A new automated tool for the spectral classification of OB stars[★]

E. Kyritsis^{1,2} , G. Maravelias^{3,2} , A. Zezas^{1,2} , P. Bonfini^{4,2} , K. Kovlakas^{1,2} , and P. Reig^{2,1}

¹ Physics Department, and Institute of Theoretical and Computational Physics, University of Crete, 71003 Heraklion, Greece
e-mail: ekyritsis@physics.uoc.gr

² Institute of Astrophysics, Foundation for Research and Technology-Hellas, 71110 Heraklion, Greece

³ Institute for Astronomy, Astrophysics, Space Applications and Remote Sensing, National Observatory of Athens, 15236 Penteli, Athens, Greece

⁴ Computer Science Department, University of Crete, 71003 Heraklion, Greece

Received 23 December 2020 / Accepted 13 September 2021

ABSTRACT

Context. As an increasing number of spectroscopic surveys become available, an automated approach to spectral classification becomes necessary. Due to the significance of the massive stars, it is of great importance to identify the phenomenological parameters of these stars (e.g., the spectral type), which can be used as proxies to their physical parameters (e.g., mass and temperature).

Aims. In this work, we aim to use the random forest (RF) algorithm to develop a tool for the automated spectral classification of OB-type stars according to their sub-types.

Methods. We used the regular RF algorithm, the probabilistic RF, which is an extension of RF that incorporates uncertainties, and we introduced the KDE – RF method which is a combination of the kernel-density estimation and the RF algorithm. We trained the algorithms on the equivalent width (EW) of characteristic absorption lines measured in high-quality spectra (signal-to-noise (S/N) ≥ 50) from large Galactic (LAMOST, GOSSS) and extragalactic surveys (2dF, VFTS) with available spectral types and luminosity classes. By following an adaptive binning approach, we grouped the labels of these data in 11 spectral classes within the O2-B9 range. We examined which of the characteristic spectral lines (features) are more important for the classification based on a number of feature selection methods, and we searched for the optimal hyperparameters of the classifiers to achieve the best performance.

Results. From the feature-screening process, we find that the full set of 17 spectral lines is needed to reach the maximum performance per spectral class. We find that the overall accuracy score is $\sim 70\%$, with similar results across all approaches. We apply our model in other observational data sets providing examples of the potential application of our classifier to real science cases. We find that it performs well for both single massive stars and for the companion massive stars in Be X-ray binaries, especially for data of similar quality to the training sample. In addition, we propose a reduced ten-features scheme that can be applied to large data sets with lower $S/N \sim 20\text{--}50$.

Conclusions. The similarity in the performances of our models indicates the robustness and the reliability of the RF algorithm when it is used for the spectral classification of early-type stars. The score of $\sim 70\%$ is high if we consider (a) the complexity of such multiclass classification problems (i.e., 11 classes), (b) the intrinsic scatter of the EW distributions within the examined spectral classes, and (c) the diversity of the training set since we use data obtained from different surveys with different observing strategies. In addition, the approach presented in this work is applicable to products from different surveys in terms of quality (e.g., different resolution) and different formats (e.g., absolute or normalized flux), while our classifier is agnostic to the luminosity class of a star, and, as much as possible, it is metallicity independent.

Key words. stars: early-type – stars: massive – X-rays: binaries – methods: statistical – stars: emission-line, Be

1. Introduction

Massive stars, although rare, can be characterized as cosmic engines (Maeder et al. 2008) since they play a leading role in the evolution of their host galaxies and the Universe in general. During their short but very interesting lives, as well as their extremely violent deaths, these stars are excellent laboratories for studying of a wide range of astrophysical phenomena relevant to many different aspects of the study of the Cosmos.

Massive stars spend the longest part of their life in the main sequence as OB-spectral-type stars, mainly emitting UV photons and forming H II regions. Hence, they can be used as valuable tracers of star formation (e.g., Steidel et al. 1996; Preibisch & Zinnecker 2007). Due to their strong stellar winds,

they are useful probes of present-day element abundances in the Galaxy (e.g., Gies & Lambert 1992; Przybilla et al. 2013), but also in external galaxies beyond the Local Group (e.g., Urbaneja et al. 2005; Castro et al. 2012). Because of their high luminosities, massive stars are ideal stellar objects for the determination of extragalactic distances (e.g., Kudritzki & Urbaneja 2012). Through their violent deaths, massive stars are the progenitors of objects such as supernovae, black holes, and neutron stars (e.g., Poelarends et al. 2008; Smith 2014). Furthermore, they can be the origin of phenomena such as gamma-ray bursts (GRBs) or the recently discovered gravitational wave sources produced by a merger of two compact objects (e.g., Langer 2012; Abbott et al. 2017). Although, the evolution of massive stars is still unclear, besides their initial mass there are also other significant factors that can determine their fate such as mass loss (Massey 2003; Smartt 2009; Vink & Gräfenor 2012), rotation (Müller & Vink 2014; Lee et al. 2016), metallicity (Bouret et al. 2013), and binarity (Sana et al. 2013, 2017).

[★] Full Table 6 is only available at the CDS via anonymous ftp to cdsarc.u-strasbg.fr (130.79.128.5) or via <http://cdsarc.u-strasbg.fr/viz-bin/cat/J/A+A/657/A62>

Traditionally, spectroscopy is the basic tool for studying the nature of massive stars, because their spectrum reflects their physical parameters (e.g., temperature, gravity, and chemical abundances). Although the detailed analysis of a stellar spectrum using fitting techniques provides the most complete array of information, this approach is generally computationally intensive and it requires high-quality data. On the other hand, the knowledge of a phenomenological parameter, such as the spectral type, is very useful for statistical studies of large samples of massive stars. In particular, the spectral type and the luminosity class of a star are correlated with its physical parameters (e.g., temperature, mass, and radius). Based on this information, we can determine ages and formation scenarios for these systems.

Spectral classification is traditionally performed through the visual examination of a spectrum (Morgan et al. 1943). However, since the entire process is based on the presence or absence of diagnostic spectral lines, it is qualitative in nature and suffers from subjectivity, despite the power of the human eye as a pattern recognition classifier. Furthermore, it is an extremely time-consuming method, unable to handle a large volume of data. Nowadays, the continuously expanding volume of astronomical data sets, both in size and complexity, have commenced the era of big data in astronomy (Pesenson et al. 2010). In particular, big spectroscopic surveys for OB stars such as IACOB (Simón-Díaz et al. 2011a), NoMaDS (Pellerin et al. 2012), and GOSSS (Maíz Apellániz et al. 2011), have started producing a large volume of data, which is obviously extremely difficult to classify using traditional spectral classification. As a result, the best way to address these problems is the development of automated methods based on quantitative measurements of spectral features.

Significant work has been done in this direction, based either on the assessment of specific criteria (e.g., spectral features) or on pattern recognition (e.g., von Hippel et al. 1994). The former approach imitates what a human classifier actually does when visually examining a spectrum and estimating the presence of spectral lines. The latter is based on searching for the spectrum that best resembles the observed one in a library of template spectra for different spectral types (Duan et al. 2009; Scibelli et al. 2014). Furthermore, there are works that combine photometric properties with measurements of spectral features (e.g., Allende Prieto et al. 2006; Gkouvelis et al. 2016). Even though these methods improve the automated spectral classification, they face some challenges. Criteria-evaluation techniques have difficulty accounting for all spectral features within a wide range of spectral types and their sub-types. On the other hand, template matching methods need observed data of similar quality with the templates in terms of wavelength range, resolution, and signal-to-noise (S/N). However, this is not typical for data obtained with different instruments and/or observational strategies. Furthermore, they are computationally intensive, limiting their use in large data sets.

As a result, the next choice for the automation of spectral classification is the use of machine-learning algorithms, which are divided into two groups: supervised and unsupervised. The former use known and labeled data as inputs, while the latter try to learn the various groups of the data from the data itself (for a full overview, see Baron 2019). Recently, thanks to improvements on software development and computational power, an increasing number of studies apply machine learning to a variety of problems in astronomy, taking advantage of its computational speed, as well as, its ability to handle large volumes of data (e.g., Mahabal et al. 2008; Laurino et al. 2011; Castro et al. 2018; Pearson et al. 2019; Clarke et al. 2020; Arnason et al. 2020).

More specifically, a number of previous studies tried to develop classifiers for the spectral classification of stars. Navarro et al. (2012) proposed a classifier for spectra with low S/N based on artificial neural networks, while Mahdi (2008) used probabilistic neural networks for the development of their classifier. In addition, Sharma et al. (2020) identified convolutional neural networks as the best-performing method among two supervised algorithms and one unsupervised one. Also, Kheirdastan & Bazarghan (2016) compared probabilistic neural networks and a K-means cluster, as well as a support vector machine to show that neural networks perform better than the others. Finally, Li et al. (2019) used a random forest algorithm and investigated the spectral feature evaluation problem. Although the aforementioned studies propose automated spectral classification methods, all of them present a more coarse classification from O-type to M-type spectral range without focusing on the OB stars. However, a finer spectral classification of these stars is very important since the difference on the spectral lines is significant among sub-spectral types because the temperature decreases drastically.

In addition, some of these methods rely on Balmer lines. This is a problem for Oe/Be stars, which account for up to 10–15% of B-type stars (McSwain & Gies 2005; Martayan et al. 2007; Maravelias et al. 2017). These stars are particularly important since they are considered to occupy the upper-end of stellar rotation velocity distribution with velocities close to their breakup limit. The reason we cannot use Balmer lines for the classification of these stars is that the strength and profile of the lines varies due to the changing environments of these stars. In addition, Be stars are the donor stars in Be X-Ray Binaries (BeXBs), the main subcategory of high-mass X-ray binaries (HMXBs; Reig 2011). The assignment of a spectral type on the optical companions of these systems is very important. Based on the spectral-type–mass distributions of HMXBs, we can investigate differences between different populations (e.g., HMXBs in the Galaxy and the Magellanic Clouds) and understand how the star-formation history and/or metallicity can influence their formation rates (Linden et al. 2010; Tzanavaris et al. 2013; Antoniou & Zezas 2016; Antoniou et al. 2010). As a result, there is a need for the development of alternative classification schemes that do not rely on the Balmer series. In order to overcome the limitation of using generic spectral classification tools in early-type stars, and particularly Oe/Be stars (which require special treatment), we developed a spectral classifier tailored to these very interesting stars.

This paper is organized as follows. In Sect. 2, we describe the training sample and processing of the data. In Sect. 3, we present an overview of the algorithms used in this work as well as their implementation and their optimization. In Sect. 4, we present the performance of the algorithms, and in Sect. 5 we discuss these results and apply our model to a number of science cases. Finally, in Sect. 6 we summarize our conclusions and discuss possible future improvements.

2. Data collection and processing

2.1. Spectroscopic data

In this work, we used a collection of different spectroscopic surveys targeting the Galaxy and the Magellanic Clouds. We include the Large Magellanic Cloud (LMC; 1/2.5 of solar metallicity) and the Small Magellanic Cloud (SMC; 1/5 of solar metallicity; Russell & Dopita 1992) in order to develop a model that can be applied to a wide range of metallicity

environments. While each survey has a different wavelength coverage, their common range in the optical regime ($\sim 3900\text{--}4900\text{ \AA}$) includes the most characteristic spectral lines of OB stars (Walborn & Fitzpatrick 1990). The reported spectral types in these surveys were obtained through different approaches (either visual inspection or via template fitting). For our purposes, we treated these spectral type labels as the ground truth. Furthermore, aiming as much as possible for a Luminosity Class (LC) independent model that can be applied for the analysis of survey data without any pre-processing, we took into account all the available LCs I-V per spectral type from each survey.

Our Galactic sample consists of spectra obtained from the Galactic O-Star Catalog (GOSC¹ Maíz Apellániz et al. 2013), which mainly includes stars from the Galactic O-Star Spectroscopic Survey (GOSSS Maíz Apellániz et al. 2011). GOSSS is a massive spectroscopic survey of Galactic O-type stars based on high $S/N \sim 250$, resolving power ($R \sim 2500$), and blue-violet observations from both hemispheres. The observations for this survey were obtained from three different telescopes: the 1.5 m at Observatorio de Sierra Nevada (OSN), the 2.5 m duPont telescope at Las Campanas Observatory, and the 3.5 m telescope at Calar Alto (CAHA). They have a spectral coverage of $\sim 3900\text{--}5100\text{ \AA}$.

Also, we obtained spectra from the catalog of OB stars from the Large Sky Area Multi-Object Fiber Spectroscopic Telescope (LAMOST²; Liu et al. 2019). LAMOST is a reflecting four-meter Schmidt telescope, and its unique design has allowed it so far (DR5), to observe almost eight million Galactic stellar spectra with spectral coverage of $\sim 3700\text{--}9000\text{ \AA}$, S/N between 20–300, and a resolving power of $R \sim 1800$ (Cui et al. 2012; Zhao et al. 2012).

Our extragalactic sample consists of spectra obtained from the 2dF survey of the Small Magellanic Cloud, which is an extensive spectroscopic survey of O-, B-, and A-type stars in the $\sim 3900\text{--}4800\text{ \AA}$ wavelength range, with a resolving power of $R \sim 1500$ and S/N of $\sim 20\text{--}150$ (Evans et al. 2004). The observations for this survey were obtained using the multiobject mode of the 2dF spectrograph that was mounted at the top end of the Anglo-Australian Telescope (AAT).

We also obtained spectra from the VLT-FLAMES Tarantula Survey (VFTS), which was a Large Programme at the European Southern Observatory (ESO). The VFTS obtained multi-epoch optical spectroscopy of over 800 OB-type stars in the 30 Doradus region of the Large Magellanic Cloud (LMC). An overview of the observations is given by Evans et al. (2011). The survey primarily used the Medusa-Giraffe mode of the Fibre Large Array Multi-Element Spectrograph (FLAMES) on the Very Large Telescope (VLT). Three of the standard Giraffe settings were used, giving coverage of $3960\text{--}5071\text{ \AA}$ at $R \sim 7500$ (LR02, LR03 settings) and $6442\text{--}6817\text{ \AA}$ at $R = 16\,000$ (HR15N setting). The spectra used here are those from the study by Evans et al. (2015) that provided classifications and radial velocities of the B-type stars from the VFTS. In their study, the LR02 and LR03 spectra were co-added for the apparently single stars (and the single-lined binaries, see Sect. 2 of their paper for further details) and then smoothed and rebinned to an effective resolving power of $R \sim 4000$ for classification; it is these classification B-type spectra that we used here.

For the construction of a uniform data set for the training and the evaluation of the algorithms, we handled each survey

separately. For the GOSC data, we collected all the publicly available spectra in the wavelength range we are interested in ($\sim 3900\text{--}4900\text{ \AA}$), their spectral types, and their LCs. For a number of spectra, the LC information was not available, thus we decided to reject 17 out of 584 objects. For the 2dF data, we followed the same procedure as for GOSC resulting in the collection of 700 spectra. For the VFTS, we rejected seven spectra without known LCs out of 431 objects. In addition, we rejected 51 objects for which their spectral types had an uncertainty larger than 1 spectral type (i.e., B0.5-B2, B1-B3, B0.7-B1.5, B1-B2, and B2-B3) and one object with spectral type A0, as well as two objects without exact spectral types (classified as early B and mid-late B).

Finally, the entire LAMOST catalog of OB stars consists of 22901 spectra of 16.032 stars in the O9-F8 spectral-range and I-V luminosity class. Initially, for stars with multiple exposures we kept only the ones with the highest g -band S/N . Then, we applied a series of quality cuts (g -band $S/N > 50$, removed objects with bad quality flags, missing fluxes). The remaining sample includes 5185 stars with spectral types OB, their subtypes and their LCs within the range I-V. On top of that, we applied an additional cut in order to reduce the contamination by A-type stars misclassified as B9 or B9.5, which is a caveat stated in the original catalog. In particular, we cross-correlated all B9-B9.5 stars with templates of B9 and A0 stars built by manually choosing two high-quality LAMOST spectra of each category. The cross-correlated spectra were first normalized by fitting a high-order polynomial to their continuum, and the template that resulted in the lowest χ^2 was considered as the most appropriate. This way, we identified 292 A0 stars within the B9-B9.5 stars. After removing them from the LAMOST sample, we continued with the remaining 4893 spectra.

In order to construct a homogeneous sample in terms of S/N , we measured the S/N of each spectrum in exactly the same way. For this reason, by visually inspecting a large fraction of randomly selected spectra from all available surveys, we defined a continuum region within the $4220\text{--}4280\text{ \AA}$ wavelength range free of strong lines. Afterwards, we calculated the S/N for each spectrum by using the following formula:

$$S/N = \frac{\bar{F}_{4220-4280}}{\sigma_{F_{4220-4280}}}, \quad (1)$$

where the $\bar{F}_{4220-4280}$ is the mean flux of the spectrum within the selected spectral region and the $\sigma_{F_{4220-4280}}$ is the standard deviation of the flux in the same wavelength range. The S/N distribution per survey is presented in Fig. 1.

The mode for the distribution of the S/N for the overall sample ('All') is ~ 42 and the 95th percentile is ~ 159 . Thus, we defined a S/N cut at 50 and we excluded spectra with $S/N < 50$ from our analysis. Despite it being higher than the mode, we chose this value as the minimum S/N of an input spectrum for two reasons: (a) after visual inspection of a large number of the available spectra, we found that, in general, spectra with $S/N \sim 40$ are not good enough to measure weak spectral lines; (b) by choosing a S/N cut much higher than 50, we lose a large fraction of the available data for the training of our classifier, as is shown clearly in Fig. 1. As a result, we opted for this specific S/N cut as a compromise between good quality spectra and adequate number of data.

After the previous S/N filtering, our final sample contains 5329 objects with the following demographics:

1. 549 galactic sources from the GOSC in a spectral type range between the O2-O9.7 types and LC I-V.

¹ <http://ssg.iaa.es/>

² <http://www.lamost.org>

Fig. 1. S/N distribution per survey as it is measured in a clear continuum region within the wavelength range 4220–4280 Å. The statistical mode for “All” distribution is ~ 42 , and the 95th percentile is ~ 159 . Based on this result, we continued our analysis by only taking into account spectra with $S/N > 50$.

- 3822 galactic sources from the LAMOST catalog in a spectral range between the O9-B9.5 types and LC I–V.
- 590 SMC sources from the 2dF survey in a spectral range between the O5-B9 types and LC I–V.
- 368 LMC sources from the VFTS in a spectral range between the B0-B9 types and LC I–V.

Since our final sample encompasses spectra from both Galactic and extragalactic sources, we de-redshifted all of them to their rest-frame wavelength. The two extragalactic surveys, VFTS and 2dF, were corrected assuming the line-of-sight velocity of the LMC and the SMC from the SIMBAD database (Wenger et al. 2000). The LAMOST spectra are already redshift-corrected, so we did not apply any further correction. Finally, for the GOSC spectra after visual inspection, we found that they are corrected for velocity shifts.

2.2. Classification scheme

To build an appropriate spectral type classification scheme for our algorithm, we relied on the classification criteria for B-type stars in the SMC, as they were defined in the previous works of Evans et al. (2004) and Maravelias et al. (2014, see Table 2). According to this scheme, we constructed a sample of 17 characteristic spectral lines by including both He I lines, He II lines (strong indicators for early type O-stars), and metal lines such as Mg II, Si I, Si II, etc. (indicators of late type B-stars). In general, despite the fact that Balmer lines are widely used for the spectral classification of OB stars, we intentionally did not consider them in our analysis. Oe and Be stars and those which are found as companions in BeXBs create a circumstellar disk that can be detected observationally from the presence of strong Balmer emission lines. In particular, the variability of the disk’s size and geometry (in timescales of months) results in a variable Balmer series (Porter & Rivinius 2003; Reig 2011). When the circumstellar disk is absent the Balmer lines are in absorption but when the circumstellar disk is fully developed the Balmer lines are in strong emission. As a result, we cannot take advantage of these lines since their presence does not depend only on the star’s tem-

Fig. 2. *Top panel:* initial spectral types distribution. The number of types is high enough, with some of them having a low number of objects. This distribution is not effective for training our algorithms, especially for the underrepresented classes. *Bottom panel:* final labels after the application of the adaptive binning. The number of spectral classes is decreased from 32 to 11.

perature but also on the disk’s size. Finally, although Table 2 in Maravelias et al. (2014) includes the Ca IIK/3928 spectral line, we did not use it because the wavelength range of our surveys did not cover a large enough bluewards region for reliable continuum subtraction. However, this does not significantly affect the classification scheme because Ca IIK/3928 is not a strong discriminator between O and B spectral types in comparison with other spectral lines included in our classification scheme.

2.3. Spectral type binning

In the top panel of Fig. 2, we present the initial spectral type distribution of the stars in our sample. For a number of spectral types, our sample only included a few objects (e.g., O2, O3, O3.5 or O9.2, and B0.2). These underrepresented types had fewer than 50 objects each, meaning that the machine learning algorithm cannot be trained efficiently. On the other hand, it is very hard to distinguish between adjacent spectral types of mid (B3–B5) or late (B5–B8) B-type stars (see e.g., Table 2 of Maravelias et al. 2014). Although in this case we have enough objects per class for the training of the algorithm, the available spectral lines cannot discriminate between neighboring spectral types (Evans et al. 2004; Maravelias et al. 2014; Gray & Corbally 2009). To overcome these limitations, we defined our classification classes following an adaptive binning scheme based on the resolution of

Table 1. Adaptively binned spectral classes considered in our analysis, and the original spectral types.

Original spectral types	Grouped spectral classes
O2-O6.5	O2-O6
O7-O7.5	O7
O8-O8.5	O8
O9-O9.7	O9
B0-B0.7	B0
B1-B1.5	B1
B2-B2.5	B2
B3-B4	B3-B4
B5-B7	B5-B7
B8	B8
B9-B9.5	B9

our classification scheme, while ensuring that there are at least 100 in each bin. The final scheme followed in our analysis is presented in Table 1. This way, we reduced the number of spectral types examined from 32 to 11 without losing the physical meaning of the grouping (e.g., grouping all earliest types together) and we increased the number of objects even for the classes with a very small number of stars in our sample.

In the bottom panel of Fig. 2, we present the final spectral class distribution of our sample as it was formed after the adaptive binning.

2.4. EW and uncertainty measurements

To quantify the intensity of the selected features, we measured the EW of each spectral line from the one-dimensional spectrum of each object in our data set. The EW is defined as the width of the continuum region of a spectrum that contains the same flux as the spectral line examined, and it is given by

$$EW = - \int_{\lambda_1}^{\lambda_2} \frac{F_{\text{cont}}(\lambda) - F_{\text{line}}(\lambda)}{F_{\text{cont}}(\lambda)} d\lambda = (\lambda_1 - \lambda_2) - \int_{\lambda_1}^{\lambda_2} \frac{F_{\text{line}}(\lambda)}{F_{\text{cont}}(\lambda)} d\lambda, \quad (2)$$

where λ_1 and λ_2 are the initial and final wavelength over which the line flux is calculated, and F_{cont} and F_{line} are the continuum and spectral-line flux density, respectively.

To measure the EW, we used the *Sherpa* fitting package v 4.13.0 (Freeman et al. 2001) to model the characteristic spectral lines and their local continuum regions by fitting a Gaussian line and a polynomial, respectively. *Sherpa* is a very useful tool that allows us to measure the flux and the amplitude either of absorption or emission spectral lines using complex models, and to determine their parameters and their corresponding uncertainties while accounting for uncertainties on the data.

First, we visually inspected 100 randomly selected spectra of different spectral types and LC from each survey to define the regions used for the spectral fit, taking care to excise strong spectral features from the continuum bands. We decided to use slightly different continuum bands for the VFTS, LAMOST, and GOSC surveys to avoid spectral features at the edge of the continuum bands, which can be particularly strong in the higher resolution spectra. For the 2dF survey, we used the same bands as those used for the LAMOST survey. In Table 2, we present the full scheme of the selected spectral lines with their central wavelengths and their corresponding spectral ranges, as well as the continuum regions that were used in the fitting process.

For the VFTS, 2dF, and GOSC surveys in which the spectra are normalized and they did not provide uncertainties, we adopted the standard deviation of the data in a relatively clear continuum region in the 4220–4280 Å wavelength range as a measurement error. The continuum of these spectra was fit with a constant. On the other hand, for the unnormalized LAMOST spectra, which also provided measurement errors, we modeled the continuum with a first-order polynomial, and we used the values provided by the survey as flux errors. We fit each spectral line separately except for the Si IV/4088 Å, Si IV/4116 Å, He I/4121 Å, and Si IV/4130 Å spectral lines. Since these lines are very close to H δ line, resulting in strong blending particularly for late-type B-stars, they were fitted in a single model also including the H δ line. In order to account for the strong wings of the H δ line, we used two Gaussian components in the later type stars, one accounting for the narrow core of the line and one for its broad wing. The position of each line was at its rest-frame wavelength ± 2 Å in order to avoid confusion with other neighboring lines (e.g., the Si III/4553 Å and the Si III/4550 Å lines). The final model for this spectral region consisted of four Gaussian lines, one for each spectral feature of interest, two Gaussians for the H δ line, and a zeroth- or first-order polynomial for the continuum. To estimate of the model parameter uncertainties, we used the *covar* tool, which computes confidence intervals for the specified model parameters in the data set using the covariance matrix of the fit. The entire fitting process was based on an automated pipeline developed for this purpose.

After fitting each spectral feature of interest, we calculated the EW using the *eqwidth* routine provided by *Sherpa*. This tool calculated the distribution of EW for each spectral line by evaluating the line and the continuum flux based on drawings of the model parameters from the fit covariance matrix. We performed the calculation for 1000 draws and we adopted the median of the EW distribution and its 68% percentile as its uncertainty.

Since in many cases some of the spectral lines were not detected, after the completion of the spectral fit we applied a post-processing step to identify non-detections. With regard to non-detections, we characterized the fits that had failed because of unconstrained parameters and all successfully measured spectral lines with an $EW/S/N < 3 \left(\frac{EW}{\sigma_{EW}} < 3 \right)$ or $EW > 0$ (emission lines). Instead of removing these cases from the training of our classifier, we included this information in our analysis by replacing their EW with the same constant (magic number) which had a different value for each line. In Fig. 3, we present the fraction of non-detections per spectral type and per spectral line. This way, we do not reduce the size of our sample by removing a large number of objects with a few problematic lines, while at the same time we do include the diagnostically important information of the non-detection of some lines in our analysis. In addition, by selecting each magic number to be relatively far from the EW distribution of the successfully measured spectral lines, we ensure that the actual detections and non-detections do not overlap.

3. Building the machine-learning models

3.1. Algorithms

In this work, we used the widely used supervised learning RF algorithm (for a thorough review, see Louppe 2014). RF is a well-known ensemble method used for both classification and regression tasks (e.g., Cariles et al. 2010; Vasconcellos et al. 2011).

Table 2. Spectral lines with their wavelength ranges, their central wavelengths and continuum regions that have been used for the calculation of their EWs, during the spectral line fitting.

Line ID	Spectral line			Continuum blue		Continuum red	
	λ_{central} (Å)	λ_{start} (Å)	λ_{end} (Å)	λ_{start} (Å)	λ_{end} (Å)	λ_{start} (Å)	λ_{end} (Å)
He I	4009	4004	4016	3990	4002	4035	4060
He I+He II	4026	4017	4034	3990	4002	4035	4060
Si IV	4088	4084	4091	4055	4075	4150	4170
H δ	4100	4095	4110	4055	4075	4150	4170
Si IV	4116	4113	4118	4055	4075	4150	4170
He I	4121	4118	4125	4035	4060	4150	4170
Si II	4130	4125	4135	4035	4060	4150	4190
He I	4144	4140	4150	4035	4060	4150	4170
He II	4200	4190	4207	4150	4190	4245	4260
Fe II	4233	4229	4237	4205	4225	4238	4260
He I	4387	4382	4392	4365	4380	4398	4415
O II	4416	4412	4421	4398	4411	4440	4460
He I	4471	4462	4477	4440	4460	4495	4535
Mg II	4481	4477	4488	4440	4460	4490	4505
He II	4541	4537	4547	4510	4535	4590	4620
Si III	4553	4548	4558	4485	4507	4600	4620
O II+C III	4645	4635	4655	4600	4625	4660	4685
He II	4686	4679	4692	4660	4670	4730	4745

Notes. We used slightly different continuum bands for VFTS, LAMOST and GOSC surveys because the different spectral resolutions of each of them can cause the presence of undesirable spectral features within the predefined continuum regions, adding noise in the data. Here, we present the continuum regions that selected for LAMOST data since their population dominates among the other surveys. This is the final scheme of features that is used through this work, except for the H δ /4100 Å spectral line, which is used only for the fit.

Fig. 3. Fraction of the non-detections per spectral type and per spectral line.

Ensemble methods combine either different supervised learning algorithms or the information of a single algorithm that was trained on different subsets of the training set. The building blocks of the RF algorithm are the decision trees. A decision tree is a nonparametric model built during the training stage that describes the relation between the features and the labels using a set of continuous nodes in a tree-like structure. The initial node of a tree is called the root node, and the final nodes, in which the predicted label is calculated, are called terminal nodes. The input training data follow a top-to-bottom flow through the tree, which is split into children nodes according to the condition at each node. This condition is defined as the feature and its corresponding value that maximizes the class separation between the children nodes. This ‘best’ splitting threshold can be computed by using different metrics such as the Gini impurity, entropy, or information gain, with the most common choice being the first one. In a two-class classification problem (can be generalized to multiclass problems), the Gini impurity is defined as the probability of a randomly selected object to be misclassified if it is assigned with a label that is randomly drawn from the distribution of the labels in the group (Reis et al. 2019).

When we try to classify an unseen object with a decision tree, the object will be propagated through the tree, and, based on the values of its features and the conditions in the nodes, it will reach a terminal node with a specific label assigned to the object. However, the prediction power of a unique decision tree is limited since it is vulnerable to overfitting and cannot be generalized well to unseen data sets (Breiman 2001). Solving for that, the RF algorithm constructs a number of decision trees, and during the training process of these trees it uses randomly selected subsets of the full data set based on a bootstrap method. Furthermore, during this process, random subsets of the features are used in each node of each decision tree to find the appropriate conditions in the nodes to build each tree (i.e., determine its nodes). The final prediction of the RF algorithm is in the form of a majority vote. Each individual tree in the forest suggests a class for the examined object, and the final prediction is the class that has been proposed from the majority of the trees.

The RF algorithm has numerous advantages in comparison with other supervised algorithms. Firstly, due to the randomness on which it is built, the correlation between different trees is reduced. Hence, the structure and the conditions in the nodes of each tree are very different from tree to tree, resulting in a model that can be generalized well to new unseen data and consequently lead to improved performance. In addition, it can handle either categorical or numerical features with no need to normalize or scale the data, and it is effective in multiclass problems such as ours. On top of that, it can properly handle the non-detections that are fed to it through the magic number since the RF algorithm only needs a value for each feature in order to make a decision at each node. However, the most important benefit of the RF algorithm for our goal is that its tree-based logic resembles the traditional spectral classification technique.

Nevertheless, the RF algorithm is unable to take into account uncertainties of the data during the training process. To overcome this limitation, an alternative probabilistic RF (PRF) approach was proposed by Reis et al. (2019). The input of the PRF algorithm at each node is a quadruplet that contains the label and its corresponding uncertainty, alongside with the value of a feature and its uncertainty. The PRF treats these values as random variables, described by a normal probability distribution. In particular, the features become probability density functions (PDFs) with the mean value being the value of the feature and the variance being the square of the feature uncertainty, whereas the

labels become probability mass functions (PMFs), where each label is assigned to an object with some probability. Since the spectral types provided with our data set do not include uncertainties for all surveys, we used the PRF accounting only for the uncertainties in the features. In the ideal case of negligible uncertainties, the PRF converges to the original RF.

We used the implementation of the RF classifier `sklearn.ensemble.RandomForestClassifier()` provided by the `scikit-learn`³ version 0.23.2 (Pedregosa et al. 2011) package for Python 3. For the PRF we used the publicly available package at GitHub repository⁴. After processing the data described in Sect. 2, we randomly shuffled our complete data set and split it into training and test data sets. We considered 70% of the full set (3730/5329 spectra) for the training data and 30% (1599/5329 spectra) for the testing data. This is a common training/test ratio in machine learning that avoids overfitting (Ksoell et al. 2018).

3.2. Hyperparameter optimization

A crucial step in building a machine learning model is the selection of appropriate values for the hyperparameters. Hyperparameters are those that control the learning process, and their values are set at the initialization of the training process. An inappropriate choice of these values can lead to underfitting or overfitting of the model. The RF algorithm includes a large number of hyperparameters. We chose to optimize our model for the most important ones, which are:

- `n_estimators`: the number of trees in the forest.
- `max_depth`: the maximum number of levels in each decision tree.
- `min_samples_split`: the minimum number of samples required to split an internal node.
- `min_samples_leaf`: the minimum number of samples required to be at a leaf node.
- `max_leaf_nodes`: this hyperparameter affects the way that trees grow in order to have the best result. The best nodes are defined as showing a relative reduction in impurity.
- `max_samples`: the number of samples to consider when the algorithm draws from the training data set to train each base estimator.

To optimize these hyperparameters, we trained and evaluated the RF model over a range of values (with others fixed to their default). After each iteration, we obtained the corresponding cross-validation accuracy score. The k -fold cross-validation (CV) accuracy score or rotation estimator (Kohavi 1995) is a technique used in order to take into account the entire data set and not only a part of it (when splitting into a training and test subsets). The initial data set is split into k smaller data sets (“folds”). For each of the k -folds, a model is trained using $k-1$ of the folds as a training set based on a stratified method (i.e., preserving the fractional representation of all classes in the original sample). Then, the resulting model is evaluated on the remaining data and reports an accuracy value. The final cross-validation score is the average over the folds and its uncertainty is the standard deviation of these values. In our work we, chose to use a $k=5$ for the number of folds. In Fig. 4, we plot the validation curve for each hyperparameter; that is, the CV accuracy score for different values of the accuracy for the hyperparameter, as well as its standard deviation.

³ <https://scikit-learn.org/stable/>

⁴ <https://github.com/ireis/PRF>

Fig. 4. Validation curves of the most important RF hyperparameters, including the standard deviation of each CV score (gray shaded region). Based on these curves we selected the optimal values in the way described in Sect. 3.2. The final values of the hyperparameters are tabulated in Table 3.

The `n_estimators` increases quickly and beyond a threshold it remains almost constant. The same behavior is seen for the `max_depth` and the `max_samples`. In contrast, the `min_samples_leaf` and the `min_samples_split` show higher scores for lower values (close to their default), while gradually their scores decrease for larger values. Finally, the `max_leaf_nodes` seems to be consistent with its default value (None), which means that the higher the number of leaf nodes is, the higher the accuracy of the model. The RF algorithm has two more important hyperparameters: `class_weight` and `max_features`. The former is useful for imbalanced data sets, while the latter controls the maximum number of features to consider when the algorithm is looking for the best split. We set the `class_weight` hyperparameter to “balanced_subsample”, which means that during the training of the algorithm, the weight per class will be inversely proportional to the number of objects belonging to this class. Moreover, `max_features` will be equal to $\sqrt{N_{\text{features}}}$, where N_{features} is the total number of features, and this is commonly used as default value for the RF algorithm.

In the previous test, when a range of values of a hyperparameter is tested all others remain constant. To check if there are any other combinations, we applied a grid search technique, which finds the best set of hyperparameter values trying all possible combinations over a predefined range of values. The model for each combination of values is evaluated based on k-fold CV test. It is a powerful and accurate method for tuning the hyperparameters but it is expensive in terms of computational time. Consequently, to reduce the computational time, we applied the grid search method using the range of values selected by the previous validation test. For `n_estimators` and `max_depth`, we selected the region where the algorithm’s performance becomes constant. For `min_samples_leaf` and `min_samples_split`, we took into account a smaller region close to their default values since the validation curve indicated that the score remains constant at its maximum value over a very broad parameter range. Finally, for the `max_samples` we considered all the val-

ues within 0.1–0.9. Furthermore, we did not consider the hyperparameter `max_leaf_nodes` for further analysis with the grid search method. This is because, from the corresponding validation curve, it is obvious that the default value is the optimal one. The results of the grid search are presented in Table 3.

Since the PRF algorithm is based on the RF algorithm, its most important hyperparameters are equivalent to the regular RF algorithm. These are `n_estimators`, `max_depth`, and `max_features`. In addition, the PRF includes one more hyperparameter called `keep_proba`, which is the threshold of the propagation probability in the nodes, below which the nodes are pruned from the tree. This pruning reduces the computational time of the algorithm, without affecting its performance. By using the validation curves for these hyperparameters again, we saw that they exhibit similar behavior to that in the regular RF. As a result, we did not apply an additional grid search for them, and we considered exactly the same values as those derived from the RF hyperparameters’ optimization. For `keep_proba`, we set the value of 5%, which is an optimal value according to Reis et al. (2019), but it is also consistent with the result of its validation curve. In Table 3, we present the optimal values (last column) as they were derived from the hyperparameters’ optimization for both RF and PRF models. In addition, the range of hyperparameter values and the number of steps that were considered during tuning of both models are tabulated. The optimal values in this table are the ones used to construct our final working models.

3.3. Feature selection

Besides the optimization of the most important hyperparameters, another way to improve a model’s performance is to investigate if there is any specific combination of features that can result in a better score. For this reason we used a sequential feature selection (SFS) algorithm in order to assess if we can reduce the error of the model when it is applied in other samples by removing irrelevant features that can introduce noise.

SFS algorithms are a family of search algorithms that are used to identify the minimal set of features that provide optimal information of a model. They are based on an iterative process, where the model’s performance is evaluated for different feature subsets, the size of which is predefined by the user. For our analysis, we used the sequential forward floating selection (SFFS) (see Pudil et al. 1994) provided by the `MLxtend` (Machine learning extensions; Raschka 2018) library for Python.

Using the above algorithm, we tested the full set of features in Table 2, except for the $H\delta$ spectral line. Due to the similarity of RF and PRF, we ran the SFFS only for the best RF configuration, as determined from the hyperparameters’ optimization. Since a priori we do not know the optimal size of the best feature subset, we added an extra step. We repeatedly executed the SFFS algorithm for all possible sizes, namely within the 1–17 range. Then, we plotted the size of these subsets; that is, the number of the features versus the CV score (based on the accuracy metric) that was suggested from the RF model evaluation, and we present it in Fig. 5.

As shown in Fig. 5, the SFFS algorithm suggests that after ten features the k-fold accuracy of the model remains constant, which means that any other addition of features does not change the model’s performance. Furthermore, for a small number of features, the accuracy is low, which is expected, since the model has not enough information to classify correctly the objects in each class. We see that after ten features the CV score already reaches a value of 0.68, indicating that the classifier performs

Table 3. Optimal hyperparameter values as they were derived after the optimization procedure for both RF and PRF models, as well as the value ranges and steps in which the optimal value was searched for.

RF model					
Hyperparameter	Validation curve range	Step	Grid search range	Step	Optimal value
n_estimators	10–950	50	100–950	50	450
max_depth	1–49	1	10–49	1	15
min_samples_split	2–59	1	2–9	1	5
min_samples_leaf	1–29	1	1–9	1	4
max_leaf_nodes	2–59	1	–	–	‘default’
max_samples	0.1–0.9	0.1	0.1–0.9	0.1	0.6
max_features	–	–	–	–	‘auto’
class_weight	–	–	–	–	‘balanced_subsample’
PRF model					
Hyperparameter	Validation curve range	Step	Optimal value		
n_estimators	10–950	50	450		
max_depth	1–49	1	15		
keep_proba	0–1	0.01	0.05		
max_features	–	–	‘auto’		

Notes. The optimal values define the final models used throughout this work.

Fig. 5. Number of features versus the CV accuracy score of RF algorithm. As it is shown the SFFS algorithm suggests that after ten features the model’s performance does not change significantly. However, we considered the entire set of spectral lines (17) since even the weaker lines can improve the performance of particular spectral classes.

well even with ten features. Although this ten-feature scheme is comprised of the strongest spectral lines (e.g., He I, He II) and the addition of extra features improves the CV score only by a factor $\sim 2\%$, we do not consider it as the main classification scheme. Instead we adopted the entire classification scheme (17 spectral lines). This choice is driven, from an astrophysical perspective, by the fact that for distinguishing individual spectral classes the examination of specific, weaker lines is needed. As a result, the adoption of a classification scheme with only the strongest lines would cause our classifier to underperform in spectral classes that rely on those weaker lines (e.g., the B8.0 class relies on the comparison of the He I lines with the weak Si II/4128+4130 (see Table 2 in Maravelias et al. 2014). However, in Appendix B we further discuss this ten-feature scheme as an alternative approach that can be applied to large data sets with lower S/Ns (~ 20 –50), where the measurement of weaker lines is not reliable.

To demonstrate that we can trust the adopted classification scheme with 17 features and understand the behavior of the SFFS algorithm, in Fig. 6 we plot the EW distribution per spectral line and per spectral class only for the well-measured spectral lines (detections). From our experience with optical spectra, we expect that the SFFS algorithm will consider the spectral lines with well-separated EW distributions between different spectral classes as the most useful features. In contrast, spectral lines with strong overlap between their EW distributions are expected as being less useful during the training of the algorithm. Furthermore, the He II lines, which are present only in hot O stars, are expected to be distributed around values lower than 0, while for later type stars their EW values will be closer to zero; or, when the line is non-detected they will be replaced by the corresponding magic number. On the other hand, the distributions of lines such as the He I/4471 Å and the Mg II/4481 Å will show an anticorrelation with the former decreasing from the early-B to late-B, and the latter increasing from early-B to late-B. Indeed, a comparison between Fig. 6 and the SFFS results (Fig. 5) confirms that the majority of the spectral lines within our classification scheme have well-separated values per spectral class (e.g., He II/4686 Å, He I/4026 Å, O II+C III/4645 Å, or Mg II/4481 Å, etc.), and at the same time these lines are those that the SFFS algorithm considers as important (i.e., before it reaches its plateau). In contrast, the weaker spectral lines such as Fe II/4233 Å and Si IV/4088 and Si IV/4116 Å provide no significant contribution to the algorithm’s performance (incorporated in the model after the 10 features) due to the low separability of their EW values.

As a final test of the credibility of the SFFS results, we compared them with a different method of feature selection that estimates how much each feature contributes during the training process and is included in sklearn’s RF implementation. Specifically, when a tree is trained, the decrease of each feature’s Gini impurity can be computed. Afterwards, this impurity decrease can be averaged across all the trees in the forest, and the features are ranked according to the above measurement. The result of this feature importance quantification is shown in Fig. 7. The most important features are consistent with the spectral lines that have been selected by the SFFS algorithm until

Fig. 6. Distribution of EW values per spectral class and per spectral line for the full set of features as it is presented in Table 2. The best-separated distributions of spectral lines coincide with the result of SFFS analysis (see Sect. 3.3). Based on these distributions, as well as the result of the SFFS analysis, we constructed the final scheme of features that it is used throughout this work (Table 2).

it reaches its plateau where any other addition of features does not significantly improve the performance of the algorithm. Furthermore, in both methods, the Fe II/4233 Å, Si IV/4088 Å, and Si IV/4116 Å lines appear to be insignificant for training the RF model; this result is also supported by their EWs distributions. In addition, the high-ranked spectral lines are in accordance with the SFFS results and consistent with the behavior of their EW distributions per spectral class. Overall, the results of the comparison between these two different techniques are consistent. Therefore, we are confident that the scheme presented in Table 2 is robust.

4. Results

4.1. Evaluating the performance of the best RF and PRF models

After optimizing the algorithms and training the models, we evaluated their performance on the test set. For the evaluation of the algorithms, we used the confusion matrix, which is widely used in problems of statistical classification (e.g., Mahabal et al. 2017; Arnason et al. 2020). The confusion matrix has a table layout that visualizes the performance of a classifier using the test subset. Each row of this table represents a true class, while

Fig. 7. Feature of relative importance of spectral lines used in this work and its standard deviation.

each column represents a predicted class. Therefore, the confusion matrix shows the number of objects in each class versus the number of objects predicted by the model to belong to a particular class. In the ideal case, where the prediction ability of the model is 100% correct, we would expect all objects to be on the diagonal of the matrix. For the interpretation of the confusion matrix and the assessment of the performance of a model, we used basic terms and metrics, which are defined in Table 4 (c.f., Pearson et al. 2019). Accuracy and Recall are the basic metrics that we use in this work since we are interested in how accurate our model is, as well as how correct the predicted spectral types are.

In the top left and top right panels of Fig. 8, we present the confusion matrices of the RF and PRF models, respectively. Furthermore, in Table 5 we present the performance metrics per class for both algorithms, as well as the number of stars that were used as a test sample for each class.

As can be seen from the confusion matrices, both the RF and the PRF models have an overall accuracy of $\sim 70\%$. This means that their performance is similar, although their training method is different in terms of including measurement uncertainties. In addition, the k-fold validation for both models gives a CV score of $69 \pm 2\%$ for RF and $70.0 \pm 0.5\%$ for PRF. To further assess the uncertainty of our classifiers we repeated the k-fold CV test three times with a different random split of the data set in k-folds at each time. Afterwards, we calculated the average accuracy and the standard deviation across all folds and all repeats. This repeated k-fold returns effectively the same scores ($69 \pm 1\%$ and $70 \pm 1\%$, respectively). These scores are high if we consider the complexity of such multi-class classification problems and the fact that for the training we used a diverse sample including data from a variety of surveys with different observational setups, which also suffer from non-negligible measurement errors for the spectral features. In addition, the labels used for the training also include an intrinsic uncertainty. The fact that the majority of misclassified stars belongs to neighboring classes, combined with the low standard deviation of the CV scores, indicates the reliability and the stability of both algorithms.

A careful examination of the metrics for both algorithms (Table 5), reveals that their behavior is similar not only in terms of overall score but also with respect to the scores per class. More specifically, the two models show relatively poor performance in the case of early B-type stars. In contrast, they perform better when they classify spectral types of O and late B-type stars.

Since we are interested in the true positives, we focused on the recall metric. Comparing the values of the same classes in both models we see that the highest scores are for class O2-O6 (0.97 for both RF and PRF) and the lowest scores are for class B8 (0.50 and 0.34, respectively). In addition, high scores are reached for the O7, O8, O9, and B0, as well as for B3-B4 and B9, while for the remaining classes (e.g., B1, B2, B5-B7) the scores are lower, indicating a difficulty for both models to distinguish these spectral classes.

4.2. Development of an independent method for incorporating uncertainties

The regular RF model does not handle measurement uncertainties on the features during the training process, contrary to the PRF. Furthermore, both methods are sensitive to the sampling of the feature distributions for each class. In our case, this is a problem since many classes are significantly underrepresented. While the imbalance between the different classes can be accounted for with the `class_weight='balanced_subsample'` flag, this cannot account for the scarce sampling of the feature parameter space in the poorly represented classes. With this in mind, we devised an alternative approach to the standard data augmentation methods where we incorporate in the sampling process a kernel density estimation (KDE).

A KDE is a nonparametric way to estimate the probability density function (PDF) of a data set without any previous knowledge of its distribution. More specifically, the KDE uses a mixture consisting of a predefined kernel component per data point. The sum of the components for all the data points essentially results in a nonparametric estimator of the density of the overall sample. In our case, we adopted a Gaussian kernel.

Following our previous training/test sample split (70/30 %), we first defined the training sample. Then, based on this sample we created a multidimensional histogram per spectral class (11 in total), which includes the EW for the spectral lines in the considered features (see Table 2).

Essentially, each subset is a different seventeen-dimensional dataset encapsulating the joint distribution of the EW of the characteristic spectral lines based on the measured EW of the objects of each class. Then, we applied a separate Gaussian KDE per subset (i.e., class) and we determined 11 different seventeen-dimensional PDFs, providing the joint distribution of the EW for all features of interest per class. The bandwidth (smoothing parameter) of the applied kernels was chosen to be equal to the medians of the uncertainties of the EW values for all the features in each class. Once we had calculated the PDFs, we drew 1000 sets of EWs per class. This way, we created a new balanced data set of 11 000 mock data representing objects very similar to the real ones, making it the proper training data set for the RF algorithm. This approach has the following advantages: (a) it smooths the sparsely sampled feature distributions of the underrepresented classes; (b) it accounts for the correlated behavior of the different features within each class and between classes (which could increase the efficiency of the model); and (c) it resolves the bias resulting from the significant imbalance of the training data for different classes.

For the training, we used the same optimized RF configuration (see Table 3), since the data distribution on which the model was optimized is the parent distribution of the artificial data set, and consequently we do not expect significant differences in the values of the hyperparameters. As a final step, we tested the performance of our method on the same test set that we used for the

Table 4. Definition of the terms and the metrics that are used in this work for the valuation of the classifiers.

Term	Definition	Formula
True positive (TP)	An object is predicted from the classifier to belong to a spectral class, and it actually belongs to this class.	
True negative (TN)	An object is predicted from the classifier to not belong to a spectral class, and it actually does not belong to this class.	
False positive (FP)	An object is predicted from the classifier to belong to a spectral class, and it actually does not belong to this class.	
False negative (FN)	An object is predicted from the classifier to not belong to a spectral class, and it actually belongs to this class.	
Accuracy	Number of objects that are predicted correctly from the classifier, over the total number of the tested sample	$(TP+TN) / (TP+TN+FP+FN)$
Precision	Number of objects that are predicted correctly from the classifier to belong to a spectral class, over the total number of objects predicted by the classifier to belong to this class.	$TP / (TP+FP)$
Recall	Number of objects that correctly predicted from the classifier to belong to a spectral class, over the total number of objects that actually belong to this class.	$TP / (TP+FN)$
F1-score	The harmonic mean of the Recall and Precision. It is as an overall performance metric for the classifiers	$2TP / (2TP+FP+FN)$

evaluation of the RF and PRF. In the bottom panel of Fig. 8, we present the final confusion matrix of our method, while the performance metrics per class are tabulated in Table 5. This method will be referred to as the KDE-RF.

As we see, the overall accuracy of our method is $\sim 70\%$, and the corresponding k-fold CV score is $69.0 \pm 0.5\%$. In addition, the repeated k-fold gives a score of $68 \pm 1\%$. Furthermore, in terms of the performance of the model for the individual spectral classes, we see again that the classifier better predicts the early-O-type and late-B-type stars in comparison to the early B types. The highest and lowest recall score is for the O2-O6 (0.97) and for B8 (0.60) classes, respectively. In general, our method demonstrates the same performance within the errors as the other methods we discussed in Sect. 4.1, but recall is slightly better in the underperforming classes (e.g., B8). This similarity of the performance for the three different approaches once again indicates the stability and the reliability of the methods.

In order to test to what degree the performance of the model is affected by the correlation between the EW of the spectral lines for each star, we performed the following Monte Carlo analysis. For each star, we drew 100 samples of EW values per line from Gaussian distributions with means and standard deviations equal to those derived from the spectral fit analysis. This way, we preserved the correlation between the strength of the different spectral lines for each star belonging to each spectral class we considered. We then performed the RF analysis using this sample of 373 000 simulated objects, which gave an overall accuracy of 69% and a confusion matrix virtually identical to those presented in Fig. 8. This indicates that the dominant factor determining the behavior of the RF classifier is the intrinsic scatter (spread) of the EWs for each spectral class.

5. Discussion

5.1. Summary of the results

We used three different RF-based methods to assess if this algorithm can be used for the spectral classification of OB-type stars into their subtypes. We find that all the models have similar performance, both in their overall score, and also with respect to

the individual spectral classes. Moreover, the fact that for almost all the classes the misclassified objects belong to neighboring classes indicates the success of our classifier. With such classification problems, the classes are contiguous and ordered based on the continuous nature of the underlying physical parameter, which in the case of spectral types is the temperature. For the direct comparison of the three methods, in Fig. 9 we present how their basic metrics vary per spectral class. Essentially, we plot the values given in Table 5, and as we can see, the precision, recall, and F1-score exhibit almost the same variation with respect to the different classes for all models. Focusing on the F1-score, which is an average estimation for both precision and recall, we see that the best performing classes are O2-O6, O7, O8, O9, B0, B3-B4, B5-B7, and B9 with a score > 0.65 , while the B1, B2, and B8 classes reach scores close to 0.60, 0.40, and 0.5, respectively. This means that all models are able to classify the majority of the spectral classes with relatively high accuracy, indicating that the RF algorithm can be used for such multiclass classification problems.

The reason behind the lower F1 score for these classes can be seen by examining the EW distributions (Fig. 6) of the specific spectral lines that distinguish each class from its neighboring classes. In particular, according to Table 2 from Maravelias et al. (2014), the main spectral lines that help to discriminate between stars of spectral classes B1 and B2 are Si IV/4088, Si IV/4116 and Si III/4553 Å. These lines show a strong overlap of their EW values, not only for the B1 and B2 spectral classes, but also for all the examined classes. That means that the RF classifier cannot learn anything from these features resulting in poor performance for these spectral classes. Exactly the same factor is responsible for the poor performance of the B8 class, although in this case the characteristic spectral lines that characterize it are He I/4144 and Si II/4128+4130. As is shown from their EW distributions, there is no clear separation of the B8 class from its adjacent classes B5-B7 and B9. The lack of useful information from these spectral lines, resulting from the strong overlap in all spectral classes, is also reflected in Fig. 7, where we see that the Si IV/4088-4116 and Si II/4128+4130 lines are among the less diagnostically important. Although He I/4144 does not

Fig. 8. *Top left panel:* confusion matrix of the best RF model applied to the test sample. *Right panel:* confusion matrix of the PRF best model applied to the same data set. *Bottom panel:* confusion matrix of the KDE-RF method. The overall accuracy is the same for all algorithms, 70%, with the majority of misclassified objects belonging to neighboring classes, indicating the reliability of the algorithms.

help to distinguish between the B8 and B9 classes, we see that its importance is higher due to the clear separation of the EW values for earlier spectral classes.

The predicted class of the RF algorithm is the one with the highest probability in the returned probability distribution. In Fig. 10, we plot the number of test sources predicted correctly and incorrectly versus the probability inferred from the RF algorithm. The two distributions have medians of 0.64 and 0.49, respectively, indicating that, in general, correctly predicted objects have higher probabilities than the incorrectly predicted ones. Specifically, above the threshold of 0.64, we have 558

true positives and 102 false negatives which means that the accuracy is $\sim 85\%$ in the reliable classification regime. In contrast, in the lower reliability probability range 0.49–0.64, the true positives are 334 and the false negatives are 139, corresponding to an accuracy of $\sim 71\%$. Below the threshold of 0.49, we have 224 true positives and 241 false negatives, or $\sim 48\%$ accuracy. With this in mind, we define three different confidence levels for the predicted spectral classes based on the aforementioned probability distributions: (a) above 0.64 (strong candidate); (b) 0.48–0.64 (good candidate); (c) below 0.48 (candidate).

Table 5. Precision, recall, and F1-score metrics per spectral type as determined by the best models of RF, PRF, and our method applied to a test sample of 1599 stars.

Class	RF			PRF			KDE-RF			Test sample
	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score	
O2-O6	0.93	0.97	0.95	0.95	0.97	0.96	0.93	0.97	0.95	58
O7	0.81	0.71	0.76	0.83	0.71	0.77	0.8	0.69	0.74	35
O8	0.81	0.79	0.8	0.82	0.74	0.78	0.78	0.82	0.79	38
O9	0.87	0.9	0.89	0.84	0.9	0.87	0.88	0.9	0.89	90
B0	0.79	0.77	0.78	0.81	0.82	0.82	0.77	0.8	0.79	111
B1	0.6	0.62	0.61	0.7	0.6	0.64	0.61	0.64	0.62	47
B2	0.6	0.61	0.61	0.73	0.48	0.58	0.59	0.67	0.63	67
B3-B4	0.72	0.73	0.73	0.7	0.74	0.72	0.73	0.74	0.73	282
B5-B7	0.65	0.69	0.67	0.61	0.74	0.67	0.68	0.63	0.65	375
B8	0.52	0.5	0.51	0.56	0.34	0.42	0.5	0.58	0.54	195
B9	0.73	0.69	0.71	0.71	0.72	0.71	0.72	0.67	0.7	301

Fig. 9. Comparison of the three models used in this work. Stars and triangles correspond to the results from the RF and PRF algorithms, respectively. The results from our method are shown with open circles. We find that all the methods are consistent, indicating the robustness of the RF algorithm for the automated spectral classification of OB stars. The majority of spectral classes are predicted with a score higher than 65%, while the lower performance of some classes is either due to their inadequate representation in the initial sample or due to overlap of the EW distributions of their characteristic spectral lines with those of other classes (see Sect. 5.1).

5.2. Sensitivity of classification on measurement and label uncertainty

In our analysis, we used the PRF algorithm for the first time (to our knowledge) in order to assess its performance in real observational data and compare it with the regular RF. The PRF was developed to incorporate the feature uncertainties during the training stage, and it is expected to exhibit an improvement of the order of 10% in accuracy compared to RF (Reis et al. 2019). However, our results show that the PRF and KDE-RF do not provide an improvement over the RF. The reason is that the performance of the algorithms is not driven by statistical scatter due to measurement errors. Instead, is mainly driven by the intrinsic

Fig. 10. Probability distributions between the sources that were predicted correctly (green dashed line) and incorrectly (red dashed line). The medians of the two distributions are 0.64 and 0.49, respectively. Based on these distributions, we define three flags regarding the reliability of our results. In particular, suggested spectral types with probability >0.64 are flagged as strong candidates, while those with probability <0.48 are flagged as candidates. For the intermediate range, $0.48-0.64$, they are flagged as good candidates.

scatter of the EW values within the range of the examined spectral classes, resulting in strong overlap between the EW distributions. This can also be seen from the Fig. A.1, where the median measurement uncertainties per spectral class are significantly lower than the scatter of the measurements. Furthermore, working on a real data set of observational data (spectra) rather than a synthetic one (as used in the development of the PRF method; Reis et al. 2019), may result in systematic errors due to noise, which corresponds to the use of different detectors and surveys.

Besides the measurement uncertainties, in this type of classification problems there are also uncertainties in the labels. The visual inspection of a stellar spectrum and the estimation of its spectral type is a relatively subjective process. It is very common for different people examining the same spectrum to conclude on slightly different spectral types. This label uncertainty exists even in spectral types that are assigned to stars through a template comparison or other automated procedure. Although with the PRF algorithm these label uncertainties can be included in

Fig. 11. Precision, recall, and F1 score as a function of the volume of the total simulated data used for the training of our RF-based model developed in Sect. 4.2. The total number of mock data divided by the total number of classes gives the number of mock data per class. The model’s performance does not depend on the size of the training sample, and it is mainly constrained from the parent distribution of the real data.

the training stage, we did not apply it in this work because they are not provided in most of the spectroscopic surveys we used. Only the LAMOST survey Liu et al. (2019) includes this information, claiming that all the spectra that were classified with the MKCLASS software package (Gray & Corbally 2014) are correct, with one subtype uncertainty. Because there is no easy way to work around this problem and the rejection of the other three surveys would limit our sample (especially with earlier type stars or lower metallicities; i.e., for SMC and LMC), we opted not to implement this option. Nonetheless, our results as illustrated in the confusion matrices (Fig. 8) show that the classification uncertainty is $\sim \pm 1$ class for later spectral classes (i.e., B5-B7, B8, B9) or better.

5.3. Volume of simulated training data

During the development of the KDE-RF method for the incorporation of uncertainties, we trained our model using a new balanced data set of mock data (see Sect. 4.2). In particular, we chose to randomly produce feature values for 1000 stars per spectral type, thus creating an overall training sample of 11 000 simulated data. However, an interesting question that arises is how the size of the sample affects the model’s performance.

To address this question, we followed a similar method to Clarke et al. (2020) by adjusting the mock training sample and assessing the model’s performance. We ranged the sample size from 100 to 2000 objects in steps of 100 objects for each of the 11 classes. In Fig. 11, we present how the basic metrics per class vary as a function of the number of the produced simulated data. The precision, recall, and F1 score are almost constant for all the spectral classes, indicating that the performance of our model does not depend on the volume of the mock data used for its training. Based on the above metrics, we see that even a ~ 1000 strong sample (i.e., ~ 100 per class) provides adequate sampling of the parent joint distribution of EW (as determined from the 17-D KDE). This could be the result of intrinsic scatter (or uncertainties) that smear the distribution and hide any complexity, in combination with smooth variations in the line strength as the temperature of a stellar atmosphere gradually decreases toward

later spectral types. The small variations in the precision, recall, and F1-score values are due to the inherent stochasticity in the training procedure of the RF algorithm (random selection of features and samples for the construction of each decision tree in the RF). For the same reason, the results will not be identical in each realization of the algorithm, but they will definitely be close to the average score value that the best model can reach.

5.4. Application of the RF best model on different science cases

The ultimate goal of this work is the application of the developed method on spectroscopic data sets of candidate OB-type stars for fast and reliable spectral classification. For this reason, we applied our best RF configuration to a number of science cases involving data of different qualities, providing examples of potential usage. During the data preparation for the analysis, we measured their EWs (following the same procedures described in Sect. 2.4).

For the first application, we used data from the IACOB project (Simón-Díaz et al. 2011a,b, 2015). In particular, we used spectra⁵ from DR1 and DR2, which includes high-resolution and good S/N (≥ 200) spectra from 328 Galactic massive OB-type stars with luminosity classes I–V. Although the IACOB spectra are of high quality, their spectral types are not well-determined since they come from different sources. With this in mind, we decided to revisit the spectral types of the IACOB project and suggest spectral-type classification based on the consistent application of our method. In Table 6, we present the results, and for comparison we also tabulate the spectral types provided from IACOB. We find that for 215 out of 328 stars, the suggested spectral classes from our model agree with those reported in the IACOB database. This score of $\sim 66\%$ is very close to the maximum accuracy that our classifier can reach. Furthermore, for the remaining 113 stars, the algorithm predicts neighboring spectral classes. In addition, regarding the confidence level of the predictions, we find that $\sim 42\%$ of the suggested spectral classes are flagged as “strong candidates”, $\sim 30\%$ as “good candidates”, and only $\sim 28\%$ are flagged as “candidates”. The results are very encouraging since they confirm that our model works properly on totally blind data sets, indicating the robustness and the reliability of our classifier.

Our second application is focused on Galactic BeXRBs. We used 12 spectra of the donor stars for nine well-studied BeXRB systems obtained with the 1.3 m telescope at the Skinas Observatory⁶. As is shown in Table 6, our model predicts spectral classes very close to those derived from manual classification. Notably, if we take into account vague uncertainties in our model, we can say with safety that the RF model works adequately. In addition, for each of the systems, 4U2206+54, IGR J06074+2205, and 1A 0535+262 were observed in two instances, and thus we tested our model on spectra of different qualities. For IGR J06074+2205 and 4U2206+54, we obtain the same correct prediction from both sources, while for 1A 0535+262 we obtain incorrect predictions from both observations of the source. However, these misclassifications are due to poor measurements, even for the most characteristic spectral lines of the B0 spectral class (e.g., He II/4541, He II/4686 etc):

⁵ The wavelength range of the observations is $\sim 3700\text{--}7000\text{ \AA}$ and the associated resolving power R of the spectra is $25000/46000/65000$ and 85000 for FIES and HERMES spectrographs, respectively.

⁶ The wavelength range of the observations is $\sim 3616\text{--}5717\text{ \AA}$ and the resolving power is $R \sim 2500$.

Table 6. Final products of our model's applications.

(1) IACOB ID	(2) Publ.	(3) RF	(4) Prob.	(5) Confidence level	(6) Probability per class										
					B0	B1	B2	B3-B4	B5-B7	B8	B9	O2-O6	O7	O8	O9
HD1279	B8	B5-B7	0.419	Candidate	0.014	0.0	0.001	0.029	0.419	0.316	0.193	0.006	0.007	0.009	0.005
HD1337	O9	O9	0.9036	Strong candidate	0.044	0.001	0.0	0.001	0.001	0.0	0.0	0.0	0.008	0.042	0.904
HD1438	B8	B8	0.4052	Candidate	0.002	0.0	0.001	0.011	0.255	0.405	0.325	0.0	0.0	0.0	0.001
HD1976	B5	B3-B4	0.8691	Strong candidate	0.001	0.006	0.034	0.869	0.084	0.001	0.005	0.0	0.0	0.0	0.0
HD2626	B7	B5-B7	0.6808	Strong candidate	0.002	0.001	0.01	0.107	0.681	0.098	0.101	0.0	0.0	0.0	0.0
HD2905	B0	B1	0.4515	Candidate	0.266	0.451	0.039	0.08	0.047	0.002	0.01	0.0	0.003	0.004	0.096
HD3240	B7	B5-B7	0.3528	Candidate	0.002	0.0	0.006	0.03	0.353	0.342	0.263	0.0	0.002	0.001	0.001
HD3369	B5	B3-B4	0.7726	Strong candidate	0.015	0.017	0.065	0.773	0.117	0.0	0.011	0.0	0.0	0.0	0.002
HD4142	B5	B3-B4	0.7976	Strong candidate	0.007	0.007	0.045	0.798	0.128	0.002	0.011	0.0	0.0	0.0	0.001
HD4382	B8	B5-B7	0.3811	Candidate	0.001	0.002	0.004	0.015	0.381	0.347	0.251	0.0	0.0	0.0	0.0
HD4636	B9	B5-B7	0.4275	Candidate	0.004	0.001	0.0	0.017	0.427	0.337	0.189	0.005	0.006	0.005	0.009
HD4727	B5	B3-B4	0.3826	Candidate	0.191	0.114	0.194	0.383	0.094	0.004	0.011	0.0	0.001	0.001	0.007
HD14272	B8	B5-B7	0.3388	Candidate	0.001	0.0	0.005	0.027	0.339	0.318	0.31	0.0	0.0	0.0	0.0
HD14818	B2	B1	0.3782	Candidate	0.057	0.378	0.108	0.321	0.109	0.005	0.015	0.0	0.001	0.001	0.004
HD14947	O4	O2-O6	0.9622	Strong candidate	0.001	0.0	0.0	0.0	0.001	0.001	0.002	0.962	0.021	0.002	0.01
HD15137	O9	O9	0.9234	Strong candidate	0.065	0.001	0.0	0.003	0.001	0.0	0.0	0.0	0.0	0.007	0.923
HD15318	B9	B9	0.8522	Strong candidate	0.001	0.003	0.0	0.005	0.037	0.101	0.852	0.001	0.0	0.0	0.0
HD15558	O4	O2-O6	0.8823	Strong candidate	0.005	0.002	0.0	0.0	0.001	0.001	0.002	0.882	0.038	0.003	0.065
HD15570	O4	O2-O6	0.876	Strong candidate	0.004	0.001	0.001	0.001	0.003	0.011	0.007	0.876	0.053	0.012	0.031
HD16046	B9	B9	0.8961	Strong candidate	0.0	0.002	0.0	0.005	0.024	0.071	0.896	0.002	0.0	0.0	0.0
HMXBs ID	Publ.	RF	Prob.	Confidence level	Probability per class										
					B0	B1	B2	B3-B4	B5-B7	B8	B9	O2-O6	O7	O8	O9
1A 0535+262_1st	B0	O2-O6	0.3274	Candidate	0.107	0.009	0.001	0.013	0.023	0.021	0.02	0.327	0.243	0.113	0.123
1A 0535+262_2nd	B0	B5-B7	0.2817	Candidate	0.215	0.068	0.006	0.148	0.282	0.09	0.094	0.012	0.028	0.015	0.042
SAX J2103.5+4545	B0	O9	0.4197	Candidate	0.185	0.029	0.009	0.056	0.078	0.012	0.02	0.003	0.053	0.136	0.42
RX J0440.9+4431_1st	B0	B0	0.3124	Candidate	0.312	0.169	0.031	0.185	0.118	0.034	0.041	0.015	0.019	0.023	0.052
RX J0440.9+4431_2nd	B0	B0	0.2393	Candidate	0.239	0.044	0.001	0.03	0.064	0.043	0.044	0.08	0.103	0.143	0.207
IGR J06074+2205_1st	B0	B0	0.3432	Candidate	0.343	0.245	0.045	0.127	0.085	0.048	0.035	0.012	0.014	0.009	0.038
IGR J06074+2205_2nd	B0	B0	0.3229	Candidate	0.323	0.281	0.159	0.203	0.018	0.003	0.002	0.0	0.002	0.002	0.007
RX J0240.4+6112	B0	O9	0.63067	Good candidate	0.277	0.022	0.008	0.009	0.006	0.001	0.001	0.0	0.008	0.039	0.631
2S0114+650	B1	B1	0.2458	Candidate	0.222	0.246	0.023	0.086	0.125	0.053	0.081	0.03	0.024	0.035	0.077
IGR J21343+473	B1	B0	0.3145	Candidate	0.314	0.239	0.309	0.106	0.012	0.002	0.002	0.0	0.0	0.0	0.015
RX J0146.9+6121	B1	B0	0.3594	Candidate	0.359	0.093	0.004	0.035	0.062	0.03	0.053	0.018	0.016	0.035	0.295
4U2206+54_1st	O9	O9	0.3997	Candidate	0.231	0.025	0.002	0.016	0.023	0.011	0.02	0.067	0.072	0.134	0.4
4U2206+54_2nd	O9	O9	0.497	Candidate	0.258	0.011	0.002	0.016	0.007	0.002	0.002	0.011	0.074	0.121	0.497

Notes. Only 20 objects from IACOB survey are tabulated. The full version is available at the CDS. Columns description: (1) Object's ID, (2) Published spectral type, (3) Predicted spectral class, (4) Probability of predicted spectral class, (5) Confidence level flag by following the scheme described at Sect. 5.1, (6) Probability per class. In the case of HMXBs spectra, 1st and 2nd indicate multiple spectra of the same object considered in our analysis.

almost all the spectral lines were replaced with magic numbers, indicating the low quality of these spectra. To conclude, our classifier performs well even in the special class of Oe and Be stars.

The above applications of the RF model prove that our proposed spectral-class classifier works well on various instances of real-world observational data constrained only by the quality range of the data on which it was trained. For this reason, we suggest to potential users of our tool that its performance is optimal for data with $S/N > 50$, and it requires spectral line measurements in the $\sim 3900\text{--}4900\text{ \AA}$ wavelength range and $\sim 1500\text{--}4000$ resolution.

6. Conclusions

The main results from this work can be summarized as follows:

1. We developed a new automated tool for the spectral classification of OB stars into their sub-classes, which is agnostic to the luminosity class of a star, and as much as possible, trained in a wide range of metallicities. The well-known RF algorithm, which is closer to the traditional way of optical spectroscopy in comparison with other automated methods is at the core of this new tool. In addition, we incorporated the feature uncertainties in our analysis by using the new PRF algorithm, as well as by developing our KDE-RF method based on data simulation. This method is based on optical
2. spectra of O-B type stars in the $\sim 3900\text{--}4900\text{ \AA}$ wavelength range and $> \sim 50\text{ S/N}$ range.
3. Comparing the three models, we find a similar accuracy score of $\sim 70\%$. This score is high if we consider the complexity of such multiclass classification problems (i.e., 11 classes), the intrinsic scatter of the EW distributions within the examined spectral classes and the diversity of the training set (resulting from the use data obtained from different surveys with different observing strategies). Given the lack of differences in the scores of the three methods, no method is considered superior. However, the proposed KDE-RF method is helpful for the data augmentation of sparsely populated samples such as those for the early-type stars. The similarity in the performance of our models indicates the robustness and the reliability of the RF algorithm when it is used for spectral classification of early-type stars that are isolated or companions in BeXRBs.
4. Using two different feature selection methods alongside the EW distribution per spectral line of our sample, we find that the full set of 17 spectral lines is needed to reach the maximum performance per spectral class. However, in Appendix B we propose a reduced ten-features scheme that can be applied to large data sets with lower S/Ns ($\sim 20\text{--}50$) and where the measurement of weaker spectral lines is not reliable.
5. In terms of the F1-score, which is an average estimation for both precision and recall, we see that the best performing

classes are the O2-O6, O7, O8, O9, B0, B3-B4, B5-B7, and B9, with a score >0.65 , while the B1, B2, and B8 classes reach scores close to 0.60, 0.40, and 0.5, respectively. This means that all models are able to classify, with relatively high accuracy, the majority of the spectral classes, indicating that the RF algorithm can be used for such multiclass classification problems.

5. We used the PRF algorithm for the first time and find that although it is more computationally intensive, it performs well when it is applied directly to real-world observational data. The general accuracy score does not change in our case. The uncertainty in the values of the features does not play such an important role since the intrinsic scatter of the EW distribution is much higher.
6. We applied our model to a number of real-world data, providing examples of possible usage of the automated spectral classification in different environments. We included either cases of isolated massive stars or companion stars in HMXBs in our applications in order to assess the applicability of our classifier in range of science cases. We find that the model works well, especially for data of similar quality to the training sample. The code for the application of our models is available through the GitHub repository⁷.

The main directions for improvement to increase the performance of our model are: (a) the use of a larger sample with better statistics in the under-represented classes, as well as better representation of different metallicity populations; and (b) the use of the RF algorithm as a regression tool in order to determine stellar temperatures, which is the underlying parameter driving spectral types. A reliable tool for spectral classification, such as the one we developed in this work, is particularly important for the analysis of data from large surveys that target thousands of stars (or galaxies) and produce millions of spectra (e.g., OWN, NoMaDS, CAFE-BEANS, etc.).

Acknowledgements. We thank the anonymous referee for his/her constructive comments and suggestions that helped to improve this work and the clarity of the manuscript. We thank Dr. Chris Evans and Dr. Philip Dufton for providing us the VFTS rectified data used in this work. In addition, we thank Dr. Itamar Reis for the useful comments regarding the PRF algorithm. EK acknowledges support from the Public Investments Program through a Matching Funds grant to the IA-FORTH. GM acknowledges funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement number 772086). The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 617001, and the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie RISE action, Grant Agreement n. 691164 (ASTROSTAT). We wish to thank the "2019 Summer School for Astrostatistics in Crete" for providing training on the statistical methods adopted in this work. We also thank Jeff Andrews for organizing the SMAC (Statistical methods for Astrophysics in Crete) seminar series.

References

- Abbott, B. P., Abbott, R., Abbott, T. D., & Acernese, F. 2017, *Phys. Lett. Rev.*, **119**, 161101
- Allende Prieto, C., Beers, T. C., Wilhelm, R., et al. 2006, *ApJ*, **636**, 804
- Antonioni, V., & Zezas, A. 2016, *MNRAS*, **459**, 528
- Antonioni, V., Zezas, A., Hatzidimitriou, D., & Kalogera, V. 2010, *ApJ*, **716**, L140
- Arnason, R. M., Barmby, P., & Vulic, N. 2020, *MNRAS*, **492**, 5075
- Baron, D. 2019, ArXiv e-prints [arXiv:1904.07248]
- Bouret, J. C., Lanz, T., Martins, F., et al. 2013, *A&A*, **555**, A1
- Breiman, L. 2001, *Mach. Learn.*, **45**, 5
- Carliles, S., Budavári, T., Heinis, S., Priebe, C., & Szalay, A. S. 2010, *ApJ*, **712**, 511
- Castro, N., Urbaneja, M. A., Herrero, A., et al. 2012, *A&A*, **542**, A79
- Castro, N., Protopapas, P., & Pichara, K. 2018, *AJ*, **155**, 16
- Clarke, A. O., Scaife, A. M. M., Greenhalgh, R., & Griguta, V. 2020, *A&A*, **639**, A84
- Cui, X.-Q., Zhao, Y.-H., Chu, Y.-Q., et al. 2012, *Res. Astron. Astrophys.*, **12**, 1197
- Duan, F.-Q., Liu, R., Guo, P., Zhou, M.-Q., & Wu, F.-C. 2009, *Res. Astron. Astrophys.*, **9**, 341
- Evans, C. J., Howarth, I. D., Irwin, M. J., Burnley, A. W., & Harries, T. J. 2004, *MNRAS*, **353**, 601
- Evans, C. J., Taylor, W. D., Hénault-Brunet, V., et al. 2011, *A&A*, **530**, A108
- Evans, C. J., Kennedy, M. B., Dufton, P. L., et al. 2015, *A&A*, **574**, A13
- Freeman, P., Doe, S., & Siemiginowska, A. 2001, in *Astronomical Data Analysis*, eds. J. L. Starck, & F. D. Murtagh, *SPIE Conf. Ser.*, **4477**, 76
- Gies, D. R., & Lambert, D. L. 1992, *ApJ*, **387**, 673
- Gkouvelis, L., Fabregat, J., Zorec, J., et al. 2016, *A&A*, **591**, A140
- Gray, R. O., & Corbally, C. J. 2009, *Stellar Spectral Classification*
- Gray, R. O., & Corbally, C. J. 2014, *AJ*, **147**, 80
- Kheirastan, S., & Bazarghan, M. 2016, *Ap&SS*, **361**, 304
- Kohavi, R. 1995, in *International Joint Conference on Artificial Intelligence (IJCAI)* (Morgan Kaufmann) 1137
- Ksoll, V. F., Gouliermis, D. A., Klessen, R. S., et al. 2018, *MNRAS*, **479**, 2389
- Kudritzki, R.-P., & Urbaneja, M. A. 2012, *Ap&SS*, **341**, 131
- Langer, N. 2012, *ARA&A*, **50**, 107
- Laurino, O., D'Abrusco, R., Longo, G., & Riccio, G. 2011, *MNRAS*, **418**, 2165
- Lee, U., Mathis, S., & Neiner, C. 2016, *MNRAS*, **457**, 2445
- Li, X.-R., Lin, Y.-T., & Qiu, K.-B. 2019, *Res. Astron. Astrophys.*, **19**, 111
- Linden, T., Kalogera, V., Sepinsky, J. F., et al. 2010, *ApJ*, **725**, 1984
- Liu, Z., Cui, W., Liu, C., et al. 2019, *ApJS*, **241**, 32
- Loupe, G. 2014, ArXiv e-prints [arXiv:1407.7502]
- Maeder, A., Meynet, G., Ekström, S., Hirschi, R., & Georgy, C. 2008, *A&A*, **250**, 3
- Mahabal, A., Djorgovski, S. G., Williams, R., et al. 2008, in *Classification and Discovery in Large Astronomical Surveys*, ed. C. A. L. Bailer-Jones, *AIP Conf. Ser.*, **1082**, 287
- Mahabal, A., Sheth, K., Gieseke, F., et al. 2017, *IEEE Symp. Ser. on Computational Intelligence (SSCI)*, Honolulu, HI, USA, 2757
- Mahdi, B. 2008, *Bull. Astron. Soc. India*, **36**, 1
- Maíz Apellániz, J., Sota, A., Walborn, N. R., et al. 2011, in *Highlights of Spanish Astrophysics VI*, eds. M. R. Zapatero Osorio, J. Gorgas, J. Maíz Apellániz, J. R. Pardo, & A. Gil de Paz, 467
- Maíz Apellániz, J., Sota, A., Morrell, N. I., et al. 2013, in *Massive Stars: From alpha to Omega*, 198
- Maravelias, G. 2014, *Investigation of the High-Mass X-ray Binary populations in the Small Magellanic Cloud* (Ph.D. Thesis)
- Maravelias, G., Zezas, A., Antoniou, V., & Hatzidimitriou, D. 2014, *MNRAS*, **438**, 2005
- Maravelias, G., Zezas, A., Antoniou, V., Hatzidimitriou, D., & Haberl, F. 2017, in *The Lives and Death-Throes of Massive Stars*, eds. J. J. Eldridge, J. C. Bray, L. A. S. McClelland, & L. Xiao, 329, 373
- Martayan, C., Floquet, M., Hubert, A. M., et al. 2007, *A&A*, **472**, 577
- Massey, P. 2003, *ARA&A*, **41**, 15
- McSwain, M. V., & Gies, D. R. 2005, *ApJS*, **161**, 118
- Morgan, W. W., Keenan, P. C., & Kellman, E. 1943, *An Atlas of Stellar Spectra, with an Outline of Spectral Classification*
- Müller, P. E., & Vink, J. S. 2014, *A&A*, **564**, A57
- Navarro, S. G., Corradi, R. L. M., & Mampaso, A. 2012, *A&A*, **538**, A76
- Pearson, W. J., Wang, L., Trayford, J. W., Petrillo, C. E., & van der Tak, F. F. S. 2019, *A&A*, **626**, A49
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Mach. Learn. Res.*, **12**, 2825
- Pellerin, A., Maíz Apellániz, J., Simón-Díaz, S., & Barbá, R. H. 2012, *Am. Astron. Soc. Meet. Abstracts*, **219**, 224.03
- Pesenson, M. Z., Pesenson, I. Z., & McCollum, B. 2010, *Adv. Astron.*, **2010**, 350891
- Poelarends, A. J. T., Herwig, F., Langer, N., & Heger, A. 2008, *ApJ*, **675**, 614
- Porter, J. M., & Rivinius, T. 2003, *PASP*, **115**, 1153
- Preibisch, T., & Zinnecker, H. 2007, in *Triggered Star Formation in a Turbulent ISM*, eds. B. G. Elmegreen, & J. Palous, *IAU Symp.*, **237**, 270
- Przybilla, N., Nieva, M. F., Irrgang, A., & Butler, K. 2013, in *EAS Publications Series*, eds. G. Alecian, Y. Lebreton, O. Richard, & G. Vauclair, *EAS Publ. Ser.*, **63**, 13
- Pudil, P., Novovičová, J., & Kittler, J. 1994, *Pattern Recogn. Lett.*, **15**, 1119
- Raschka, S. 2018, *J. Open Source Soft.*, **3**
- Reig, P. 2011, *Ap&SS*, **332**, 1
- Reis, I., Baron, D., & Shahaf, S. 2019, *AJ*, **157**, 16
- Russell, S. C., & Dopita, M. A. 1992, *ApJ*, **384**, 508
- Sana, H. 2017, in *The Lives and Death-Throes of Massive Stars*, eds. J. J. Eldridge, J. C. Bray, L. A. S. McClelland, & L. Xiao, *IAU Symp.*, **329**, 110

⁷ <https://github.com/EliasKyritsis/SpectralClassifier>

- Sana, H., de Koter, A., de Mink, S. E., et al. 2013, [A&A](#), **550**, [A107](#)
- Scibelli, S., Newberg, H. J., Carlin, J. L., & Yanny, B. 2014, [ApJS](#), **215**, [24](#)
- Sharma, K., Kembhavi, A., Kembhavi, A., et al. 2020, [MNRAS](#), **491**, [2280](#)
- Simón-Díaz, S., Castro, N., Garcia, M., Herrero, A., & Markova, N. 2011a, [Bull. Soc. R. Sci. Liege](#), **80**, [514](#)
- Simón-Díaz, S., Garcia, M., Herrero, A., Maíz Apellániz, J., & Negueruela, I. 2011b, in [Stellar Clusters & Associations: A RIA Workshop on Gaia](#), 255
- Simón-Díaz, S., Negueruela, I., Maíz Apellániz, J., et al. 2015, in [Highlights of Spanish Astrophysics VIII](#), eds. A. J. Cenarro, F. Figueras, C. Hernández-Monteagudo, J. Trujillo Bueno, & L. Valdivielso, 576
- Smartt, S. J. 2009, [ARA&A](#), **47**, [63](#)
- Smith, M. D. 2014, [MNRAS](#), **438**, [1051](#)
- Steidel, C. C., Giavalisco, M., Pettini, M., Dickinson, M., & Adelberger, K. L. 1996, [ApJ](#), **462**, [L17](#)
- Tzanavaris, P., Fragos, T., Tremmel, M., et al. 2013, [ApJ](#), **774**, [136](#)
- Urbaneja, M. A., Herrero, A., Bresolin, F., et al. 2005, [ApJ](#), **622**, [862](#)
- Vasconcellos, E. C., de Carvalho, R. R., Gal, R. R., et al. 2011, [AJ](#), **141**, [189](#)
- Vink, J. S., & Gräfener, G. 2012, [ApJ](#), **751**, [L34](#)
- von Hippel, T., Storrie-Lombardi, L. J., Storrie-Lombardi, M. C., & Irwin, M. J. 1994, in [The MK Process at 50 Years: A Powerful Tool for Astrophysical Insight](#), eds. C. J. Corbally, R. O. Gray, & R. F. Garrison, [ASP Conf. Ser.](#), **60**, [289](#)
- Walborn, N. R., & Fitzpatrick, E. L. 1990, [PASP](#), **102**, [379](#)
- Wenger, M., Ochsenbein, F., Egret, D., et al. 2000, [A&AS](#), **143**, [9](#)
- Zhao, G., Zhao, Y.-H., Chu, Y.-Q., Jing, Y.-P., & Deng, L.-C. 2012, [Res. Astron. Astrophys.](#), **12**, [723](#)

Appendix A: An alternative bands-based method for measuring EW and its uncertainty

Although the spectral fitting method (Sec. 2.4) provides the most reliable measurements of the EW, it is computationally intensive and it often requires optimization for the particular data set it is applied to. For this reason, we also consider the more commonly used bands-based method. In this method the EW of a given line is obtained from measurement of the intensity in three bands: one centered on the spectral line of interest and two used to measure the continuum on the blue and the red sides of the spectral line, respectively. For the calculation of the EW, the basic formula in Equation 2 can be transformed as

$$EW = -\left(d \times N - \sum_{i=1}^N \frac{F_{line_i}}{F_{cont_i}} d\right) = -\left(d \times N - d \times \frac{1}{C} \sum_{i=1}^N F_{line_i}\right), \quad (A.1)$$

since in our case the spectral density is measured in individual pixels and their relation to the wavelength is given by the dispersion of the spectrum (d). The flux of the continuum is considered constant (C) over the wavelength range of N pixels of the line. Thus, we can measure the total flux of the spectral line as the sum of the pixel values included within the λ_1 - λ_2 wavelength region. To estimate the continuum flux density at the location of the line, we interpolated between the flux density at the red and the blue continuum regions based on Equation A.2:

$$C = C_{blue} + \frac{C_{red} - C_{blue}}{\lambda_{red} - \lambda_{blue}}(\lambda_{line} - \lambda_{blue}), \quad (A.2)$$

where C_{blue} and C_{red} are the average values for the continuum flux density at the blue and red sides (in $\text{\AA}/\text{pix}$), and λ_{blue} , λ_{red} , and λ_{line} are the central wavelengths of the blue, red continuum, and the line regions, respectively. Appropriate wavelength ranges for the spectral lines and their continuum regions were defined after visual inspection to ensure that they did not include other lines or artifacts (Maravelias 2014). Next, we calculated the mean intensity at the center of each continuum side, and by following the above procedure, we calculated the continuum flux density at the central wavelength of each line. In addition, when we measured the average continuum flux densities, we applied a sigma clipping method with an accepting threshold of 1σ and by iterating once. In this way, we measured a pure continuum region that was not biased from outliers (i.e., spectral lines or artefacts). As a final step, this mean continuum value was used as an input in Equation A.1 to measure the EW of each spectral line.

For the calculation of the EW uncertainty, we applied error propagation to Equation A.1:

$$\delta EW = \sqrt{\left(\frac{(\sum_{i=1}^N F_{line_i})\delta C}{C^2}\right)^2 + \left(\frac{(\sum_{i=1}^N \delta F_{line_i}^2)^{1/2}}{C}\right)^2}. \quad (A.3)$$

The calculation of the total continuum uncertainty (δC) and the total flux uncertainty of the line (δF_{line}) depends on the availability of flux uncertainties in the initial spectral data. Thus, for the 2dF, VFTS, and GOSC data, from which the information of flux uncertainty was missing, the second term in Equation A.3 was equal to zero, and the continuum uncertainty was calculated from the standard error of the mean of the flux of the pixels in the continuum region:

$$\delta C = \sqrt{\left(\frac{\sigma_{C_{blue}}}{\sqrt{N_{C_{blue}}}}\right)^2 + \left(\frac{\sigma_{C_{red}}}{\sqrt{N_{C_{red}}}}\right)^2}. \quad (A.4)$$

On the other hand, for the LAMOST data that provide flux uncertainty for each spectral bin, the total flux uncertainty of the line was calculated as

$$\delta F_{line} = \sqrt{\sum_{i=1}^N \delta F_i^2}, \quad (A.5)$$

where δF_i is the flux uncertainty of N -th pixel in the line. Finally, the total continuum uncertainty was calculated as

$$\delta C = \sqrt{\left(\frac{\sigma_{\delta C_{blue}}}{\sqrt{N_{C_{blue}}}}\right)^2 + \left(\frac{\sigma_{\delta C_{red}}}{\sqrt{N_{C_{red}}}}\right)^2}, \quad (A.6)$$

where we used the standard deviation of the flux uncertainties of the blue and red continuum regions, respectively, instead of the standard deviation of the flux values that we used in Equation A.4. We applied this method to the entire data set of the spectra using the same spectral lines and corresponding bands presented in Sect. 2.4.

In Fig. A.1, we present the one-to-one comparison of the EW measurements per spectral line and per spectral class between the spectral fit and the bands method. With black crosses, we depict the median EW uncertainties per spectral class. As shown, the two methods are consistent mainly for the strongest spectral lines (e.g., He II, He I/4471 $\text{\AA}$, He I/4144 $\text{\AA}$, etc.). On the other hand, for weaker spectral lines (e.g., He I/4009 $\text{\AA}$, He I + He II/4026 $\text{\AA}$, etc.) the bands-based method seems to overestimate the EW value, since in these cases the measurements are dominated by continuum noise. This is particularly prominent in the case of the Si IV/4088 $\text{\AA}$ and Si IV/4116 $\text{\AA}$ spectral lines. These two lines are relatively weak and close to the H δ Balmer line, resulting in strong blending, especially for late B-type stars with strong Balmer lines. This effect makes the measurement of these two lines very unreliable because the bands-based method actually measures the broad H δ wings instead of the Si IV spectral lines. However, this is solved with the spectral line fitting method presented in Sect. 2.4. Regarding the EW uncertainties measured from both methods, we find that the fit method estimates the EW error slightly more accurately than the bands method. The most important thing, however, is that the scatter of the data is not driven by the uncertainties of the measurements (which are significantly lower), but from the intrinsic scatter of the EW values that hampers the separation of adjacent spectral classes based on the EW of characteristic spectral lines.

To compare the performance of the RF algorithm, we trained it on the same data set but with EW measurements obtained from both methods. We used the 15 spectral lines that can be measured either from the fit method or from the bands-based method as features, and we excluded Si IV/4088 $\text{\AA}$ and Si IV/4116 $\text{\AA}$, for which the bands-based method provides unreliable measurements. In Fig. A.2, we present the results of this comparison. As shown from the confusion matrices of the two methods, they present similar behavior in terms of general accuracy, as well as in terms of the recall score per spectral class. In addition, the k-fold validation for both gives exactly the same CV score, $69 \pm 1\%$. The consistency between the measurements of both methods is further proof that the score of the RF algorithm is driven by the intrinsic scatter of the EW values.

To conclude, we present an alternative method for measuring the EW of characteristic spectral lines. Although in weak spectral lines it tends to overestimate the EW, in general it demonstrates similar performance to the basic fit method we presented in the main paper. However, the fit method is of course the most

Fig. A.1. One-to-one comparison of the EW measurements per spectral line and per spectral class between the two methods. Black crosses depict the median EW uncertainty per spectral class.

accurate method that can be applied to such of problems. We suggest that potential users of our tool use this bands-based method if they are interested in the classification of large data

sets, since the time needed to measure the EW with spectral fitting of the lines is significantly higher than with the bands-based method.

Fig. A.2. Top right panel shows the normalized confusion matrix of the best RF model trained on EW values obtained from the fit method, while the left panel shows the normalized confusion matrix of the RF best model trained on values obtained from bands-based method. The k-fold validation for both gives exactly the same CV score, $69 \pm 1\%$, indicating that any failure of the RF algorithm to correctly predict all the spectral classes is not due to inaccurate measurements but due to the intrinsic scatter of the EW within this range of spectral classes.

Appendix B: A reduced feature scheme for the application of the RF-bands method

Although in our analysis we considered the full scheme of lines, aiming to receive the maximum performance of our model per spectral class, in Table B.1 we present an alternative scheme that can also be used for the spectral classification of OB stars in green. In particular, it comprises the ten most important spectral features resulting from the SFFS analysis in Sect. 3.3. This reduced scheme includes the strong spectral lines He I/4471 and Mg II/4481, the most important He II lines (indicators of O-type stars), as well as He I lines, which are good indicators of the early- and mid-B-type stars, while the weak Si IV/4088-4116 are not included. The absence of Si IV lines is particularly important because one can use the bands-based method, which, as we discussed in Appendix A, overestimates their EW measurements. Furthermore, given the comparison between the bands-based and the fit method for the reduced scheme (see Fig. A.1), we would expect the bands-based EW measurements to yield a score similar to $\sim 68\%$, as is suggested by the SFFS algorithm for the fit method. The overall performance of $\sim 68\%$ suggests that it can be applied to large data sets with lower S/Ns ($\sim 20 - 50$), where the measurement of weaker lines is not reliable.

Table B.1. Reduced classification scheme, which can be measured with the bands-based method.

Line ID	λ_{central} (Å)
He I	4009
He I+He II	4026
Si IV	4088
Si IV	4116
He I	4121
Si II	4128+4130
He I	4144
He II	4200
Fe II	4233
He I	4387
O II	4416
He I	4471
Mg II	4481
He II	4541
Si III	4553
O II+C III	4645
He II	4686