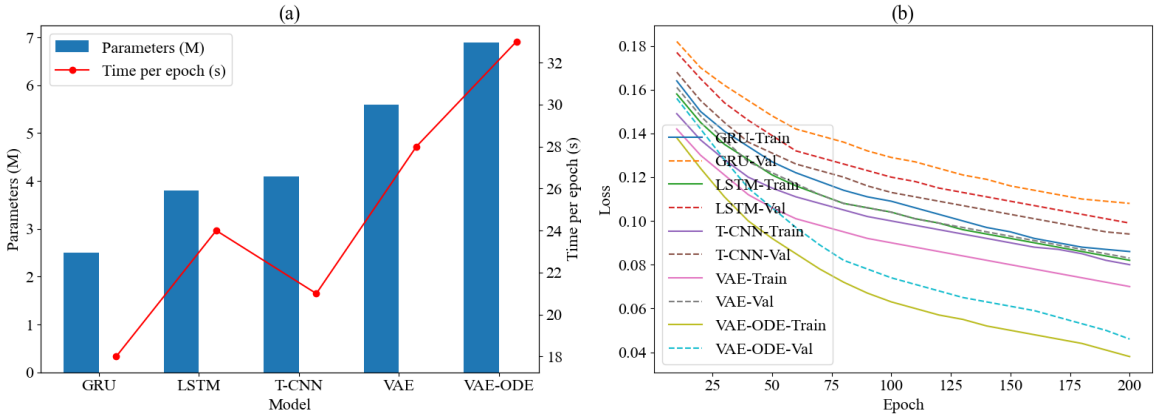

Figure 7. Comparison of model complexity and convergence. (a) Comparison of parameter quantity and training time; and (b) training convergence curve.

Figure 7 shows that the GRU has 2.5 million parameters and a single-epoch training time of 18 seconds. The LSTM, due to its increased memory cell structure, has 3.8 million parameters and a training time of 24 seconds. The Temporal CNN has 4.1 million parameters and a training time of 21 seconds, indicating that the differences in computational paths between the gating structure and convolutional mechanism directly affect model overhead. The VAE, in latent distribution modeling, has 5.6 million parameters and a training time of 28 seconds, demonstrating the additional computational burden imposed by the generative network structure. After VAE-neural ODE further introduces dynamic equations into the latent space, the parameter count rises to 6.9 million, and the training time reaches 33 seconds. This increase is primarily due to the computational overhead introduced by the differential equation solver's continuous-time evolution. During the convergence process, the validation loss of the VAE-Neural ODE drops to 0.078 by epoch 50 and stabilizes at 0.046 after epoch 200. The GRU and LSTM maintain validation losses of 0.108 and 0.099, respectively, in the same epochs, demonstrating that the ODE module effectively enhances the efficiency of continuous-time fitting. Overall, despite the VAE-neural ODE's higher complexity, its faster loss reduction and lower convergence level enable it to demonstrate superior training stability in dynamic prediction tasks.

5. Conclusions

The joint VAE-neural ODE model proposed in this study integrates generative representation learning with continuous-time latent dynamics to predict employee satisfaction. By using the VAE to learn compact latent distributions from high-dimensional, multi-source HR data and the Neural ODE to characterize smooth state evolution in continuous time, the framework provides an end-to-end approach for modeling the nonlinear temporal variation of employee satisfaction. Experimental results show that the proposed method outperforms competing models in predictive accuracy and error control, achieving 94.7% accuracy, an MAE of 0.031, and an RMSE of 0.052 for the test set, while maintaining 87.2% accuracy and an MAE of 0.050 in cross-quarter forecasting. In addition, the convergence behavior indicates that the model can be effectively optimized under the adopted training setting. These findings suggest that the proposed framework can serve as a feasible technical approach

for predicting employee satisfaction in complex, dynamic environments and provide quantitative evidence for data-informed human resource analysis. The results also indicate that the continuous-time latent formulation offers potential value for downstream tasks, such as scenario-based analysis and policy assessment. Future studies may extend this framework by introducing controllable HR interventions, such as training intensity, workload adjustment, or incentive changes, as time-dependent inputs to the latent dynamics, thereby enabling controlled neural ODE modeling for what-if analysis and dynamic decision support under practical operational constraints.

Use of Generative-AI tools declaration

The author declares he has not used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of interest

The author declares no conflicts of interest in relation to this paper.

References

1. C. H. Chang, H. W. Lin, W. H. Tsai, W. L. Wang, C. T. Huang, Employee satisfaction, corporate social responsibility and financial performance, *Sustainability*, **13** (2021), 9996. https://doi.org/10.3390/su13189996
2. B. Miethlich, M. Beliakova, L. Voropaeva, O. Ustyuzhina, T. Yurieva, Internal corporate policy: CSR and employee satisfaction, *Employ. Respons. Rights J.*, **35** (2023), 127–141. https://doi.org/10.1007/s10672-022-09406-5
3. P. Ruiz-Palomino, R. Morales-Sánchez, R. Martínez-Cañas, Corporate sustainability, ethics and employee satisfaction, *Sustainability*, **13** (2021), 11964. https://doi.org/10.3390/su132111964
4. E. Uchida, Y. Kino, Study on the relationship between employee satisfaction and corporate performance in Japan via text mining, *Procedia Comput. Sci.*, **192** (2021), 1730–1739. https://doi.org/10.1016/j.procs.2021.08.178
5. T. V. D. Voordt, P. A. Jensen, The impact of healthy workplaces on employee satisfaction, productivity and costs, *J. Corp. Real Estate*, **25** (2023), 29–49. https://doi.org/10.1108/jcre-03-2021-0012
6. T. Zhang, J. Zhang, S. Tu, An empirical study on corporate ESG behavior and employee satisfaction: A moderating mediation model, *Behav. Sci.*, **14** (2024), 274. https://doi.org/10.3390/bs14040274
7. I. N. Latifah, A. A. Suhendra, I. Mufidah, Factors affecting job satisfaction and employee performance: A case study in an Indonesian sharia property companies, *Int. J. Prod. Perform. Manage.*, **73** (2024), 719–748. https://doi.org/10.1108/ijppm-03-2021-0132
8. S. Tabassum, S. Alam, M. S. Akter, S. Ahmad, Factors influencing employee satisfaction in workplace, *Int. J. Eng. Appl. Sci. Technol.*, **7** (2023), 61–67. https://doi.org/10.33564/ijeast.2023.v07i12.010
9. J. Kim, P. S. Chang, S. B. Yang, I. Choi, B. Lee, A comparative analysis of job satisfaction prediction models using machine learning: A mixed-method approach, *Data Technol. Appl.*, **59** (2025), 41–60. https://doi.org/10.1108/dta-10-2023-0697

10. M. E. Zhuang, W. T. Pan, Data modelling in human resource management: influencing factors of employees' job satisfaction, *Math. Probl. Eng.*, **2022** (2022), 3588822. https://doi.org/10.1155/2022/3588822
11. P. Kang, Predicting HR management strategies and employee satisfaction in enterprises based on deep generative models, *Appl. Math. Nonlinear Sci.*, **10** (2025), 1–16. https://doi.org/10.2478/amns-2025-0043
12. F. Rustam, I. Ashraf, R. Shafique, A. Mehmood, S. Ullah, G. S. Choi, Review prognosis system to predict employees job satisfaction using deep neural network, *Comput. Intell.*, **37** (2021), 924–950. https://doi.org/10.1111/coin.12440
13. J. S. Sekhon, H. Sadawarti, Prediction of employee satisfaction level using optimized neural network, *Int. J. Adv. Res. Eng. Technol.*, **12** (2021), 129–143. https://doi.org/10.34218/IJARET.12.1.2021.011
14. M. Y. M. Adib, S. Chakraborty, M. T. Waishy, M. H. K. Mehedi, A. A. Rasel, BiLSTM-ANN based employee job satisfaction analysis from glassdoor data using web scraping, *Procedia Comput. Sci.*, **222** (2023), 1–10. https://doi.org/10.1016/j.procs.2023.08.139
15. M. Qadir, I. Noreen, A. A. Shah, Bi-LSTM deep learning approach for employee churn prediction, *J. Inf. Commun. Technol. Robot. Appl.*, **12** (2021), 1–10. https://doi.org/10.51239/jictra.v0i0.255
16. T. Kanchinadam, Z. Meng, J. Bockhorst, V. Singh, G. Fung, Graph neural networks to predict customer satisfaction following interactions with a corporate call center, *arXiv:2102.00420*, 2021. https://doi.org/10.48550/arXiv.2102.00420
17. Z. Yue, M. R. A. Razak, L. Li, Dynamic context-weighted embeddings: A novel approach to predictive modelling, *IET Commun.*, **19** (2025), e70036. https://doi.org/10.1049/cmu2.70036
18. X. Li, L. Zhang, D. Li, D. Guo, Construction and simulation of a strategic HR decision model based on recurrent neural network, *J. Math.*, **2022** (2022), 5390176. https://doi.org/10.1155/2022/5390176
19. H. K. Sriram, Generative AI and neural networks in human resource management: Transforming payroll, workforce insights, and digital employee payments through AI innovations, *Adv. Consum. Res.*, **2** (2025), 250–260.
20. H. Du, S. Du, W. Li, Probabilistic time series forecasting with deep non-linear state space models, *CAAI Trans. Intell. Technol.*, **8** (2023), 3–13. https://doi.org/10.1049/cit2.12085
21. H. Han, Z. Liu, M. B. Barrios, J. Li, Z. Zeng, N. Sarhan, et al., Time series forecasting model for non-stationary series pattern extraction using deep learning and GARCH modeling, *J. Cloud Comp.*, **13** (2024), 2. https://doi.org/10.1186/s13677-023-00576-7
22. L. Zhang, R. Wang, Z. Li, J. Li, Y. Ge, S. Wa, Time-series neural network: a high-accuracy time-series forecasting method based on Kernel filter and time attention, *Information*, **14** (2023), 500. https://doi.org/10.3390/info14090500
23. M. A. Khan, S. A. Iqbal, M. S. Khan, M. G. Hafez, Factor-bridging algorithm for the prediction of job satisfaction: Developing country perspective, *J. King Saud Univ. Comput. Inf. Sci.*, **35** (2023), 101743. https://doi.org/10.1016/j.jksuci.2023.101743
24. R. Wu, X. Zhao, Z. Li, Y. Xie, The role of employee personality in employee satisfaction and turnover: Insights from online employee reviews, *Pers. Rev.*, **53** (2024), 1581–1611. https://doi.org/10.1108/pr-04-2023-0309

25. S. Fraihat, Q. Shambour, M. A. Al-Betar, S. N. Makhadmeh, Variational autoencoders-based algorithm for multi-criteria recommendation systems, *Algorithms*, **17** (2024), 561. https://doi.org/10.3390/a17120561
26. N. Milano, M. Casella, R. Esposito, D. Marocco, Exploring the potential of variational autoencoders for modeling nonlinear relationships in psychological data, *Behav. Sci.*, **14** (2024), 527. https://doi.org/10.3390/bs14070527
27. H. Yang, R. Rang, L. Xing, L. Zhang, H. Cai, M. Guo, et al., V-GMR: A variational autoencoder-based heterogeneous graph multi-behavior recommendation model, *Appl. Intell.*, **54** (2024), 3337–3350. https://doi.org/10.1007/s10489-024-05360-x
28. M. Celia, S. Monaco, D. Apiletti, A comparative study of neural ordinary differential equations and neural operators for modeling temporal dynamics, *Neural Comput. Applic.*, **37** (2025), 25319–25338. https://doi.org/10.1007/s00521-025-11580-0
29. Z. Chang, S. Liu, R. Qiu, S. Song, Z. Cai, G. Tu, Time-aware neural ordinary differential equations for incomplete time series modeling, *J. Supercomput.*, **79** (2023), 18699–18727. https://doi.org/10.1007/s11227-023-05327-8
30. F. Zhan, X. Zhou, S. Li, D. Jia, H. Song, Learning latent ODEs with graph RNN for multi-channel time series forecasting, *IEEE Signal Process. Lett.*, **30** (2023), 1432–1436. https://doi.org/10.1109/LSP.2023.3320439

AIMS Press