Research in Astronomy and Astrophysics, 25:104008 (13pp), 2025 October
https://doi.org/10.1088/1674-4527/adf713
CSTR: 32081.14.RAA.adf713
© 2025. National Astronomical Observatories, CAS and IOP Publishing Ltd. All rights, including for text and data mining, AI training, and similar technologies, are reserved. Printed in China.

The Discovery of New Hot Subdwarf Stars in LAMOST DR12

Yuhang Zhang[1,7], Jiangchuan Zhang[2,7], Yude Bu[1], Zhenxin Lei[3,4], Beining Yang[1], Ke Wang[1], Jingzhen Sun[1], Qinqin Li[1], Siqi Wang[1], Mengmeng Zhang[1], Shanshan Li[1], Zhenping Yi[5,6], Xiaoming Kong[5,6], and Meng Liu[5,6]
[1] School of Mathematics and Statistics, Shandong University, Weihai 264209, China; buyude@sdu.edu.cn
[2] School of Space Science and Technology, Shandong University, Weihai 264209, China
[3] Physics Department, Xiangtan University, Xiangtan 411105, China
[4] Key Laboratory of Stars and Interstellar Medium, Xiangtan University, Xiangtan 411105, China
[5] School of Mechanical, Electrical and Information Engineering, Shandong University, Weihai 264209, China
[6] Shandong Key Laboratory of Intelligent Electronic Packaging Testing and Application, Shandong University, Weihai 264209, China
Received 2025 April 26; revised 2025 June 23; accepted 2025 July 4; published 2025 September 12

Abstract

Hot subdwarf stars are important celestial objects in the study of stellar physics, but the population remains limited. The LAMOST DR12-V1, released in 2025 March, is currently the world's largest spectroscopic database, holding great potential for the search of hot subdwarf stars. In this study, we propose a two-stage deep learning model called the hot subdwarf network (HsdNet), which integrates multiple advanced techniques, comprising a binary classification model in stage one and a five-class classification model in stage two. HsdNet not only achieves high precision with 94.33% and 94.00% in the binary and the five-class classification stages, respectively, but also quantifies the predicted uncertainty, enhancing the interpretability of the classification results through visualizing the model's key focus regions. We applied HsdNet to the 601,217 spectra from the LAMOST DR12-V1 database, conducting a two-stage search for hot subdwarf candidates. In stage one, we initially identified candidates using the binary classification model. In stage two, the five-class classification model was used to further refine these candidates. Finally, we confirmed 1008 newly identified hot subdwarf stars. The distribution of their atmospheric parameters is consistent with that of known hot subdwarf stars. These efforts are expected to significantly advance the research on hot subdwarf stars.

Key words: methods: statistical – techniques: spectroscopic – (stars:) subdwarfs – methods: data analysis

1. Introduction

Hot subdwarf stars are located between the main sequence (MS) and the white dwarf (WD) cooling sequence in the Hertzsprung–Russell diagram (Heber 2016, 2024). Based on the strength of the hydrogen (H) Balmer and helium (He) lines, hot subdwarf stars can be further subdivided into He-poor hot subdwarf stars (e.g., sdB/sdOB/sdO) and He-rich hot subdwarf stars (e.g., He-sdB/He-sdOB/He-sdO) (Moehler et al. 1990; Geier et al. 2017). Most hot subdwarf stars are thought to be in the He-core or He-shell-burning stage. They are located at the extreme blue end of the horizontal branch (HB), also known as the extreme horizontal branch (EHB) stars (Heber 1986). The effective temperatures of hot subdwarf stars typically range from 20,000 to 70,000 K, with surface gravities usually between 5.0 and 6.5 (Heber 2009, 2016, 2024).

Hot subdwarf stars play a vital role in astronomical studies. Hot subdwarf stars typically have low masses around 0.5 $M_{\odot}$ and very thin H-rich envelopes (e.g., $\leqslant$0.01 $M_{\odot}$; Heber 2009). Stars with such thin envelopes inevitably experience significant mass loss during their early evolutionary stages. Understanding the driving mechanism behind this substantial mass loss constitutes an important scientific problem. The formation mechanisms of hot subdwarf stars mainly consist of single-star evolution and binary evolution. Single-star evolution channels include the evolution during the HB and post-HB phases (Dorman et al. 1993; Van Winckel 2003), the He-core flash process (Miller Bertolami et al. 2008), the He mixing model during the red giant branch phase (Sweigart 1997; Tailo et al. 2015), and more. The high incidence of hot subdwarf stars in binaries has made binary evolution a focus of current research (Copperwheat et al. 2011). Binary evolution channels include the merger of two He-WDs (Zhang & Jeffery 2012), the merger of a He-WD with a low-mass MS companion (Zhang et al. 2017; Meng & Luo 2021), stable Roche-lobe overflow (Han et al. 2002; Götberg et al. 2018; Li et al. 2022), common-envelope ejection (Han et al. 2003; Ge et al. 2024; Geier et al. 2024), and more. Moreover, hot subdwarf stars also play a significant role in other astrophysical domains, including revealing the origin of the ultraviolet upturn in elliptical galaxies (Han et al. 2007), probing progenitor star of Type Ia supernovae (Wang & Han 2010; Pelisoli et al. 2021), serving as gravitational wave sources (Kupfer et al. 2018; Lin et al. 2024; Yang et al. 2025), probing the Galactic gravitational potential and putting constraints on the mass of the Galactic

[7] Co-first author.

dark matter halo (Geier et al. 2015; Németh et al. 2016), testing chemical diffusion theory (Michaud et al. 2008, 2011), and more. The primary prerequisite for advancing these studies is to discover more hot subdwarf stars, which makes it necessary for us to conduct a large-scale search based on the latest survey data.

The Large Sky Area Multi-Object Fiber Spectroscopic Telescope (LAMOST) is a large-aperture (effective aperture of 3.6–4.9 m) and wide-field (5°) telescope independently designed by Chinese astronomers (Cui et al. 2012). This telescope can collect 4000 spectra in a single exposure, having the highest spectral acquisition efficiency in the world. The wavelength coverage is 370–900 nm, with a resolution $R \sim 1800$, and a limiting magnitude 19^m (Zhao & Zhao 2012). Since 2014, the LAMOST dataset has been the largest spectral dataset ever obtained (Yan et al. 2022). In 2024 March, LAMOST DR11-V1 released 25.12 million spectra, which was 2.2 times the total number of spectra released by other spectroscopic surveys worldwide at that time. In 2025 March, LAMOST DR12-V1 expanded the capacity to 28.07 million spectra, providing an unprecedented opportunity for a large-scale search of hot subdwarf stars.

Searching for hot subdwarf stars from large-scale survey data typically involves two steps: first, selecting candidates based on color cuts, followed by further confirmation through visual inspection or template matching of spectra (Geier et al. 2011; Vennes et al. 2011; Németh et al. 2012; Geier et al. 2015; Kepler et al. 2015, 2016; Luo et al. 2016). With the release of Gaia DR2 (Gaia Collaboration et al. 2018) and the accumulation of LAMOST spectra, some astronomers have begun combining the photometric magnitudes and parallaxes provided by Gaia with LAMOST spectra to search for hot subdwarf stars (Lei et al. 2018, 2019, 2020, 2023; Luo et al. 2019, 2021).

However, such methods are only suitable for the cross-matched sources between Gaia and LAMOST. Faced with the massive spectra from LAMOST DR12, we aim to directly search for hot subdwarf stars from LAMOST spectra without relying on the photometric magnitudes and parallaxes provided by Gaia. In recent years, machine learning techniques have been applied in astronomy. Bu et al. (2019), Tan et al. (2022), and Cheng et al. (2024) used convolutional neural network (CNN) techniques to search for hot subdwarf stars from LAMOST DR4, DR7, and DR8, respectively, discovering about 200 new hot subdwarf stars in total. However, these works have certain limitations. On the one hand, these algorithms cannot quantify the classification uncertainty, and CNNs are limited by the locality of convolutional kernels, focusing on local features of the spectra, which may neglect their global sequence information. On the other hand, LAMOST DR12 has already accumulated the largest number of spectra, and we urgently need to design more practical and effective algorithms based on the latest data.

In this study, we propose an advanced two-stage deep learning algorithm, the hot subdwarf network (HsdNet), which can not only accurately search for high-confidence candidates from massive spectral data but also quantify the classification uncertainty. We have also applied HsdNet to the LAMOST DR12-V1 database, which shows great potential for expanding the population of hot subdwarf stars. The organization of this paper is as follows. In Section 2, we describe the design motivation behind HsdNet. In Section 3, we present the dataset used for model training, validation and testing in our study. In Section 4, we train HsdNet and evaluate its performance from many perspectives. In Section 5, we conduct a two-stage search for hot subdwarf candidates, and confirm the newly identified hot subdwarf stars among these candidates. Finally, we summarize our work in Section 6.

2. Method

2.1. HsdNet Overview

Although deep learning has been applied in astronomy, the vast majority of algorithms for searching peculiar stars based on spectra still rely on CNNs (Tan et al. 2022, 2023; Cai et al. 2023; He et al. 2024). Prior studies usually used one-dimensional convolutional kernels of size 3 or 5, much smaller than spectral original length, limiting CNNs to local feature extraction due to the kernel locality. Spectral lines exhibit ordered patterns along the wavelengths, which fundamentally characterizes spectral data as sequential data structures with inherent ordering dependencies. Sequence modeling algorithms, particularly the recurrent neural networks (RNNs) and their variants, have been effectively applied to spectral analysis, including spectral classification (Xu et al. 2019), stellar parameter estimation (Helmy et al. 2025), and more. In this study, we integrate these two different mainstream neural network design concepts to propose HsdNet, which consists of four modules: the input module, the PRSE module, the GRU module, and the output module, as shown in Figure 1.

The input module is designed to expand the channel number of the original spectra. Let $x_{B,C,L} \in \mathbb{R}^{B\times C\times L}$ be a three-dimensional tensor where the dimensions represent batch size B, channel number C, and spectrum length L, respectively. As shown in Equation (1), the input module expands the original spectral channel from 1 to 32, where CBP denotes a one-dimensional convolutional layer, a batch normalization layer, and a parametric rectified linear unit (PReLU) activation function (Ding et al. 2018).

$$x_{B,32,L} = \mathrm{Input}(x_{B,1,L}) = \mathrm{CBP}(x_{B,1,L}). \quad (1)$$

The PRSE module consists of three stacked PRSE blocks, which employ multiple local convolutional kernels to extract local spectral features. Each PRSE block contains two one-dimensional convolutional layers, a skip connection (He et al. 2016), and a squeeze-and-excitation (SE) module (Hu et al.

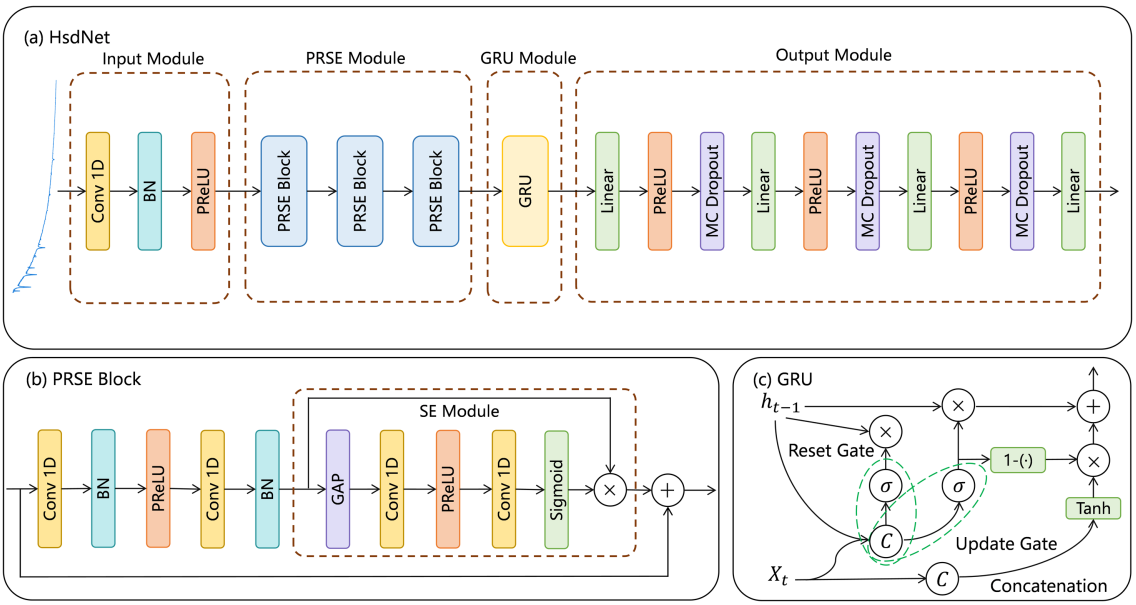


Figure 1. Overview of HsdNet. The normalized spectra are fed into the network. The input module performs channel expansion. The PRSE module extracts key spectral line features using three PRSE blocks, as detailed in panel (b). The GRU module models the flux dependencies across wavelengths, as detailed in panel (c). The output module produces the final predicted results of the network, which contains four linear layers. Between the linear layers, a Monte Carlo dropout layer is incorporated. During the forward propagation process, the network computes the predicted values and evaluates the loss, whereas during the backward propagation process, it optimizes the predicted performance by calculating gradients and updating parameters.

2018). After each block, the length of the feature is halved while the channel number is doubled, as follows:

$$x_{B,256,L/8} = \mathrm{PRSE}(\mathrm{PRSE}(\mathrm{PRSE}(x_{B,32,L}))). \tag{2}$$

The SE module is a type of channel attention mechanism that assigns weights to different channels, enabling selective learning of important channels to enhance the network's capabilities (Hu et al. 2018). Given an input tensor $x_{B,C,L}$, the attention weights for each channel are computed as:

$$w_{B,C,1} = \sigma(F_{[W_1,W_2]}(g(x_{B,C,L}))), \tag{3}$$

where $g(x_{B,C,L}) = \frac{1}{L}\sum_{i=1}^{L} x_{B,C,i}$ represents the channel-wise global average pooling operator, and σ denotes the sigmoid activation function. Let $y = g(x_{B,C,L})$, then $F_{[W_1,W_2]}$ takes the form:

$$F_{[W_1,W_2]}(y) = W_2\,\mathrm{PReLU}(W_1 y). \tag{4}$$

To maintain reasonable model complexity, the dimensions of W_1 and W_2 are set to $C \times (C/r)$ and $(C/r) \times C$ respectively, where the reduction ratio r is set to 16.

The gated recurrent unit (GRU) module captures global dependencies in spectral features. As a gated recurrent neural network, the GRU mitigates the long-term memory loss and gradient explosion problems in traditional RNNs (Chung et al. 2014; Dey & Salem 2017). The input tensor $x_{B,256,L/8}$ is processed by the GRU module to fuse feature with 256 channels into $x_{B,1,L/8}$ as follows:

$$x_{B,1,L/8} = \mathrm{GRU}(x_{B,256,L/8}). \tag{5}$$

The output module employs linear layers to map abstract features to different classes, as follows:

$$x_{B,S} = \mathrm{Output}(x_{B,1,L/8}) = \mathrm{Linears}(x_{B,1,L/8}), \tag{6}$$

where S represents the total number of final predicted classes. As illustrated in Figure 1, except for the last linear layer, we apply a Monte Carlo dropout (Gal & Ghahramani 2016) layer after each linear layer to enhance the network's generalization capabilities, with a dropout probability set to 0.5.

2.2. Monte Carlo Dropout

Standard neural networks attempt to find deterministic optimal parameters. As shown in Equation (7), given a training set D, they usually seek point estimates of parameters through maximum likelihood estimation:

$$w^{\mathrm{MLE}} = \underset{w}{\mathrm{argmax}}\ \log P(D|w). \tag{7}$$

Such models cannot capture the predicted uncertainty, while the optimal parameters often make the model overconfident. In the classification tasks, the final softmax output is usually interpreted as the predicted probability, where high predicted probability with substantial model uncertainty can lead to significant bias from true results. In contrast, Bayesian neural networks attempt to learn optimal parameter distribution $P(w|D)$. For unknown data $\hat{x}$ and label $\hat{y}$, the predicted

distribution of $\hat{y}$ in the Bayesian neural networks is:

$$P(\hat{y}|\hat{x}) = \mathbb{E}_{P(w|D)}[P(\hat{y}|\hat{x}, w)]. \quad (8)$$

Taking the expectation is equivalent to integrating over infinitely many models sampled from $P(w|D)$. The core challenge in Bayesian deep learning is finding the optimal posterior distribution $P(w|D)$, which can be approximated through variational inference with reparameterization tricks (Blundell et al. 2015; Kingma et al. 2015). However, this approach incurs high computational complexity.

Embedding a Monte Carlo dropout layer between the hidden layers of a fully connected neural network achieves effects comparable to variational inference without increasing computational complexity, significantly improving the convergence efficiency of the network (Maeda 2014; Gal & Ghahramani 2016). Furthermore, Gal & Ghahramani (2016) demonstrated that the dropout optimization process was equivalent to minimizing the Kullback–Leibler divergence between a variational distribution and the posterior of a full deep Gaussian process. For a deep neural network with L hidden layers, we sample T sets of vectors from the Bernoulli distribution $\{\mathbf{z}_1^t,\ldots,\mathbf{z}_L^t\}_{t=1}^T$ to implement stochastic dropout with $\mathbf{z}_i^t = [z_{i,j}^t]_{j=1}^{K_i}$, where K_i denotes the number of neurons in the ith hidden layer. The weights of the L hidden layers are $[\boldsymbol{w}_1^t,\ldots,\boldsymbol{w}_L^t]_{t=1}^T$, yielding the approximate predicted distribution $P(\hat{y}|\hat{x})$:

$$P(\hat{y}|\hat{x}) \approx \frac{1}{T}\sum_{t=1}^{T} P(\hat{y}|\hat{x}, \boldsymbol{w}_1^t, \ldots, \boldsymbol{w}_L^t). \quad (9)$$

3. Data

3.1. Dataset Setup

To design and evaluate the classifier, we constructed two sample sets: a positive sample set consisting of hot subdwarf stars and a negative sample set consisting of non-hot subdwarf stars. For the positive sample set, we cross-matched the hot subdwarf stars provided by Culpan et al. (2022), Tan et al. (2022), and Cheng et al. (2024) with the LAMOST DR12-V1 catalog, and screened out these spectra with signal-to-noise ratios (S/N) $\leqslant$ 10 in either the g-band or r-band as well as the redshift anomalies. Finally, we obtained 2626 spectra corresponding to 1668 hot subdwarf stars.

Referring to Bu et al. (2019), Tan et al. (2022), and Cheng et al. (2024), we selected four types of stars to construct the negative sample set: WDs, blue horizontal-branch stars (BHBs), A-type stars, and B-type stars. Since there are relatively very few O-type stars in the LAMOST DR12-V1 database, we did not consider them when constructing the negative sample space. Among them, WDs were obtained by cross-matching the catalog from Kong & Luo (2021) with the LAMOST DR12-V1 catalog. BHBs were obtained by cross-matching the catalog from Vickers et al. (2021) with the LAMOST DR12-V1 catalog. Similar to the selection criteria for the positive samples, we screened out the above spectra with S/N $\leqslant$ 10 in either the g-band or r-band as well as redshift anomalies, and randomly selected 3000 spectra for each class respectively. We used LAMOST_uid to remove duplicates between the A- and B-type spectra satisfying the above criteria in the LAMOST DR12-V1 catalog and known hot subdwarf stars from Culpan et al. (2022), Tan et al. (2022), Cheng et al. (2024), as well as other negative samples. Finally, we randomly selected 3000 A- and B-type spectra, respectively. Simultaneously, we used the spectral analysis tool XTGRID (Németh et al. 2012) to derive atmospheric parameters for these spectra, further confirming the absence of hot subdwarf stars among these A- and B-type stars. Each spectrum in the positive and negative sample sets has a unique spectrum identifier (obsid) in LAMOST. The specific composition of our sample set is shown in Table 1.

Table 1
Composition of Positive and Negative Sample Sets

Sample Type	Spectral Type	Number
Positive	Hot subdwarf	2626
Negative	WD	3000
	BHB	3000
	A-type	3000
	B-type	3000

3.2. Data Preprocessing

To obtain the data required for our experiments, we need to preprocess the 14,626 LAMOST spectra following these steps:

(1) *Wavelength correction with radial velocity*. We performed wavelength correction using the following equation:

$$w_{\rm new}(n) = \frac{w(n)}{1+z}, \quad (10)$$

where $w(n)$ represents the wavelengths of the original observed spectrum, and the redshift z is calculated by the LAMOST pipeline (Luo & Zhao 2001).

(2) *Interpolation*. Due to variations in the wavelength coverage of each spectrum in the original dataset, they cannot be directly input into our model for training. Therefore, interpolation is necessary for unification. Through linear interpolation, the wavelength range of each spectrum in the original dataset is adjusted to 4000–8095 Å, with a feature length of 4096, which covers important lines required for visible band spectral analysis.

(3) *Normalization*. Normalizing the spectral flux before training can eliminate the impact of outliers on model training, accelerate model convergence, and also suppress the adverse effects of imperfect spectral flux calibration. The following linear function normalization equation can adjust the flux of

each spectrum in the original dataset to the range of 0–1:

$$F' = \frac{F - F_{\min}}{F_{\max} - F_{\min}}, \tag{11}$$

where F represents the flux of the original spectrum, and $F_{\min}$ and $F_{\max}$ denote the minimum and the maximum flux of the original spectrum, respectively.

4. Modeling and Performance Analysis

4.1. Performance Metric

To quantify the generalization performance of HsdNet, we adopted accuracy, precision, recall, F1-score, receiver operating characteristic (ROC) curve, area under curve (AUC) value, and precision-recall (P-R) curve. These metrics collectively enable a comprehensive assessment of HsdNet's capability to discriminate between hot subdwarf stars and non-hot subdwarf stars, providing valuable insights into its classification efficacy.

Accuracy refers to the proportion of correctly classified samples among all samples. Precision represents the proportion of hot subdwarf stars among those predicted as hot subdwarf stars. Recall denotes the proportion of hot subdwarf stars that are correctly identified. F1-score is the harmonic mean of precision and recall, balancing both metrics. These evaluation metrics are calculated as follows:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP}+\text{FP} + \text{FN}+\text{TN}}, \tag{12}$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \tag{13}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \tag{14}$$

$$\text{F1-score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}, \tag{15}$$

where TP (true positive) represents the number of hot subdwarf stars correctly classified by HsdNet, TN (true negative) denotes the number of non-hot subdwarf stars correctly classified by HsdNet, FP (false positive) represents the number of non-hot subdwarf stars incorrectly classified as hot subdwarf stars by HsdNet, and FN (false negative) denotes the number of hot subdwarf stars incorrectly classified as non-hot subdwarf stars by HsdNet.

In ROC curves, the x-axis represents the false positive rate (FPR) and the y-axis denotes the true positive rate (TPR). Here, FPR corresponds to the rate of misclassifying non-hot subdwarf stars as hot subdwarf stars, while TPR denotes the rate of correctly classifying hot subdwarf stars. These metrics are defined as:

$$\text{FPR} = \frac{\text{FP}}{\text{TN} + \text{FP}}, \tag{16}$$

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}}. \tag{17}$$

Table 2
Performance of Different-stage Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
Binary	97.20	94.33	92.18	93.25
Five-class	92.89	94.00	94.95	94.47

The AUC quantifies the overall performance of the ROC curve. A model achieves superior discriminative capability when its ROC curve approaches the upper-left corner. For our imbalanced sample sets, the P-R curve provides complementary insights. This curve plots precision (y-axis) against recall (x-axis), illustrating the trade-off between these metrics. The P-R curve closer to the top-right corner indicates better model performance.

4.2. Classification Construction

In this section, we present the construction of HsdNet (a two-stage deep learning model), including a binary classification model and a five-class classification model. During the experiment, we randomly divided the 14,626 spectra into training, validation and test sets in an 8:1:1 ratio, which were used to train the model, determine the model and evaluate the model, respectively. The training process was carried out in batches, and the batch size was set to 128. In terms of the loss function, we adopted the cross-entropy loss, a common choice for classification tasks.

In the field of deep learning, the choice of optimization algorithms, hyperparameter configurations, and training techniques can directly affect the performance of the model. In this study, we employed the Adam optimizer (Kingma 2014). Empirically, we set the parameters for the Adam optimizer with the first-moment decay rate $\beta_1 = 0.9$, the second-moment decay rate $\beta_2 = 0.999$, and a small constant $\epsilon = 1e^{-8}$ to assure steady model learning. The initial learning rate is $5e^{-5}$ and it decreased by a factor of 5 after every 10 training epochs. To address potential overfitting, we applied L_2 regularization with a coefficient of $1e^{-3}$. These methods helped constrain the model weights to converge to smaller values, leading to more consistent and smoother predictions.

We first trained the binary classification model using hot subdwarf stars as positive samples and the remaining four-class stars (WDs, B-type stars, BHBs, and A-type stars) as negative samples. We saved the model which performed best on the validation set and evaluated it on the test set, with an overall accuracy of 97.20%, a precision of 94.33%, a recall of 92.18%, and an F1-score of 93.25% for hot subdwarf stars, as shown in the first row of Table 2.

However, as shown in Table 1, the number of negative samples is approximately 4.6 times the number of positive samples in the binary classification stage, which may affect the classification result. To mitigate the impact of class imbalance and improve the classification result, we tried to replace the cross-entropy loss with the Focal loss for training a binary

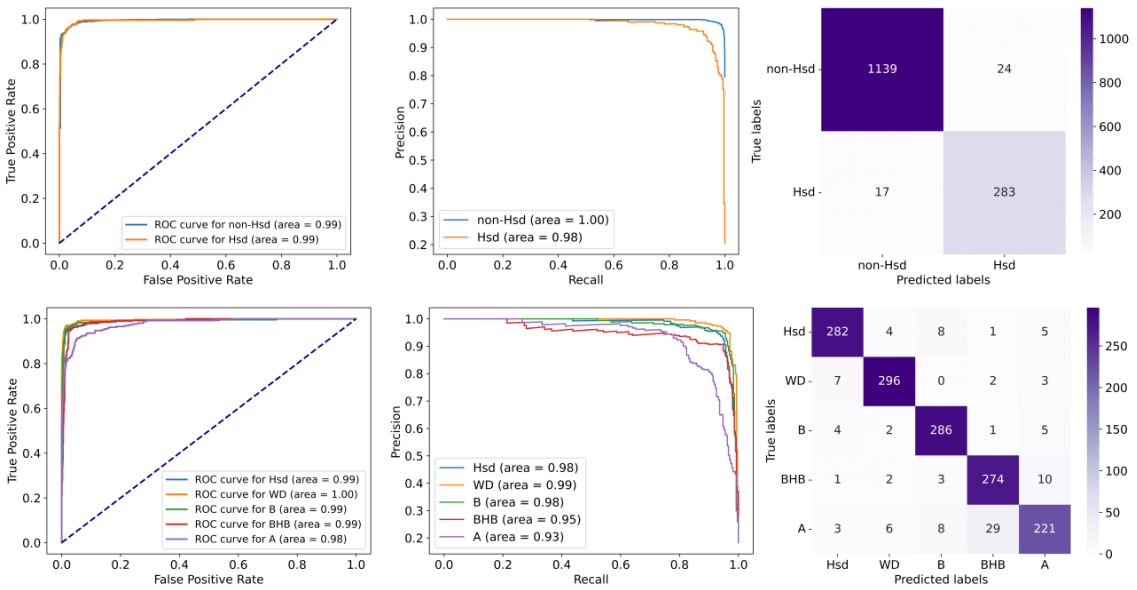


Figure 2. ROC curves, P-R curves, and confusion matrices of different-stage models. The upper panel displays (from left to right) the ROC curve, P-R curve, and confusion matrix of the binary classification model, while the lower panel presents the corresponding plots for the five-class classification model in the same order.

classification model. Originally designed to handle extreme class imbalance in object detection, the Focal loss introduces a weighting factor $\alpha_t \in [0, 1]$ for class t and a modulating factor $(1 - p_t)^\gamma$ with a tunable focusing parameter $\gamma \geqslant 0$ compared to the standard cross-entropy loss, formulated as:

$$FL(p_t) = -\alpha_t(1 - p_t)^\gamma \log(p_t), \tag{18}$$

where p_t denotes the predicted probability for class t, α_t balances class importance, and $(1 - p_t)^\gamma$ focuses hard-to-classify samples (Lin et al. 2017). During our experiments, γ was set to 2 according to the methodology outlined in Lin et al. (2017), and α_t was set to 4.6 based on the positive-to-negative ratio in our sample set. On the test set, the optimal model achieved an overall accuracy of 96.17%, a precision of 94.67%, a recall of 87.65%, and an F1-score of 91.03% for hot subdwarf stars. We observed that while the Focal loss tried to mitigate the impact of class imbalance by loss weighting and slightly improved precision, it led to significant recall degradation. Compared to the cross-entropy loss, the Focal loss demonstrated limited performance gains in our task. Therefore, we retained the optimal binary classification model trained by the cross-entropy loss.

In order to further decontaminate other-class stars to obtain high-confidence hot subdwarf candidates, we need to perform multiple-class classification on the basis of the binary classification results. Therefore, we adapted the negative sample labels and modified the number of neurons in the output layer to 5 to continue to train the five-class classification model, which was evaluated to achieve an overall accuracy of 92.89%, a precision of 94.00%, a recall of 94.95%, and an F1-score of 94.47% for hot subdwarf stars on the test set, as shown in the second row of Table 2. We also plotted the ROC curves, P-R curves and confusion matrices for both the binary and the five-class classification models, respectively, as shown in Figure 2.

Table 3
Performance Comparison with Other Methods

Metric	Task	Method			
		HsdNet (Ours)	CNN +SVM (Bu)	CNN (Tan)	Se-ResNet +SVM (Cheng)
Accuracy (%)	Binary	**97.20**	86.23	96.45	96.75
	Five-class	**92.89**	85.00	90.98	90.53
Precision (%)	Binary	**94.33**	76.07	89.67	92.99
	Five-class	94.00	87.35	93.67	**94.77**
Recall (%)	Binary	92.18	43.28	**92.76**	90.23
	Five-class	**94.95**	87.96	94.93	91.80
F1-score (%)	Binary	**93.25**	55.17	91.19	91.59
	Five-class	**94.47**	87.65	94.30	93.26

Note. The bold values in the table highlight the best performance across different methods.

4.3. Comparative Analysis

To demonstrate HsdNet's advantage over prior dedicated models for hot subdwarf star identification, we compared its performance with that of the CNN+SVM model from Bu et al. (2019), the CNN model from Tan et al. (2022), and the Se-ResNet+SVM model from Cheng et al. (2024). The results are shown in Table 3, where the bold values represent the best

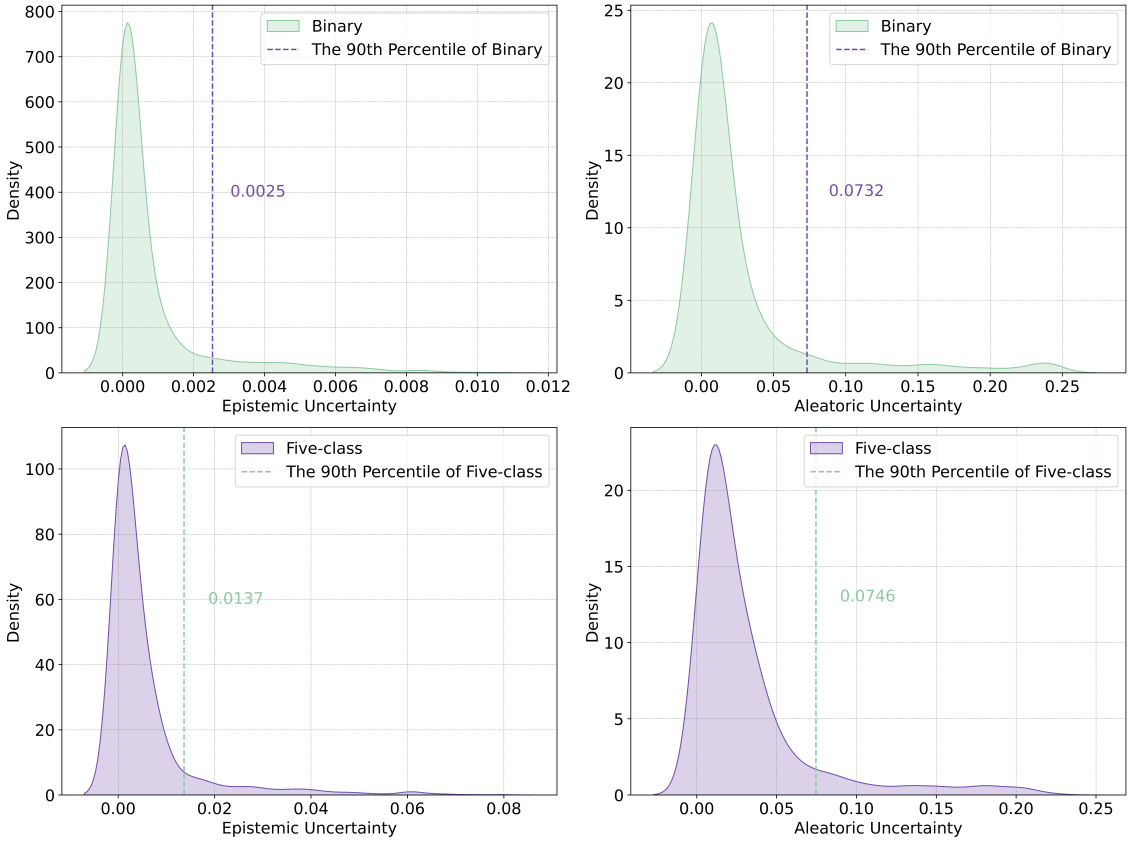


Figure 3. Predicted uncertainty distributions of HsdNet. The kernel density estimation (KDE) curves in the upper and lower panels represent the uncertainty distributions of the binary classification model and the five-class classification model, respectively. The purple dashed lines represent the 90th percentiles of the predicted uncertainty for the binary classification model, while the green dashed lines denote those for the five-class classification model.

performance. It can be seen that, whether for binary classification or five-class classification, HsdNet significantly outperforms the CNN+SVM model from Bu et al. (2019). For the CNN model from Tan et al. (2022) and the Se-ResNet +SVM model from Cheng et al. (2024), although our improvements in classification metrics are relatively modest, we have implemented approximate Bayesian inference in HsdNet. It not only provides classification results but also quantifies uncertainty, thereby enabling further assessment of the model's robustness and enhancing the credibility of astronomical decision-making, and the details can be found in Section 4.4.

4.4. Uncertainty Analysis

Predicted uncertainty in deep neural networks can generally be categorized into epistemic uncertainty and aleatoric uncertainty (also known as statistical uncertainty) (Der Kiureghian & Ditlevsen 2009). The former is associated with the network itself, meaning that additional samples can improve its performance, while the latter arises from inherent noise in the data and remains irreducible even with increased sample size. Kwon et al. (2020) proposed a method to quantify the predicted uncertainty in classification based on Bayesian neural networks. This method leverages the relationship between variance and the mean of multinomial random variables, avoiding additional parameter estimation for variance. It demonstrates good numerical stability and allows for straightforward interpretation. Specifically, the method is implemented as follows:

$$\begin{aligned} &\mathrm{Var}_q(p(y^*|x^*)) \\ &= \mathbb{E}_q[yy^T] - \mathbb{E}_q[y]\mathbb{E}_q[y]^T \\ &= \underbrace{\mathrm{diag}(\bar{p}) - \frac{1}{T}\sum_{t=1}^{T}\hat{p}_t\,\hat{p}_t^T}_{\text{Aleatoric}} + \underbrace{\frac{1}{T}\sum_{t=1}^{T}(\hat{p}_t - \bar{p})(\hat{p}_t - \bar{p})^T}_{\text{Epistemic}}, \end{aligned} \tag{19}$$

where $\mathrm{Var}_q(p(y^*|x^*))$ represents the variance of the conditional probability distribution of the target variable y^* given the input x^*, reflecting the predicted uncertainty. T denotes the number of stochastic forward passes, which is set to 100 in this paper. $\hat{p}_t$ represents the predicted probability obtained from the tth stochastic forward pass, defined as $\hat{p}_t = \mathrm{Softmax}(f_{\omega_t}(x^*))$, where $f_{\omega_t}(\cdot)$ and ω_t denote the model and its corresponding parameters during the tth stochastic forward pass, respectively.

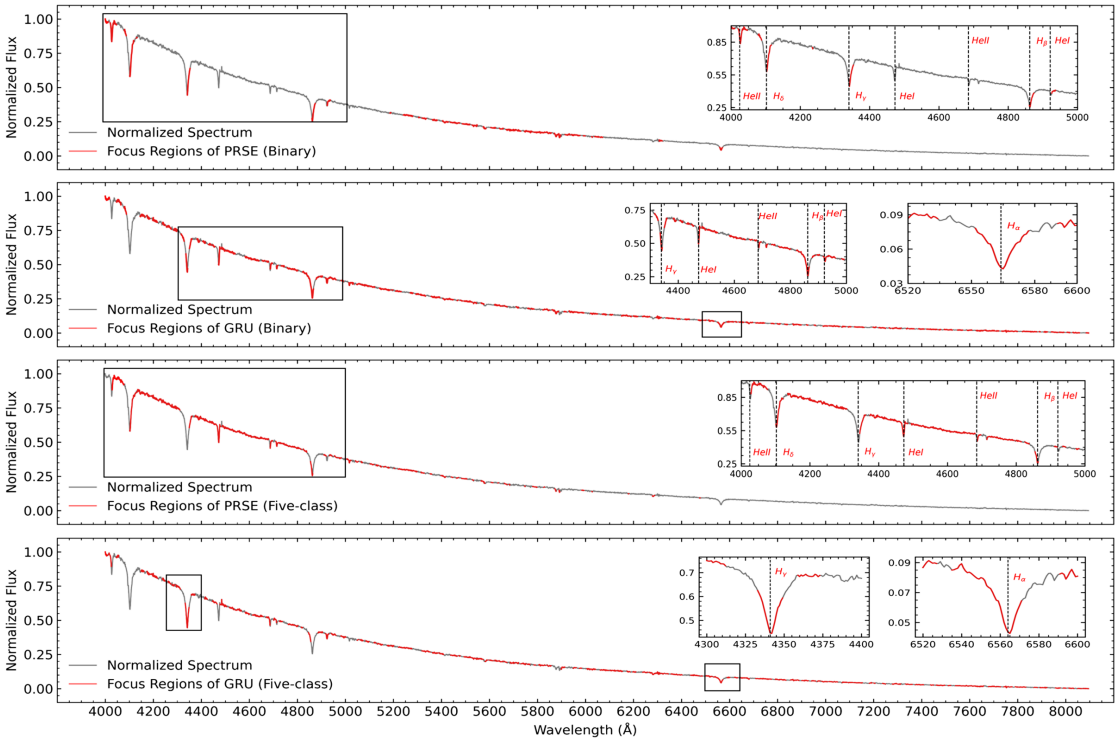


Figure 4. Focus regions (red solid lines) during hot subdwarf star identification for the binary classification model (upper two panels) and the five-class classification model (lower two panels). Black dashed lines mark some important H Balmer lines (H_δ, H_γ, H_β, and H_α) and He lines (He I and He II). The LAMOST_obsid of this spectrum is 157006179.

$\overline{p}$ is the mean predicted probability averaged over T stochastic forward passes, defined as $\overline{p} = \frac{1}{T}\sum_{t=1}^{T}\hat{p}_t$.

HsdNet implements approximate Bayesian inference by introducing Monte Carlo dropout. In this study, we calculated both epistemic and aleatoric uncertainties for both the binary and five-class classification models based on the aforementioned method. Their density distributions are shown in Figure 3, with corresponding 90th percentiles marked. For the binary classification model, 90% of samples exhibit epistemic uncertainty $\leqslant$0.0025 and aleatoric uncertainty $\leqslant$0.0732. For the five-class classification model, 90% of samples exhibit epistemic uncertainty $\leqslant$0.0137 and aleatoric uncertainty $\leqslant$0.0746. These results suggest that HsdNet demonstrates strong robustness, with the data itself possessing relatively low uncertainty, thereby the classification outcomes are highly reliable.

4.5. Interpretability Analysis

We demonstrated the interpretability of the classification outcomes by analyzing the key spectral features that HsdNet focused on when identifying hot subdwarf stars. Gradient-weighted class activation mapping (Grad-CAM; Selvaraju et al. 2017) is a visualization method for deep neural networks, highlighting the data regions most influential to the model's decision-making. Grad-CAM works by backpropagating the gradients through the network and mapping them to the input data, quantifying the impact of different regions on predictions. This technique facilitates understanding and interpreting the internal mechanisms of deep neural networks.

We applied Grad-CAM to visualize the focus regions of the PRSE and GRU modules in both the binary and the five-class classification models. As shown in Figure 4, the PRSE module emphasizes local spatial features near the H_δ, H_γ, He I, and He II lines, while the GRU module focuses more on global sequential information across the entire spectrum. In general, HsdNet can not only focus on the key spectral lines used by astronomers for manual hot subdwarf star confirmation, but also capture the dependencies between the corresponding fluxes at different wavelengths of the entire spectrum.

5. Results and Discussion

5.1. Searching for Hot Subdwarf Candidates from LAMOST DR12

From the LAMOST DR12-V1 database, we downloaded 601,217 low-resolution spectra by applying filtering criteria, including spectra with S/N $>$ 10 in both the *g*- and *r*-bands, non-anomalous redshifts, and exclusion of FGKM-type stars.

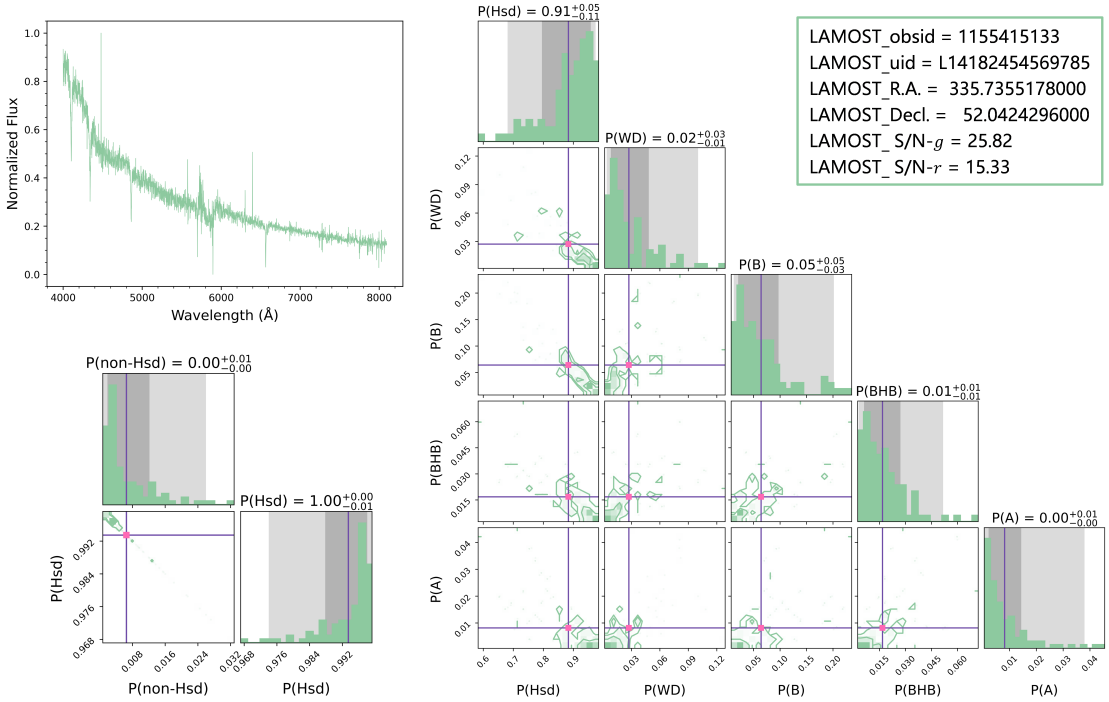


Figure 5. A hot subdwarf candidate spectrum identified through a two-stage search and its predicted probability distribution. The top-left panel displays the normalized spectrum. The top-right panel provides the unique spectrum identifier (LAMOST_obsid), g-band S/N (LAMOST_S/N-g), r-band S/N (LAMOST_S/N-r), unique source identifier (LAMOST_uid), R.A. (LAMOST_R.A.), and decl. (LAMOST_Decl.). The contour plot in the bottom-left panel illustrates the two-dimensional distribution of P(class = Hsd) and P(class = non-Hsd) from stage one. The contour plot in the bottom-right panel presents the two-dimensional distribution of P(class = Hsd), P(class = WD), P(class = B), P(class = BHB), and P(class = A) from stage two. The green histograms represent the predicted probability distribution based on 100 samples, the purple solid lines indicate the mean predicted probability for each class, and the gray shaded regions cover the second to 98th percentiles (light gray) and 16th to 84th percentiles (dark gray).

Table 4
Search Results under Different Classification Thresholds in Stage Two

Metric	Classification Threshold				
	0.5	0.6	0.7	0.8	0.9
Purity	42.24%	48.02%	54.67%	64.79%	79.57%
Completeness	89.33%	88.49%	86.51%	84.71%	77.58%

After undergoing a data preprocessing similar to that used during model training in Section 3.2, we adopted a two-stage approach to identify hot subdwarf candidates.

In stage one, we applied the binary classification model to the 601,217 spectra, identifying 13,217 spectra corresponding to 11,142 hot subdwarf candidates, including 1626 hot subdwarf stars confirmed by Culpan et al. (2022), Tan et al. (2022), Lei et al. (2023), and Cheng et al. (2024), and 9516 candidates.

In stage two, we employed the five-class classification model to the 13,217 candidate spectra. With a classification threshold (defined as the probability value or confidence level for a spectrum being a hot subdwarf star spectrum) set at 0.5, we identified 6470 spectra corresponding to 5078 hot subdwarf candidates, including 1569 previously confirmed hot subdwarf stars and 3509 candidates. In Section 5.2, we will confirm true hot subdwarf stars among the candidates. Figure 5 shows a hot subdwarf candidate spectrum identified by HsdNet, along with its two-stage predicted probability distribution.

5.2. Further Confirm New Hot Subdwarf Stars

In this section, we will confirm true hot subdwarf stars among the candidates identified in Section 5.1 by employing the spectroscopic analysis tool XTGRID. This program fits the profiles of H Balmer lines and He lines of the observed spectra with SYNSPEC synthetic spectra (version 49; Lanz & Hubeny 2007) calculated from non-local thermodynamic equilibrium TLUSTY model atmospheres (version 204; Hubeny & Lanz 2017). The best-fitting model is searched for iteratively with a successive approximation method along the steepest gradient of the χ^2 field. Parameter uncertainties are estimated by mapping the $\Delta\chi^2$ field until the 60% confidence level at the given number of free parameters is reached.

Table 5
Catalog of Partial Newly Identified Hot Subdwarf Stars

LAMOST_obsid	LAMOST_uid	LAMOST_R.A.	LAMOST_Decl.	LAMOST_S/N-g	LAMOST_S/N-r	$T_{\rm eff}$ (K)	$\log g$(cm s^{-2})	$\log(n{\rm He}/n{\rm H})$
1223311074	L16196229753596	105.87438	9.783824	80.04	65.95	41355 ± 168	5.56 ± 0.03	−0.35 ± 0.03
1223313148	L16200088236201	107.962284	8.806504	197.81	215.27	42511 ± 1658	5.36 ± 0.11	−2.39 ± 0.1
1223402097	L10033966293296	133.415152	−1.401921	34.9	48.7	38346 ± 770	6.11 ± 0.1	−2.42 ± 0.24
1223508223	L15426159663168	176.124705	16.746388	25.08	17.11	36756 ± 680	5.31 ± 0.09	−2.15 ± 0.14
1223607090	L15152473075475	204.698157	32.751565	42.5	52.68	29864 ± 732	5.1 ± 0.11	−4.0 ± 0.74
1224101063	L15888039353054	133.301872	62.745896	390.08	520.92	22661 ± 329	5.0 ± 0.08	0.09 ± 0.06
1228806094	L14321052533447	269.941355	14.655779	63.02	75.38	35225 ± 1622	5.51 ± 0.22	−2.95 ± 0.67
1228909182	L16450291443802	141.670342	40.280054	68.7	84.97	28103 ± 1020	5.1 ± 0.13	−3.57 ± 0.56
1229203075	L15339653582604	238.597213	31.159601	42.68	45.44	40102 ± 839	6.08 ± 0.1	−3.06 ± 0.58
1229403200	L15176080909139	198.120022	45.688716	83.35	94.1	39691 ± 696	5.61 ± 0.07	−3.7 ± 1.21
1229515165	L15340886090555	234.022576	30.606782	57.58	66.66	38774 ± 657	6.2 ± 0.09	−3.29 ± 0.8
1229515174	L15340947731676	234.37241	30.74829	39.6	48.06	38811 ± 739	6.1 ± 0.1	−3.7 ± 1.6
1235703054	L14749001647796	241.903009	61.861108	79.9	95.9	31570 ± 769	5.1 ± 0.12	−3.51 ± 0.61
1235703100	L14748940636988	242.088426	62.498485	125.4	143.86	33453 ± 531	5.93 ± 0.09	−4.06 ± 2.07
1235703181	L14779030969724	241.561625	62.670881	52.09	60.13	30144 ± 599	5.1 ± 0.09	−3.01 ± 0.24
1235904229	L15399964966967	179.242618	7.257803	38.2	47.5	35499 ± 686	5.35 ± 0.11	−1.54 ± 0.06
1236002144	L14970249348650	210.364884	6.019638	32.94	44.63	38976 ± 773	6.24 ± 0.11	−2.78 ± 0.42
1236110212	L14775964136043	230.169823	58.721666	45.79	57.46	36041 ± 674	5.59 ± 0.08	−4.0 ± 0.53
1236204025	L14027023786046	284.242833	39.143242	240.29	295.16	22654 ± 254	5.0 ± 0.03	−0.77 ± 0.04
1236209031	L14024721452498	285.674181	40.238672	144.57	199.31	36891 ± 530	6.08 ± 0.22	0.21 ± 0.58
1246407016	L14832785332079	212.929846	67.630907	92.72	109.67	34685 ± 1300	6.02 ± 0.25	−2.08 ± 0.39
1246514074	L15310884066884	240.297392	41.471828	110.78	143.02	39440 ± 609	6.06 ± 0.09	−3.7 ± 1.41
1246712049	L14473220571506	238.253597	14.725215	433.48	479.57	28298 ± 1571	5.98 ± 0.19	−3.81 ± 1.78
1246807207	L14044881469714	288.637879	34.440269	42.54	37.88	34656 ± 475	5.21 ± 0.06	−2.75 ± 0.19
1246815036	L14082814112651	288.460212	37.523746	40.83	17.16	31499 ± 501	5.22 ± 0.07	−2.29 ± 0.1
1246815132	L14079385707573	287.854864	36.927072	47.52	37.34	23704 ± 623	5.0 ± 0.07	−2.59 ± 0.11
1246815203	L14079301914591	287.952784	37.53848	43.18	34.85	36841 ± 324	5.54 ± 0.05	−1.41 ± 0.04
1246815231	L14079729599645	286.960111	37.467833	75.92	73.17	36148 ± 295	5.41 ± 0.04	−1.55 ± 0.06
1246913245	L14046997693888	290.087349	36.796165	92.72	139.61	24169 ± 3580	5.12 ± 0.44	−1.14 ± 0.37
1246914151	L14079066494791	285.534713	36.651404	102.27	82.94	41780 ± 4711	5.1 ± 0.23	−3.09 ± 1.58

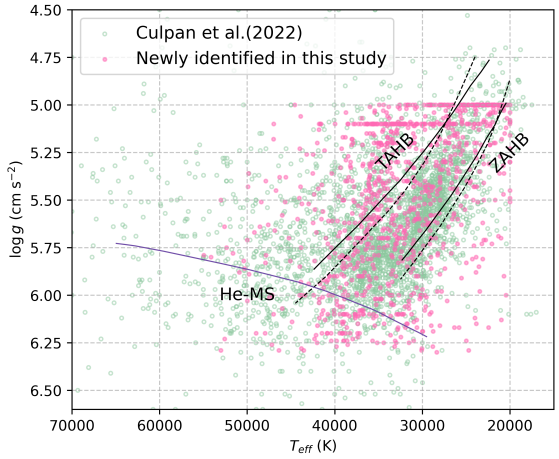


Figure 6. $T_{\rm eff}$–log g diagram for the 1008 newly identified hot subdwarf stars in this study. The zero-age HB (ZAHB) and terminal-age HB (TAHB) are plotted for two different He contents $Y = 0.24$ (black dashed lines) and $Y = 0.40$ (black solid lines) from Latour et al. (2018). The He-MS from Paczyński (1971) is marked by the purple solid line.

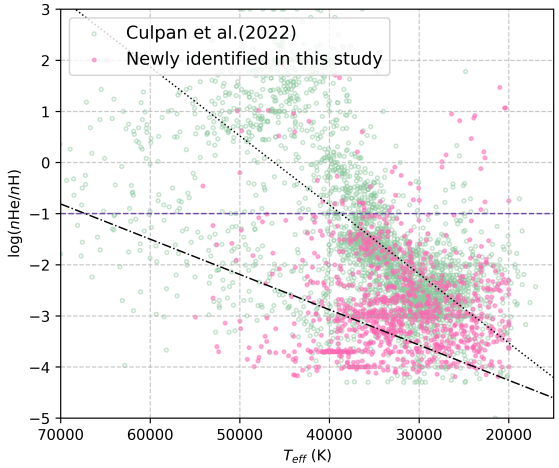


Figure 7. $T_{\rm eff}$-log(nHe/nH) diagram for the 1008 newly identified hot subdwarf stars in this study. The purple dashed line denotes the solar He abundance (e.g., log(nHe/nH) $= -1$). The black dotted line is the linear regression line for the He-rich sequence fitted by Edelmann et al. (2003), while the black dotted-dashed line is the linear regression line for the He-weak sequence fitted by Németh et al. (2012).

The atmospheric parameters of the identified candidates, such as effective temperature ($T_{\rm eff}$), surface gravity (log g), and surface He abundance (log(nHe/nH)), were obtained by the method described above. As in the previous studies (Lei et al. 2018, 2019, 2020; Cheng et al. 2024), we regarded candidates with $T_{\rm eff} \geqslant 20{,}000$ K and $5.0 \leqslant \log g \leqslant 6.3$ as hot subdwarf stars, while candidates exhibiting $6.3 < \log g \leqslant 6.5$ may be potential WDs.

Among the 9516 candidates identified in stage one, 1008 were confirmed as true hot subdwarf stars. Thus, the purity of

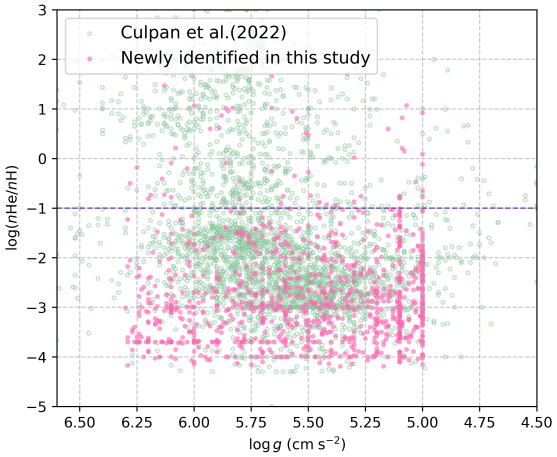


Figure 8. log g–log(nHe/nH) diagram for the 1008 newly identified hot subdwarf stars in this study. The purple dashed line denotes the solar He abundance (e.g., log(nHe/nH) $= -1$).

the candidates identified solely by the binary classification model was 23.64% (purity = (1626 + 1008)/11,142), where purity is defined as the proportion of true hot subdwarf stars among all the identified candidates. Meanwhile, the completeness of the hot subdwarf stars identified solely by the binary classification model was 91.43% (completeness = 1525/1668), where completeness refers to the proportion of the hot subdwarf stars successfully recovered by HsdNet among the original hot subdwarf stars in the training, validation, and test sets.

Then, we investigated the search results under different classification thresholds in stage two from purity and completeness perspectives. As summarized in Table 4, applying the five-class classification model to the binary-classification-identified candidates significantly improved purity without substantial loss of completeness, demonstrating the practical effectiveness of our two-stage model. However, while higher classification thresholds in stage two further increased purity, they also reduced completeness.

Table 5 presents the main information of partial newly identified hot subdwarf stars (the complete version of 1008 entries is available online at https://github.com/SDUZYH-star/HsdNet). Columns (1)–(6) display LAMOST_obsid, LAMOST_uid, LAMOST_R.A., LAMOST_Decl., LAMOST_S/N-g, and LAMOST_S/N-r. Columns (7)–(9) provide $T_{\rm eff}$, log g, log(nHe/nH), and their estimated errors.

The relationships between atmospheric parameters of the 1008 newly identified hot subdwarf stars in this study, e.g., $T_{\rm eff}$-log g, $T_{\rm eff}$–log(nHe/nH), and log g-log(nHe/nH), are shown in Figures 6–8, respectively. The green points represent known hot subdwarf stars from Culpan et al. (2022). The pink points represent the newly identified hot subwarf stars in this study. Figure 6 shows that many newly identified hot subdwarf stars are well located in the region defined by the zero-age HB

(ZAHB) and terminal-age HB (TAHB) lines, indicating that these stars are undergoing He-core burning. In addition, Figure 7 demonstrates that both the He-rich and He-weak sequences show an increasing $\log(n\mathrm{He}/n\mathrm{H})$ with T_{eff}. However, He-rich (e.g., $\log(n\mathrm{He}/n\mathrm{H}) > -1$) hot subdwarf stars are under-represented in our samples. In Figure 8, no distinct sequences are presented, our samples show a wide distribution of $\log g$ similar to that of known hot subdwarf stars. Overall, the atmospheric parameter distribution of these newly identified hot subdwarf stars exhibits a high degree of consistency with that of known hot subdwarf stars from Culpan et al. (2022), demonstrating a high level of reliability of our samples.

6. Conclusion

The LAMOST DR12-V1, released in 2025 March, is currently the world's largest spectroscopic database, holding great potential for the search of hot subdwarf stars. In this study, we propose HsdNet, an advanced two-stage deep learning model incorporating multiple techniques, comprising a binary classification model in stage one and a five-class classification model in stage two. HsdNet not only achieves high classification precision with 94.33% and 94.00% in the binary and five-class classification stages, respectively, but also quantifies the predicted uncertainty, enhancing the interpretability of the classification results through visualizing the model's key focus regions.

We applied HsdNet to the 601,217 spectra from the LAMOST DR12-V1 database to conduct a two-stage search for hot subdwarf candidates. In stage one, we initially identified candidates using the binary classification model. In stage two, the five-class classification model was used to further refine these candidates. Subsequently, through the spectral analysis tool XTGRID, we confirmed 1008 newly identified hot subdwarf stars and provided their atmospheric parameters (e.g., T_{eff}, $\log g$, and $\log(n\mathrm{He}/n\mathrm{H})$). The distribution of their atmospheric parameters is consistent with that of known hot subdwarf stars. These results highlight the practical effectiveness of HsdNet and will significantly advance research related to hot subdwarf stars.

Acknowledgments

We thank the anonymous referee for insightful comments that helped us improve the manuscript substantially. This study received funding from the Natural Science Foundation of Shandong Province with grant Nos. ZR2024MA063, ZR2022MA076 and ZR2022MA089. Additional support was provided by the science research grants from the China Manned Space Project with grant Nos. CMS-CSST-2021-B05 and CMS-CSST-2021-A08. The research was also supported by the National Natural Science Foundation of China (NSFC) under grant Nos. 11873037 and 11803016. Furthermore, the Young Scholars Program of Shandong University, Weihai (2016WHWLJH09) provided support for this research.

We also thank the LAMOST team for providing 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

Yuhang Zhang https://orcid.org/0009-0009-7973-4077
Jiangchuan Zhang https://orcid.org/0009-0006-5523-3997
Yude Bu https://orcid.org/0000-0002-9474-4734
Zhenxin Lei https://orcid.org/0000-0003-2362-6607
Jingzhen Sun https://orcid.org/0009-0000-8724-497X
Zhenping Yi https://orcid.org/0000-0001-8590-4110
Xiaoming Kong https://orcid.org/0000-0002-4764-4749
Meng Liu https://orcid.org/0000-0003-2442-2841

References

Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. 2015, in Proc. of the 32nd Int. Conf. on Machine Learning, 37, ed. F. Bach & D. Blei (JMLR), 1613, https://proceedings.mlr.press/v37/blundell15.html
Bu, Y., Zeng, J., Lei, Z., et al. 2019, ApJ, 886, 128
Cai, B., Kong, X., Shi, J., et al. 2023, AJ, 165, 52
Cheng, Z., Kong, X., Wu, T., et al. 2024, ApJS, 274, 2
Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. 2014, in NIPS 2014 Workshop on Deep Learning arXiv:1412.3555
Copperwheat, C. M., Morales-Rueda, L., Marsh, T. R., Maxted, P. F. L., & Heber, U. 2011, MNRAS, 415, 1381
Cui, X.-Q., Zhao, Y.-H., Chu, Y.-Q., et al. 2012, RAA, 12, 1197
Culpan, R., Geier, S., Reindl, N., et al. 2022, A&A, 662, A40
Der Kiureghian, A., & Ditlevsen, O. 2009, Structural Safety, 31, 105
Dey, R., & Salem, F. M. 2017, in Proc. of the IEEE 60th Int. Midwest Symposium on Circuits and Systems, ed. S. Sonkusale & A. Shrivastava (Piscataway, NJ: IEEE), 1597
Ding, B., Qian, H., & Zhou, J. 2018, in Proc. of the Chinese Control And Decision Conf. (Piscataway, NJ: IEEE), 1836
Dorman, B., Rood, R., & O'Connell, R. 1993, ApJ, 419, 596
Edelmann, H., Heber, U., Hagen, H.-J., et al. 2003, A&A, 400, 939
Gaia Collaboration, Babusiaux, C., van Leeuwen, F., et al. 2018, A&A, 616, A10
Gal, Y., & Ghahramani, Z. 2016, in Proc. of the 33rd Int. Conf. on Machine Learning, 48, ed. M. F. Balcan & K. Q. Weinberger (JMLR), 1050
Ge, H., Tout, C., Webbink, R., et al. 2024, ApJ, 961, 202
Geier, S., Fürst, F., Ziegerer, E., et al. 2015, Sci, 347, 1126
Geier, S., Heber, U., Irrgang, A., et al. 2024, A&A, 690, A368
Geier, S., Hirsch, H., Tillich, A., et al. 2011, A&A, 530, A28
Geier, S., Kupfer, T., Heber, U., et al. 2015, A&A, 577, A26
Geier, S., Østensen, R., Németh, P., et al. 2017, A&A, 600, A50
Götberg, Y., de Mink, S., Groh, J., et al. 2018, A&A, 615, A78
Han, Z., Podsiadlowski, P., & Lynas-Gray, A. 2007, MNRAS, 380, 1098
Han, Z., Podsiadlowski, P., Maxted, P., et al. 2002, MNRAS, 336, 449
Han, Z., Podsiadlowski, P., Maxted, P., et al. 2003, MNRAS, 341, 669
He, K., Zhang, X., Ren, S., & Sun, J. 2016, in Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (Piscataway, NJ: IEEE), 770
He, Y., Cao, Z., Deng, H., et al. 2024, ApJS, 274, 6
Heber, U. 1986, A&A, 155, 33
Heber, U. 2009, ARA&A, 47, 211

Heber, U. 2016, PASP, 128, 082001
Heber, U. 2024, arXiv:2410.11663
Helmy, I., Ismail, M., & Eid, D. 2025, ExA, 59, 2
Hu, J., Shen, L., & Sun, G. 2018, in Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (Piscataway, NJ: IEEE), 7132
Hubeny, I., & Lanz, T. 2017, arXiv:1706.01859
Kepler, S. O., Pelisoli, I., Koester, D., et al. 2015, MNRAS, 446, 4078
Kepler, S. O., Pelisoli, I., Koester, D., et al. 2016, MNRAS, 455, 3413
Kingma, D. 2014, arXiv:1412.6980
Kingma, D., Salimans, T., & Welling, M. 2015, in Proc. of the 29th Int. Conf. on Neural Information Processing Systems, 2, ed. C. Cortes et al. (Cambridge, MA: MIT Press), 2575
Kong, X., & Luo, A. 2021, RNAAS, 5, 249
Kupfer, T., Korol, V., Shah, S., et al. 2018, MNRAS, 480, 302
Kwon, Y., Won, J., Kim, B., et al. 2020, Computational Statistics & Data Analysis, 142, 106818
Lanz, T., & Hubeny, I. 2007, ApJS, 169, 83
Latour, M., Randall, S. K., Calamida, A., et al. 2018, A&A, 618, A15
Lei, Z., He, R., Németh, P., et al. 2023, ApJ, 942, 109
Lei, Z., Zhao, J., Németh, P., et al. 2018, ApJ, 868, 70
Lei, Z., Zhao, J., Németh, P., et al. 2019, ApJ, 881, 135
Lei, Z., Zhao, J., Németh, P., et al. 2020, ApJ, 889, 117
Li, J., Onken, C., Wolf, C., et al. 2022, MNRAS, 515, 3370
Lin, J., Wu, C., Xiong, H., et al. 2024, NatAs, 8, 491
Lin, T., Goyal, P., Girshick, R., et al. 2017, in Proc. of the IEEE Int. Conf. on Computer Vision, ed. E. Mortensen (Piscataway, NJ: IEEE), 2980
Luo, A., & Zhao, Y. 2001, ChJAA, 1, 563
Luo, Y., Németh, P., Deng, L., et al. 2019, ApJ, 881, 7
Luo, Y., Németh, P., Liu, C., et al. 2016, ApJ, 818, 202
Luo, Y., Németh, P., Wang, K., et al. 2021, ApJS, 256, 28
Maeda, S. 2014, arXiv:1412.7003
Meng, X.-C., & Luo, Y.-P. 2021, MNRAS, 507, 4603
Michaud, G., Richer, J., & Richard, O. 2008, ApJ, 675, 1223
Michaud, G., Richer, J., & Richard, O. 2011, A&A, 529, A60
Miller Bertolami, M., Althaus, L., Unglaub, K., et al. 2008, A&A, 491, 253
Moehler, S., Richtler, T., de Boer, K., et al. 1990, A&AS, 86, 53
Németh, P., Kawka, A., & Vennes, S. 2012, MNRAS, 427, 2180
Németh, P., Ziegerer, E., Irrgang, A., et al. 2016, ApJL, 821, L13
Paczyński, B. 1971, AcA, 21, 1
Pelisoli, I., Neunteufel, P., Geier, S., et al. 2021, NatAs, 5, 1052
Selvaraju, R., Cogswell, M., Das, A., et al. 2017, in Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (Piscataway, NJ: IEEE), 618
Sweigart, A. 1997, ApJL, 474, L23
Tan, L., Liu, Z., Wang, F., et al. 2023, ApJS, 268, 28
Tan, L., Mei, Y., Liu, Z., et al. 2022, ApJS, 259, 5
Tailo, M., D'Antona, F., Vesperini, E., et al. 2015, Natur, 523, 318
Van Winckel, H. 2003, ARA&A, 41, 391
Vennes, S., Kawka, A., & Németh, P. 2011, MNRAS, 410, 2095
Vickers, J., Li, Z., Smith, M., et al. 2021, ApJ, 912, 32
Wang, B., & Han, Z. 2010, A&A, 515, A88
Xu, L., Yan, Y., Yu, X., et al. 2019, RAA, 19, 135
Yan, H., Li, H., Wang, S., et al. 2022, Innov, 3, 100224
Yang, M., Yuan, H., Bai, Z., et al. 2025, A&A, 693, A322
Zhang, X., Hall, P., Jeffery, C., et al. 2017, ApJ, 835, 242
Zhang, X., & Jeffery, C. 2012, MNRAS, 419, 452
Zhao, G., & Zhao, Y. 2012, RAA, 12, 723