Research in Astronomy and Astrophysics, 25:104007 (9pp), 2025 October
https://doi.org/10.1088/1674-4527/adf711
CSTR: 32081.14.RAA.adf711
© 2025. National Astronomical Observatories, CAS and IOP Publishing Ltd. All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Enhancing Stellar Spectra with Diffusion Probabilistic Models: A Novel Approach to Denoising Low SNR Astronomical Data

Jingzhen Sun[1], Yude Bu[1,4], Jiangchuan Zhang[1], Mengmeng Zhang[1], Shanshan Li[1], Ke Wang[1], Yuhang Zhang[1], Zhenping Yi[2], Xiaoming Kong[2], Meng Liu[2], and Minglei Wu[3]
[1] School of Mathematics and Statistics, Shandong University, Weihai 264209, China; buyude@sdu.edu.cn
[2] School of Mechanical, Electrical and Information Engineering, Shandong University, Weihai 264209, China
[3] School of Computer Science and Technology, Shandong Technology and Business University, Yantai 264005, China
Received 2025 April 30; revised 2025 June 9; accepted 2025 June 15; published 2025 September 12

Abstract

Astronomical spectra are vital for deriving stellar properties, yet low signal-to-noise ratio (SNR) spectra often obscure key features, complicating accurate analysis. This study presents spec-Diffusion Probabilistic Models (DDPM), a novel deep learning approach based on DDPM, aimed at denoising low SNR spectra to improve stellar parameter estimation. Leveraging the LAMOST DR10 data set, we developed spec-DDPM using a tailored U-Net architecture (spec-Unet) to iteratively predict and remove noise. The model was trained on 28,500 low and high SNR spectral pairs and benchmarked against conventional methods, including Principal Component Analysis, wavelet techniques, and a modified DnCNN model. The spec-DDPM demonstrated superior performance, with reduced Mean Absolute Error, elevated Structural Similarity Index Measure, and enhanced spectral loss metrics. It effectively preserved critical spectral features and corrected continuum distortions. Validation experiments further confirmed its ability to improve stellar parameter estimation with reduced errors. These results underscore spec-DDPM's potential to elevate spectral data quality, offering applications in restoring defective spectra and refining large-scale astronomical surveys. This work highlights the transformative role of deep learning in astronomical data processing.

Key words: techniques: spectroscopic – methods: statistical – methods: data analysis

1. Introduction

Astronomical spectra serve as indispensable tools for unraveling the physical properties of celestial bodies, particularly stars. Through the analysis of absorption and emission lines within these spectra, astronomers can derive critical parameters such as effective temperature, surface gravity, and elemental abundances. However, observational challenges often degrade the quality of spectral data, especially for faint or distant stars, resulting in low signal-to-noise ratio (SNR) spectra. Besides, for large scale surveys such as LAMOST, we found that nearly a half of low resolution spectra ($R \approx 1800$) has SNR_g less than 20 in DR10 (as discussed in Section 3). Such noise can mask essential spectral features, leading to uncertainties in the determination of stellar parameters. Conventional denoising techniques, including Principal Component Analysis (PCA) and wavelet-based methods, have been widely applied to address this issue. Yet, these approaches frequently struggle to preserve the detailed structure of spectral lines or to adequately handle the complex noise patterns inherent in astronomical observations.

Diffusion Probabilistic Models (DDPMs) have emerged as a powerful class of generative models, achieving remarkable success in image synthesis, super-resolution, and beyond (Ho et al. 2020; Dhariwal & Nichol 2021; Yue et al. 2023). These models operate by iteratively adding noise to data and then learning to reverse this process, effectively capturing complex data distributions. Recent theoretical advancements have deepened our understanding of DDPMs, particularly through their connection to score-based generative models via stochastic differential equations (SDEs), providing a unified framework for generative modeling (Song et al. 2021). Their versatility is evident in diverse applications, including bioinformatics for protein design and drug discovery (Guo et al. 2023), and astronomy for generating realistic galaxy images (Smith et al. 2022) and conditionally producing astrophysical images (Zhao et al. 2023).

Motivated by these successes, our work adapts DDPMs to the challenge of astronomical spectral denoising. We propose spec-DDPM, a method that leverages the generative capabilities of DDPMs to enhance the quality of low SNR spectra, thereby improving the accuracy of stellar parameter estimation and elemental composition analysis. By tailoring the iterative noise addition and removal process to one-dimensional spectral data, spec-DDPM reduces noise while preserving critical spectral features, building on principles demonstrated in image processing tasks (Yue et al. 2023). This approach addresses the unique demands of astronomical spectra, offering a robust tool for precision astrophysics.

[4] Corresponding author.

The key contributions of this research are threefold:

1. *Development of spec-DDPM*. We present a tailored adaptation of DDPMs for spectral denoising, employing a Markov chain process to progressively refine the residual between low and high SNR spectra.
2. *U-Net-Based Noise Prediction*. We design a U-Net architecture, dubbed spec-Unet, to predict noise at each time step, integrating ResNet blocks and linear attention mechanisms to capture both global and local spectral characteristics (Ronneberger et al. 2015).
3. *Rigorous Comparative Analysis*. We evaluate spec-DDPM against established denoising methods-PCA, wavelet-based techniques (Haar and db1), and a modified DnCNN model-demonstrating its superior performance through metrics such as Mean Absolute Error (MAE) and a one-dimensional Structural Similarity Index Measure (SSIM-1D) (Wang et al. 2004; Zhang et al. 2017).

This paper is structured as follows: Section 2 elaborates on the spec-DDPM model, detailing its forward and reverse processes and the noise schedule. Section 3 describes the LAMOST DR10 data set utilized in this study and the preprocessing methodology (Wang et al. 2009). Section 4 presents the experimental results, comparing spec-DDPM's performance across various stellar types and effective temperatures against baseline methods, detailed analysis on feature lines restoration, and hyperparameter tuning. In Section 5, we design a parameter estimation experiment to verify the effect of spec-DDPM. We conclude with a discussion of our findings and future research directions in Section 6.

2. Methodology

2.1. spec-DDPM Model

The full structure of spec-DDPM is shown in Figure 1. We design a U-Net model, called spec-Unet, to predict noise at each time step. This model includes multiple layers of ResNet blocks and linear attention mechanisms. These components help the model capture both the overall spectrum and its small details effectively.

The design of the forward process, reverse process, and noise schedule in our spec-DDPM model is inspired by the ResShift framework proposed by Yue et al. (2023) for image super-resolution tasks. While ResShift focuses on adjusting residuals between low- and high-resolution images, we adapt its core principles to the domain of astronomical spectral denoising. Specifically, we tailor the Markov chain formulation and noise scheduling strategy to address the unique challenges of low SNR spectra, ensuring robust preservation of spectral features and accurate continuum correction. By building upon the foundational concepts of ResShift, we develop a novel approach optimized for one-dimensional spectral data, as detailed in the following subsections.

2.2. Forward Process

We define high-quality (HQ) spectra as x_0 and low-quality (LQ) spectra as y_0, with their residual given by $e_0 = y_0 - x_0$. The central concept of our method is to progressively adjust this residual e_0 over a Markov chain of length T, transitioning from y_0 to x_0. To this end, we introduce a monotonically increasing shift sequence $\{\eta_t\}_{t=1}^{T}$, where $\eta_1 \to 0$ and $\eta_T \to 1$. The transition distribution is defined as:

$$q(x_t|x_{t-1}, y_0) = \mathcal{N}(x_t; x_{t-1} + \alpha_t e_0, \kappa^2\alpha_t I), \quad t = 1, 2,\ldots,T, \tag{1}$$

with $\alpha_t = \eta_t - \eta_{t-1}$ for $t > 1$, $\alpha_1 = \eta_1$, and κ as a hyperparameter regulating noise variance, I is identity matrix. The marginal distribution at any step t is analytically derivable:

$$q(x_t|x_0, y_0) = \mathcal{N}(x_t; x_0 + \eta_t e_0, \kappa^2\eta_t I), \quad t = 1, 2,\ldots,T. \tag{2}$$

This transition design rests on two principles. First, the standard deviation $\kappa\sqrt{\alpha_t}$ ensures a smooth shift between x_{t-1} and x_t, with the expected distance bounded by $\sqrt{\alpha_t}$ within the spectral range [0, 1]:

$$\max[(x_0 + \eta_t e_0) - (x_0 + \eta_{t-1}e_0)] = \max[\alpha_t e_0] < \alpha_t < \sqrt{\alpha_t}. \tag{3}$$

The hyperparameter κ enhances adaptability. Second, the mean $x_0 + \alpha_t e_0$ governs the marginal distribution in Equation (2), with x_1 and x_T approximating $\delta_{x_0}(\cdot)$ and $\mathcal{N}(\cdot;y_0, \kappa^2 I)$, respectively, as proxies for HQ and LQ spectra. This structured Markov chain enables spectral denoising via reverse sampling conditioned on y_0.

2.3. Reverse Process

The reverse process seeks to approximate the posterior $p(x_0|y_0)$ through:

$$p(x_0|y_0) = \int p(x_T|y_0) \prod_{t=1}^{T} p_\theta(x_{t-1}|x_t, y_0) dx_{1:T}, \tag{4}$$

where $p(x_T|y_0) \approx \mathcal{N}(x_T|y_0, \kappa^2 I)$, and $p_\theta(x_{t-1}|x_t, y_0)$ is a learnable reverse transition kernel. Consistent with diffusion model conventions (Sohl-Dickstein et al. 2015; Ho et al. 2020; Song et al. 2021), we model it as:

$$p_\theta(x_{t-1}|x_t, y_0) = \mathcal{N}(x_{t-1}; \boldsymbol{\mu}_\theta(x_t, y_0, t), \boldsymbol{\Sigma}_\theta(x_t, y_0, t)). \tag{5}$$

Optimization of θ minimizes the negative evidence lower bound:

$$\min_\theta \sum_t D_{\mathrm{KL}}[q(x_{t-1}|x_t, x_0, y_0) \| p_\theta(x_{t-1}|x_t, y_0)]. \tag{6}$$

where D_{KL} refers to Kullback–Leibler divergence, which measures the difference between distributions. Using Equations (1) and (2),

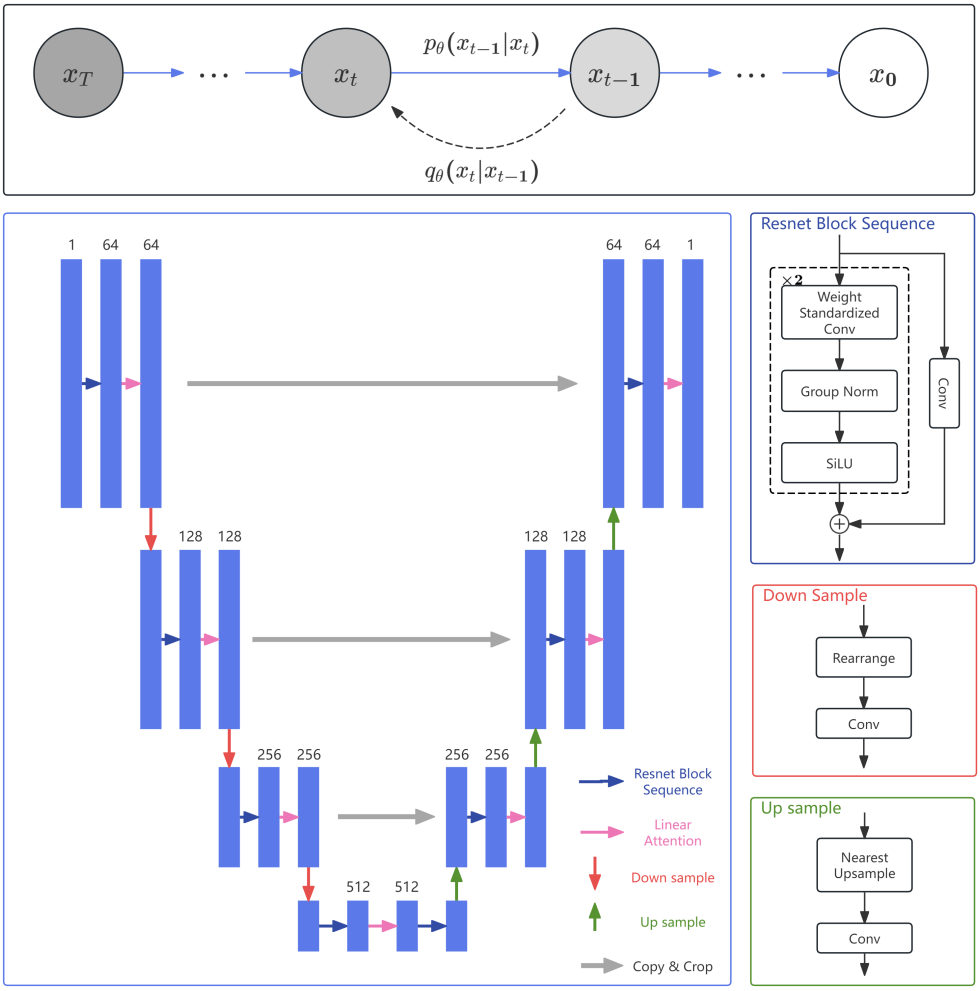


Figure 1. The structure of the spec-DDPM model. The top panel shows the overall spec-DDPM design. The bottom-left panel shows the spec-Unet network used at each time step. The bottom-right panel gives a detailed view of the key components inside the spec-Unet.

the target distribution becomes:

$$q(x_{t-1}|x_t, x_0, y_0) = \mathcal{N}\Bigg(x_{t-1}|\frac{\eta_{t-1}}{\eta_t}x_t + \frac{\alpha_t}{\eta_t}x_0, \kappa^2\frac{\eta_{t-1}}{\eta_t}\alpha_t I\Bigg). \quad (7)$$

The detailed calculation of Equation (7) can be found in Yue et al. (2023). Given the variance's independence from x_t and y_0, we set $\boldsymbol{\Sigma}_\theta(x_t, y_0, t) = \kappa^2\frac{\eta_{t-1}}{\eta_t}\alpha_t I$. $\boldsymbol{\mu}_\theta(x_t, y_0, t)$ is parameterized as:

$$\boldsymbol{\mu}_\theta(x_t, y_0, t) = \frac{\eta_{t-1}}{\eta_t}x_t + \frac{\alpha_t}{\eta_t}f_\theta(x_t, y_0, t), \quad (8)$$

where f_θ is a neural network with parameter θ that predicts x_0. Among different options, Equation (8) works best (Yue et al. 2023). Using it, Equation (6) can be simplified to:

$$\min_\theta \sum_t w_t|f_\theta(x_t, y_0, t) - x_0|_2^2, \quad (9)$$

with $w_t = \frac{\alpha_t}{2\kappa^2\eta_t\eta_{t-1}}$. Omitting w_t empirically improved outcomes (Yue et al. 2023).

2.4. Noise Schedule

Our approach utilizes κ and the shift sequence $\{\eta_t\}_{t=1}^T$ to define the noise schedule. The parameter κ adjusts overall noise intensity, while we focus on crafting $\{\eta_t\}_{t=1}^T$. From Equation (2), the noise level at x_t scales with $\kappa\sqrt{\eta_t}$, prompting us to design $\sqrt{\eta_t}$. As suggested by Song & Ermon (2019), $\kappa\sqrt{\eta_1}$ should be small (e.g., 0.04, as in LDM Rombach et al. 2022) to keep $q(x_1|x_0, y_0)$ close to $q(x_0)$. Since $\eta_1 \to 0$, we set η_1 as the smaller of $\left(\frac{0.04}{\kappa}\right)^2$ or 0.001. For the last step,

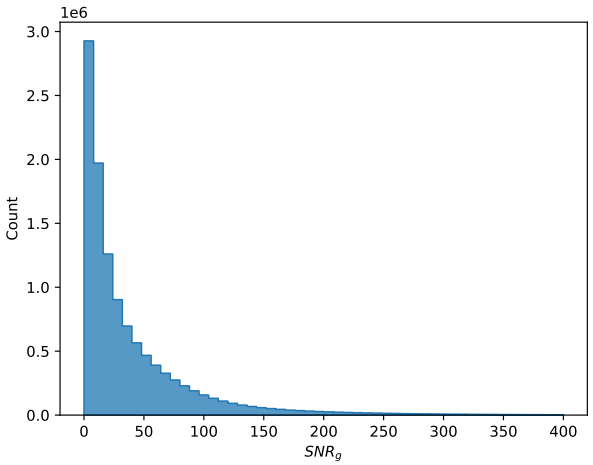


Figure 2. Histogram of the g-band signal-to-noise ratio (SNR_g) in the LAMOST DR10 dataset. The distribution is truncated at $SNR_g = 400$ for clarity.

$\eta_T = 0.999$ to approach 1. For steps t from 2 to $T-1$, we use a geometric pattern for $\sqrt{\eta_t}$:

$$\sqrt{\eta_t} = \sqrt{\eta_1} \times b_0^{\beta_t}, \quad t = 2,\ldots,T-1, \tag{10}$$

where

$$\beta_t = \left(\frac{t-1}{T-1}\right)^p \times (T-1),$$
$$b_0 = \exp\left[\frac{1}{2(T-1)} \log \frac{\eta_T}{\eta_1}\right]. \tag{11}$$

The hyperparameter p governs the growth rate. This schedule offers flexibility: smaller κ values shorten the chain, p balances fidelity and realism.

3. Dataset and Preprocessing

We utilized low-resolution spectra with a resolving power of approximately 1800 from the LAMOST Data Release 10 (DR10) (Wang et al. 2009). Spectra with low SNR were selected based on the criterion $10 < SNR_g < 15$, while high SNR spectra were chosen with $SNR_g > 50$. This selection process resulted in 28,500 pairs of low and high SNR spectra. The distribution of SNR_g in the LAMOST DR10 data set is illustrated in Figure 2, showing that a significant portion of the spectra have low SNR; specifically, 48.90% of the spectra have $SNR_g < 20$, while only 26.04% have $SNR_g > 50$.

During preprocessing, all spectra were RV corrected and interpolated onto a uniform wavelength grid ranging from 4000 to 8096 Å, producing flux vectors of 4096 dimensions. Each spectrum was then normalized by dividing it by its pseudo-continuum. Examples of paired high and low SNR spectra are presented in Figure 3. The data set was divided into training, validation, and test sets in an 8:1:1 ratio.

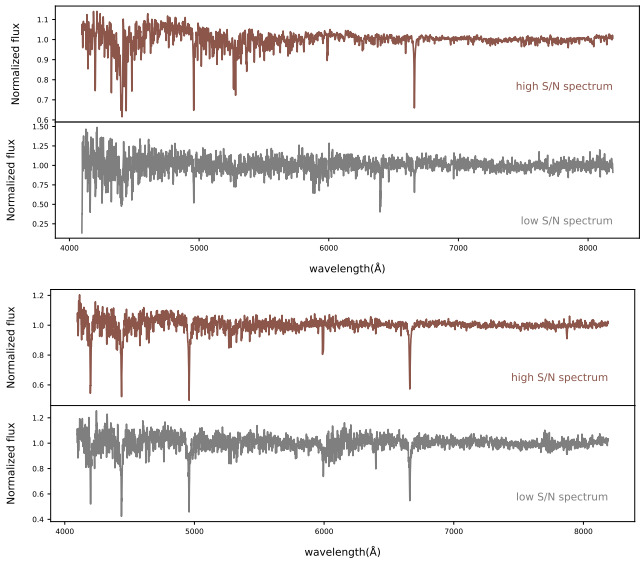


Figure 3. Examples of paired high and low signal-to-noise ratio (SNR) spectra. The top sub-panel of each panel shows high SNR spectra, while the bottom sub-panel displays the corresponding low SNR spectra from the same stars.

4. Results

4.0.1. Performance Comparison

To assess the performance of our spec-DDPM model, we compared it with several well-established denoising methods: PCA, wavelet-based denoising using Haar and db1 wavelets, and a modified version of the DnCNN model adapted for one-dimensional spectral data.

PCA proved to be suboptimal for denoising high-dimensional, low-SNR spectra. With only ten principal components, it accounted for only 40% of the cumulative variance, and even with 1000 components, this increased to only 80%. This suggests that PCA struggles to effectively capture the complex structure of spectral data.

Wavelet-based denoising was conducted using the Haar and db1 (Daubechies-1) wavelets. The Haar wavelet, valued for its simplicity and effectiveness in astronomical spectral denoising, was applied with a two-level decomposition to balance noise reduction with the preservation of spectral features. Similarly, the db1 wavelet, known for its excellent time-frequency localization and multi-resolution capabilities, decomposed the spectra into frequency sub-bands, enabling targeted noise suppression in high-frequency components.

The DnCNN model (Zhang et al. 2017), originally developed for image denoising, was adapted for 1D spectral data by modifying its architecture to process one-dimensional inputs. Despite its potential, our experiments showed that DnCNN often produced overly smoothed spectra, likely because it was trained with white Gaussian noise (WGN), which does not fully represent the complex noise patterns in astronomical spectra.

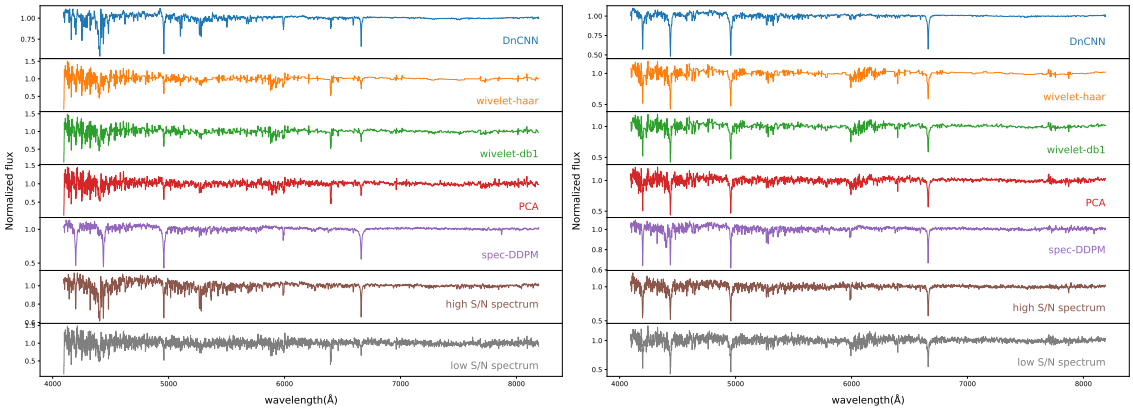


Figure 4. Denoising results of various models on two examples low signal-to-noise ratio (SNR) spectra. From top to bottom in each panel: denoised spectra from DnCNN, wavelet-Haar, wavelet-db1, PCA, spec-DDPM, followed by the original high SNR spectrum and the low SNR spectrum.

Our spec-DDPM model, employing a diffusion process with $T = 15$ time steps, demonstrated superior denoising performance, as detailed in the metrics below. Figure 4 showcases the denoising results of spec-DDPM alongside the other models on two examples low-SNR spectra. Notably, spec-DDPM not only preserves spectral lines more consistently but also corrects inaccuracies in the continuum—a capability not observed in the comparative methods. This points to potential applications beyond denoising, such as restoring defective spectra or inpainting incomplete data, which we intend to investigate in future research.

Model performance was evaluated using two metrics: MAE, a 1D adaptation of the SSIM-1D, and a spectral loss based on the Fourier transform. MAE measures the average absolute difference between the denoised and high-SNR spectra, while SSIM-1D assesses structural similarity by considering luminance, contrast, and structural correlation. The SSIM-1D index is defined as:

$$\text{SSIM-1D}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}, \quad (12)$$

where μ_x and μ_y are the sample means, σ_x and σ_y are the sample standard deviations, σ_{xy} is the sample covariance, and C_1, C_2 are stabilization constants.

While spectral loss quantifies the difference between the amplitude spectra of the denoised and high-SNR spectra, emphasizing the preservation of both high-frequency components (e.g., sharp atomic lines) and low-frequency components (e.g., broad molecular bands). The spectral loss is defined as:

$$L_{\text{spectral}}(x, y) = \frac{1}{N}\sum_{i=1}^{N} ||\text{FFT}(x_i)| - |\text{FFT}(y_i)||, \quad (13)$$

where x and y are the denoised and high-SNR spectra, respectively, FFT$(\cdot)$ denotes the one-dimensional fast Fourier transform, $|\cdot|$ is the amplitude of the complex Fourier coefficients, and N is the number of spectral bins (4096 in our case). The loss is computed using the L1 norm to ensure robustness to outliers in the frequency domain.

Table 1 presents the performance comparison of all models on the test set. The spec-DDPM outperforms the other methods, achieving the lowest MAE and spectral loss with the highest SSIM-1D, indicating exceptional denoising capability and structural fidelity.

As depicted in Figure 5, the MAE increases significantly when $T_{\text{eff}} < 5000$ K. Late-type stars, such as K and M dwarfs with effective temperatures below 5000 K, exhibit complex spectral features due to their cooler atmospheres. These spectra are dominated by molecular bands (e.g., TiO, VO) and other complexities arising from lower temperatures. Accurately modeling these features is challenging, as molecular absorption complicates continuum identification and line fitting. This complexity likely explains the higher errors observed in this temperature range. These results align with Figure 6, which indicates that the MAE for K-type stars is notably higher than for the other three stellar types.

4.0.2. Spectral Feature Lines Restoration

To evaluate the ability of our spec-DDPM model to restore critical spectral features, we analyzed its performance on key atomic lines and molecular bands within the spectral range of 4000–8096 Å. These features, including atomic lines such as Hα (6562.81 Å), Hβ (4861.33 Å), Na I D (5889.95/5895.92 Å), O I (7771.94 Å), and Mg I b (5183.60 Å/5172.68 Å), and molecular bands such as TiO (6150–6300 Å, 6650–6850 Å, 7050–7250 Å), MgH (5100–5200 Å), and CN (7900–8090 Å), are essential for characterizing stellar atmospheres and physical properties. The restoration of these features is particularly challenging in low SNR spectra, where noise can obscure sharp atomic lines and broad molecular bands.

We computed the MAE for each feature by comparing the denoised spectra from spec-DDPM and the baseline models (PCA, Wavelet-Haar, Wavelet-db1, and DnCNN) with the

Table 1
Performance Comparison of Different Denoising Models on the Test Set

Model	spec-DDPM	PCA	Wavelet-Haar	Wavelet-db1	DnCNN
MAE	0.0159	0.0435	0.0471	0.0475	0.0227
SSIM-1D	0.9917	0.7053	0.6882	0.6881	0.9750
Spectral Loss	0.8157	1.8175	2.2855	2.3022	1.0603

high-SNR ground truth spectra within the specified wavelength ranges. Additionally, a spectral loss based on the Fourier transform was calculated to assess the model's ability to preserve frequency-domain characteristics, such as the sharpness of atomic lines and the smoothness of molecular bands.

Figure 7 illustrates the MAE for five atomic lines and five molecular bands across the different denoising models. Each line, color-coded by model, represents the MAE averaged over the test set for the respective feature. The spec-DDPM model consistently achieves the lowest MAE across all features, demonstrating superior restoration of both narrow atomic lines and broad molecular bands. Notably, for strong lines like Hα and Hβ, spec-DDPM maintains high fidelity, with MAE values significantly lower than those of DnCNN, which tends to oversmooth these features. For molecular bands, such as the TiO bands prevalent in cooler stars, spec-DDPM effectively captures their broad profiles, whereas PCA and wavelet-based methods struggle to distinguish these bands from continuum noise.

The spectral loss analysis further highlights spec-DDPM's ability to preserve high-frequency components associated with atomic lines and low-frequency components linked to molecular bands. This is particularly evident in the Na I D doublet and the CN band, where spec-DDPM's frequency-domain fidelity surpasses that of the baseline models. These results underscore the model's robustness in handling the diverse spectral features encountered in astronomical data, making it a promising tool for applications requiring precise feature restoration.

4.1. Hyperparameter Tuning

Hyperparameter Tuning for the spec-DDPM model to enhance the performance of the spec-DDPM model, we conducted an extensive experimental analysis to evaluate the effects of two key hyperparameters: κ and p. These parameters are integral to the diffusion process, where κ dictates the overall noise intensity, and p modulates the growth rate of the shift sequence $\{\eta_t\}_{t=1}^{T}$. We systematically tested a range of values for $\kappa \in \{10, 20, 30, 40, 50\}$ and $p \in \{0.3, 0.5, 0.8\}$, assessing their impact on model performance through four metrics: MAE, SSIM-1D, spectral loss, and feature lines loss.

The heatmaps, presented in Figure 8, with color intensity indicating performance quality-darker shades correspond to better outcomes. To ensure consistency across visualizations, the colormap for SSIM-1D was inverted, as higher SSIM-1D values indicate superior performance, unlike the other metrics where lower values are preferable, which demonstrates that the optimal hyperparameter configuration occurs at $\kappa = 10$ and $p = 0.8$. At this setting, the model achieves the lowest MAE, spectral loss, and feature lines loss, while maximizing SSIM-1D. This configuration effectively balances noise reduction with the retention of spectral details, highlighting the critical role of tuning κ and p to optimize denoising performance in the spec-DDPM model.

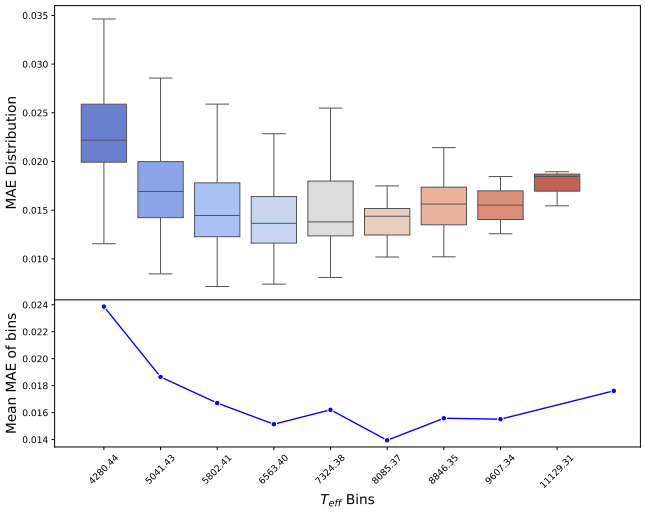


Figure 5. Effective temperature ($T_{\rm eff}$)-dependent denoising errors. Upper panel: Boxplot of MAE across $T_{\rm eff}$ bins, with boxes spanning the 25th–75th percentiles and whiskers showing the 5th–95th percentiles. Lower panel: Trend line of the binned average MAE.

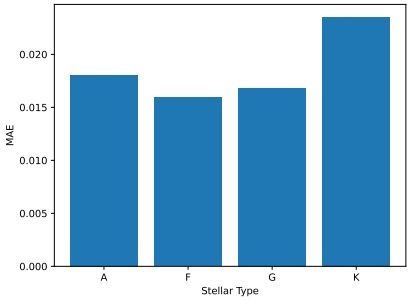


Figure 6. Bar chart showing the Mean Absolute Error (MAE) for different stellar types.

5. Validation

This section aims to validate the effectiveness of the spec-DDPM model in improving the accuracy of stellar parameter estimation under low SNR conditions. By designing a controlled experiment, we compare the performance of a Transformer-based model in estimating effective temperature ($T_{\rm eff}$), surface gravity ($\log g$), α element abundance ([α/M])

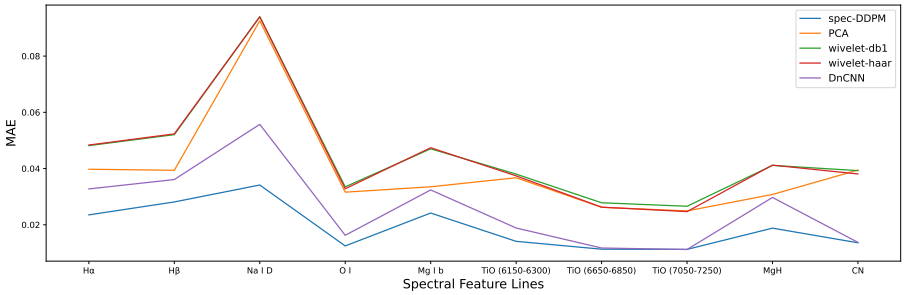


Figure 7. MAE of different denoising models for five atomic lines (Hα, Hβ, Na I D, O I, Mg I b) and five molecular bands (TiO 6150–6300 Å, TiO 6650–6850 Å, TiO 7050–7250 Å, MgH, CN) within the 4000–8096 Å spectral range. Lines are color-coded by model: spec-DDPM, PCA, Wavelet-Haar, Wavelet-db1, and DnCNN.

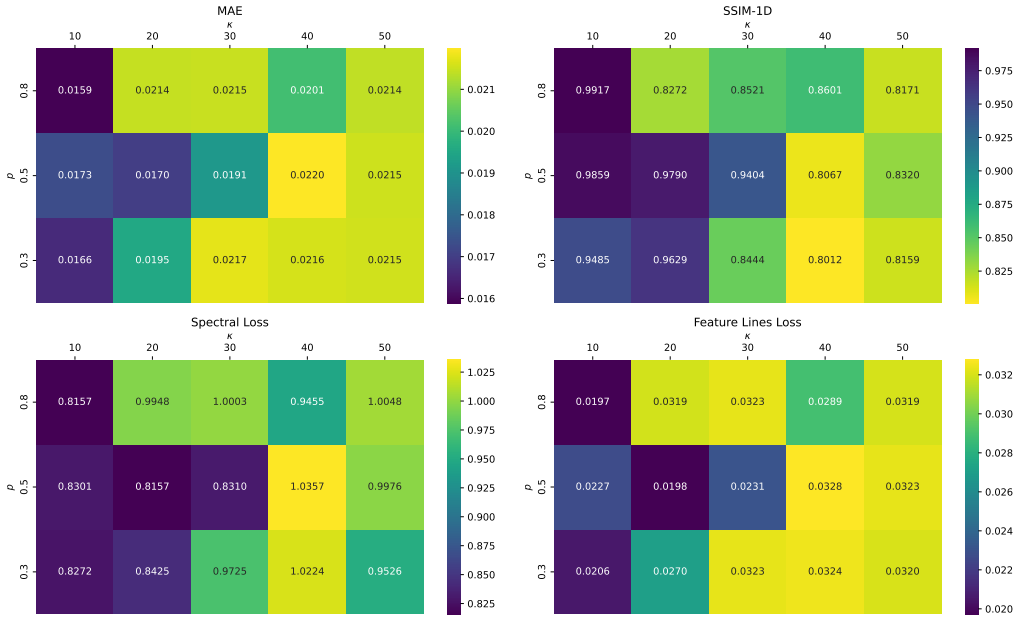


Figure 8. Heatmaps depicting the effects of hyperparameters κ and p on spec-DDPM model performance across four metrics: MAE, SSIM-1D, spectral loss, and feature lines loss. Darker shades indicate better performance, with the SSIM-1D colormap inverted for consistency with other metrics.

and metallicity ([M/H]) with and without spec-DDPM denoising preprocessing. The results demonstrate the model's ability to enhance parameter estimation precision for low-SNR spectra.

5.1. Data Preparation

The experiment utilizes data from the LAMOST DR10 low-resolution survey. We applied the following criteria to select the data set:

1. Spectra with $10 < \mathrm{SNR}_g < 15$ were chosen to focus on low-SNR conditions.
2. Ground-truth values for T_{eff}, $\log g$, $[\alpha/\mathrm{M}]$, and $[\mathrm{M/H}]$ were obtained by cross-matching with the APOGEE DR17's ASPCAP pipeline (Prieto et al. 2008; Nidever et al. 2015; Shetrone et al. 2015; García Pérez et al. 2016; Majewski et al. 2017; Smith et al. 2021; Abdurro'uf et al. 2022).
3. Label quality was controlled by restricting $\log g_{\mathrm{err}} < 50$, $[\alpha/\mathrm{M}]_{\mathrm{err}} < 0.02$, and $T_{\mathrm{eff\,err}} < 50$.

After filtering, a total of 10,968 spectra were obtained. The data set was split into training, validation, and test sets in an 8:1:1 ratio. Preprocessing of the spectra followed the methods described in Section 3. Specifically, T_{eff} was standardized to mitigate the impact of outliers, while $\log g$ and $[\mathrm{M/H}]$ underwent no additional processing.

5.2. Model and Results

A Transformer-based model was designed to estimate T_{eff}, $\log g$, $[\alpha/\mathrm{M}]$, and $[\mathrm{M/H}]$ from the input spectra. The model accepts either raw low-SNR spectra or spectra preprocessed by spec-DDPM as inputs and outputs the predicted stellar

Table 2
Parameters Estimation Performance

	[α/M] (dex)		[M/H] (dex)		log g (dex)		$T_{\rm eff}$ (K)	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
With spec-DDPM	0.0287	0.0422	0.0783	0.1114	0.1740	0.2871	87.05	149.5
Without spec-DDPM	0.0659	0.0871	0.1121	0.1598	0.2498	0.3874	104.8	170.5

parameters. The architecture leverages the Transformer's ability to capture complex patterns in spectral data, making it suitable for this task.

The experimental results are summarized in Table 2:

The results show a clear improvement in estimation accuracy when spec-DDPM denoising is applied. These reductions indicate that spec-DDPM effectively mitigates noise in low-SNR spectra, enabling the Transformer model to extract more reliable features and produce more accurate parameter estimates.

6. Conclusion

In this study, we present spec-DDPM, a novel deep learning model based on DDPMs, tailored for denoising low SNR astronomical spectra. Leveraging the generative power of DDPMs, spec-DDPM effectively enhances spectral quality by recovering critical features such as absorption lines and continuum shapes, outperforming traditional methods like PCA, wavelet-based techniques, and a modified DnCNN model. Performance evaluations using metrics such as MAE, SSIM-1D, and spectral loss confirm its superiority in preserving spectral integrity.

The validation experiments, detailed in Section 5, further demonstrate spec-DDPM's efficacy. By preprocessing low-SNR spectra from the LAMOST DR10 data set (SNR(g) between 10 and 15), spec-DDPM significantly improves the accuracy of a Transformer-based model in estimating stellar parameters—$T_{\rm eff}$, log g, [α/M] and [M/H]. As shown in Table 2, MAE reductions of approximately 17 K for $T_{\rm eff}$, 0.075 dex for log g, 0.0372 dex for [α/M], and 0.034 dex for [M/H] highlight the model's ability to mitigate noise, enabling more precise parameter estimation. These improvements underscore spec-DDPM's potential to enhance the scientific yield of large-scale spectroscopic surveys, particularly for faint or distant stars observed under suboptimal conditions.

The broader implications of this work are substantial for astrophysics. High-quality spectra are vital for deriving accurate stellar parameters, which inform studies of stellar evolution, galactic chemical enrichment, and Milky Way structure. Beyond denoising, spec-DDPM's capacity to correct continuum inaccuracies and restore defective spectra suggests applications in spectral inpainting and radial velocity measurements for exoplanet detection. However, challenges remain. The computational complexity of the diffusion process may hinder real-time applications, and while spec-DDPM performs robustly across diverse stellar types, further optimization could address spectra from extreme cases, such as very cool stars or those with unusual abundance patterns.

Future work could focus on improving computational efficiency, possibly through model compression or faster sampling techniques, and incorporating physical priors to enhance performance on specialized data sets. In summary, spec-DDPM marks a significant step forward in processing low-SNR astronomical spectra, offering a robust tool to maximize the utility of current and future surveys. As spectroscopic data sets grow, such methods will be critical for advancing stellar classification, galaxy evolution studies, and high-precision astrophysical measurements.

Acknowledgments

This study was supported by the National Natural Science Foundation of China (NSFC) under grant Nos. 11873037 and 11803016, the science research grants from the China Manned Space Project with Nos. CMS-CSST-2021-B05 and CMS-CSST-2021-A08, the Natural Science Foundation of Shandong Province under grant Nos. ZR2022MA076, ZR2022MA089 and ZR2024MA063, and the Young Scholars Program of Shandong University, Weihai, under grant No. 2016WHWLJH09 and GHfund A (202202018107).

Thanks to the LAMOST team for releasing reliable spectral data. Guo Shou Jing Telescope (the Large Sky Area Multi-Object Fiber Spectroscopic Telescope LAMOST) is a National Major Scientific Project built by the Chinese Academy of Sciences. Funding for the project has been provided by the National Development and Reform Commission. LAMOST is operated and managed by the National Astronomical Observatories, Chinese Academy of Sciences. The LAMOST website is https://www.lamost.org/.

ORCID iDs

Jingzhen Sun https://orcid.org/0009-0000-8724-497X
Yuhang Zhang https://orcid.org/0009-0009-7973-4077

References

Abdurro'uf, Accetta, K., Aerts, C., et al. 2022, ApJS, 259, 35

Dhariwal, P., & Nichol, A. 2021, in Advances in Neural Information Proc. System, 34 (Red Hook, NY: Curran Associates, Inc.), 8780, arXiv:2105.05233

García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., et al. 2016, AJ, 151, 144

Guo, Z., Liu, J., Wang, Y., et al. 2024, Nature Reviews Bioengineering, 2, 136

Ho, J., Jain, A., & Abbeel, P. 2020, in Advances in Neural Information Proc. Systems, 33 (Red Hook, NY: Curran Associates, Inc.), 6840

Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, AJ, 154, 94

Nidever, D. L., Holtzman, J. A., Allende Prieto, C., et al. 2015, AJ, 150, 173

Prieto, C. A., Majewski, S., Schiavon, R., et al. 2008, AN, 329, 1018

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. 2022, in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition *(New Orleans, LA, USA)* (Los Alamitos, CA: IEEE Computer Society), 10674

Ronneberger, O., Fischer, P., & Brox, T. 2015, in Medical Image Computing and Computer-assisted Intervention-MICCAI 2015: 18th Int. Conf., Proc., Part III 18 *(Munich, Germany)* (Berlin: Springer), 234

Shetrone, M., Bizyaev, D., Lawler, J. E., et al. 2015, ApJS, 221, 24

Smith, M. J., Geach, J. E., Jackson, R. A., et al. 2022, MNRAS, 511, 1808

Smith, V. V., Bizyaev, D., Cunha, K., et al. 2021, AJ, 161, 254

Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N., & Ganguli, S. 2015, arXiv:1503.03585

Song, Y., & Ermon, S. 2019, in Advances in Neural Information Proc. Systems, 32 (Red Hook, NY: Curran Associates, Inc.), 11918

Song, Y., Sohl-Dickstein, J., Kingma, D. P., et al. 2021, arXiv:2011.13456

Wang, D. L., Wang, J. G., & Dong, X. B. 2009, RAA, 12, 509

Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. 2004, ITIP, 13, 600

Yue, Z., Wang, J., & Loy, C. C. 2023, in Advances in Neural Information Proc. Systems, 36 (Red Hook, NY: Curran Associates, Inc.), 13294

Zhang, K., Zuo, W., Chen, Y., Meng, D., & Zhang, L. 2017, ITIP, 26, 3142

Zhao, X., Ting, Y.-S., Diao, K., & Mao, Y. 2023, MNRAS, 526, 1699