Optimizing spectroscopic follow-up strategies for supernova photometric classification with active learning

E. E. O. Ishida,¹★ R. Beck,^{2,3} S. González-Gaitán⁴, R. S. de Souza,⁵
A. Krone-Martins,⁶ J. W. Barrett,⁷ N. Kennamer,⁸ R. Vilalta⁹, J. M. Burgess¹⁰,
B. Quint¹¹, A. Z. Vitorelli,¹² A. Mahabal,¹³ E. Gangler¹ and for the COIN
collaboration

¹Université Clermont Auvergne, CNRS/IN2P3, LPC, F-63000 Clermont-Ferrand, France

²Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA

³Department of Physics of Complex Systems, Eötvös Loránd University, Pf. 32, H-1518 Budapest, Hungary

⁴CENTRA/COSTAR, Instituto Superior Técnico, Universidade de Lisboa, Av. Rovisco Pais 1, 1049-001 Lisboa, Portugal

⁵Department of Physics & Astronomy, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3255, USA

⁶CENTRA/SIM, Faculdade de Ciências, Universidade de Lisboa, Ed. C8, Campo Grande, 1749-016, Lisboa, Portugal

⁷Institute of Gravitational-wave Astronomy and School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

⁸Department of Computer Science, University of California Irvine, Donald Bren Hall, Irvine, CA 92617, USA

⁹Department of Computer Science, University of Houston, 3551 Cullen Blvd, 501 PGH, Houston, TX 77204-3010, USA

¹⁰Max-Planck-Institut für extraterrestrische Physik, Giessenbachstrasse, D-85748 Garching, Germany

¹¹SOAR Telescope, AURA-O, Colina El Pino S/N, Casila 603, La Serena, Chile

¹²Instituto de Astronomia, Geofísica e Ciências Atmosféricas, Universidade de São Paulo, São Paulo, SP, Brazil

¹³Center for Data-Driven Discovery, California Institute of Technology, Pasadena, CA 91125, USA

Accepted 2018 November 5. Received 2018 November 4; in original form 2018 April 19

ABSTRACT

We report a framework for spectroscopic follow-up design for optimizing supernova photometric classification. The strategy accounts for the unavoidable mismatch between spectroscopic and photometric samples, and can be used even in the beginning of a new survey – without any initial training set. The framework falls under the umbrella of active learning (AL), a class of algorithms that aims to minimize labelling costs by identifying a few, carefully chosen, objects that have high potential in improving the classifier predictions. As a proof of concept, we use the simulated data released after the SuperNova Photometric Classification Challenge (SNPCC) and a random forest classifier. Our results show that, using only 12 per cent the number of training objects in the SNPCC spectroscopic sample, this approach is able to double purity results. Moreover, in order to take into account multiple spectroscopic observations in the same night, we propose a semisupervised batch-mode AL algorithm that selects a set of $N = 5$ most informative objects at each night. In comparison with the initial state using the traditional approach, our method achieves 2.3 times higher purity and comparable figure of merit results after only 180 d of observation, or 800 queries (73 per cent of the SNPCC spectroscopic sample size). Such results were obtained using the same amount of spectroscopic time necessary to observe the original SNPCC spectroscopic sample, showing that this type of strategy is feasible with current available spectroscopic resources. The code used in this work is available in the [COINtoolbox](#).

Key words: methods: data analysis – methods: observational – supernovae: general.

1 INTRODUCTION

The standard cosmological model rests on three observational pillars: primordial big bang nucleosynthesis (Gamow 1948), the cosmic microwave background radiation (Spergel et al. 2007; Planck Collaboration I 2016), and the accelerated cosmic expansion

* E-mail: emille.ishida@clermont.in2p3.fr

(Riess et al. 1998; Perlmutter et al. 1999) – with type Ia supernovae (SNe Ia) playing an important role in probing the last one. SNe Ia are astronomical transients that are used as standardizable candles in the determination of extragalactic distances and velocities (Hillebrandt & Niemeyer 2000; Goobar & Leibundgut 2011). However, between the discovery of an SN candidate and its successful application in cosmological studies, additional research steps are necessary.

Once a transient is identified as a potential SN, it must go through three main steps: (i) classification, (ii) redshift estimation, and (iii) estimation of its standardized apparent magnitude at maximum brightness (Phillips 1993; Tripp 1998). Ideally, each SN thus requires at least one spectroscopic observation (preferably around maximum – items i and ii) and a series of consecutive photometric measurements (item iii). Since we are not able to get spectroscopic measurement for all transient candidates, soon after a variable source is detected a decision must be made regarding its spectroscopic follow-up, making coordination between transient imaging surveys and spectroscopic facilities mandatory. From a traditional perspective, taking a spectrum of a transient that ends up classified as an SN Ia results in the object being included in the cosmological analysis. On the other hand, if the target is classified as non-Ia, spectroscopic time for cosmology is essentially considered ‘lost’.¹

In the last couple of decades, a strong community effort has been devoted to the detection and follow-up of SNe Ia for cosmology. Classifiers (human or artificial) on which follow-up decisions are based have become increasingly efficient in identifying SNe Ia from early stages of their light-curve evolution – successfully targeting them for spectroscopic observations (e.g. Perrett et al. 2010). The available cosmology data set has grown from 42 (Perlmutter et al. 1999) to 740 (Betoule et al. 2014) in that period of time. This success helped building consensus around the paramount importance of SNe Ia for cosmology. It has also encouraged the community to add even more objects to the Hubble diagram and to investigate the systematics uncertainties that currently dominate dark energy error budget (e.g. Conley et al. 2011). Thenceforth, SNe Ia are major targets of many current, e.g. *Dark Energy Survey*² (DES), *Panoramic Survey Telescope and Rapid Response System*³ (Pan-STARRS), and upcoming surveys, e.g. *Zwicky Transient Facility*⁴ (ZTF) and *Large Synoptic Survey Telescope*⁵ (LSST). These latter new surveys are expected to completely change the data paradigm for SN cosmology, increasing the number of available light curves by a few orders of magnitude.

However, to take full advantage of the great potential in such large photometric data sets, we still have to contend with the fact that spectroscopic resources are – and will continue to be – scarce. The majority of photometrically identified candidates will never be followed spectroscopically. Full cosmological exploitation of wide-field imaging surveys necessarily requires a framework able to infer reliable spectroscopically derived features (redshift and class) from purely photometric data. Provided a particular transient has an identifiable host, redshift can be obtained before/after the event

from the host observations (spectroscopic or photometric) or even from the light curve itself (e.g. Wang, Gjergo & Kuhlmann 2015). On the other hand, classification should primarily be inferred from the light curve.⁶ This paper concerns itself with the latter.

Before we dive into the details of SN photometric classification, it is important to keep in mind that, regardless of the method chosen to circumvent these issues, photometric information will always carry a larger degree of uncertainty than those from the spectroscopic scenario. Photometric redshift estimations are expected to have non-negligible error bars and, at the same time, any kind of classifier will carry some contamination to the final SN Ia sample. Nevertheless, if we manage to keep these effects under control, we should be able to use photometrically observed SNe Ia to increase the statistical significance of our results. The question whether the final cosmological outcomes surpass those of the spectroscopic-only sample enough to justify the additional effort is still debatable. Despite a few reports in this direction using real data from the *Sloan Digital Sky Survey*⁷ (SDSS; Hlozek et al. 2012; Campbell et al. 2013) and Pan-STARRS (Jones et al. 2017), the answer keeps changing as different steps of the pipeline are improved and more data become available. Nevertheless, there seems to be a consensus in the astronomical community that we have much to gain from such an exercise.

The vast literature, with suggestions on how to improve/solve different stages of the SN photometric classification pipeline, is a testimony of the positive attitude with which the subject is approached. For more than 15 yr, the field has been overwhelmed with attempts relying on a wide range of methodologies: colour–colour and colour–magnitude diagrams (Poznanski et al. 2002; Johnson & Crotts 2006), template fitting techniques (e.g. Sullivan et al. 2006), Bayesian probabilistic approaches (Poznanski, Maoz & Gal-Yam 2007; Kuznetsova & Connolly 2007), fuzzy set theory (Rodney & Tonry 2009), kernel density methods (Newling et al. 2011), and more recently, machine learning-based classifiers (e.g. Richards et al. 2012a; Ishida & de Souza 2013; Karpenka, Feroz & Hobson 2013; Lochner et al. 2016; Möller et al. 2016; Charnock & Moss 2017; Dai et al. 2018).

In 2010, the *SuperNova Photometric Classification Challenge* (SNPCC; Kessler et al. 2010) summarized the state of the field by inviting different groups to apply their classifiers to the same simulated data set. Participants were asked to classify a set of light curves generated according to the DES photometric sample characteristics. As a starting point, they were provided with a spectroscopic sample enclosing ~5 per cent of the total data set, and for which class information was disclosed. The organizers posed three main questions: full light-curve classification with and without the use of redshift information (supposedly obtained from the host galaxy) and an early epoch classification – where participants were allowed to use only the first six observed points from each light curve. The goal of the latter was to access the capability of different algorithms to advise on spectroscopic targeting while the SN was still bright enough to allow it. A total of 10 groups replied to the call, submitting 13 (9) entries for the full light-curve classification with (without) the use of redshift information. No submission was received for the early epoch scenario.

The algorithms competing in the SNPCC were quite diverse, including template fitting, statistical modelling, selection cuts and machine learning-based strategies (see summary of all partici-

¹This is strictly for cosmological purposes; spectroscopic observations are extremely valuable, irrespective of the transient in question, though for different scientific goals.

²<https://www.darkenergysurvey.org/>

³<https://panstarrs.stsci.edu/>

⁴<http://www.ptf.caltech.edu/ztf/>

⁵<https://www.lsst.org/>

⁶Although, see Foley & Mandel (2013).

⁷<http://www.sdss.org/>

pants and result in Kessler et al. 2010). Classification results were consistent among different methods with no particular algorithm clearly outperforming all the others. The main legacy of this initiative however, was the updated public data set made available to the community once the challenge was completed. It became a test bench for experimentation, specially for machine learning approaches (Newling et al. 2011; Richards et al. 2012a; Karpenka et al. 2013; Ishida & de Souza 2013).

One particularly challenging characteristic of the SN classification problem, also present in the SNPCC data, is the discrepancy between spectroscopic and photometric samples. In a supervised machine learning framework, we have no alternative other than to use spectroscopically classified SNe as training. This turns out to be a serious problem, since one of the hypothesis behind most textbook learning algorithms relies on training being representative of target. Due to the stringent observational requirements of spectroscopy, this will never be the case between spectroscopic and photometric astronomical samples. But the situation is even more drastic for SNe, where the spectroscopic follow-up strategy was designed to target as many Ia-like objects as possible. Albeit modern low-redshift surveys try to mitigate and counterbalance this effect (e.g. ASASSN⁸ and iPTF⁹), the medium/high redshift ($z > 0.1$) spectroscopic sample is still heavily underrepresented by all non-Ia SNe types. Spectroscopic observations are so time demanding, and the rate with which the photometric samples are increasing is so fast, that the situation is not expected to change even with dedicated spectrographs (OzDES – Childress et al. 2017). This issue has been pointed out by many post-SNPCC machine learning-based analysis (e.g. Richards et al. 2012a; Karpenka et al. 2013; Varughese et al. 2015; Lochner et al. 2016; Charnock & Moss 2017; Revsbech, Trotta & van Dyk 2018). In spite of the general consensus being that one should prioritize faint objects for spectroscopic targeting, as an attempt to increase representativeness (Richards et al. 2012a; Lochner et al. 2016), the details on how exactly this should be implemented are yet to be defined.

Thus, the question still remains: how do we optimize the distribution of spectroscopic resources with the goal of improving photometric SN identification? Or, in other words, how do we construct a training sample that maximizes accurate classification with a minimum number of labels, i.e. spectroscopically classified SNe? The above question is similar in context to the ones addressed by an area of machine learning called *active learning* (AL; Cohn, Ghahramani & Jordan 1996; Balcan, Beygelzimer & Langford 2009; Settles 2012).

AL iteratively identifies which objects in the target (photometric) sample would most likely improve the classifier if included in the training data, allowing sequential updates of the learning model with a minimum number of labelled instances. It has been widely used in a variety of different research fields, e.g. natural language processing (Thompson, Califf & Mooney 1999), spam classification (DeBarr & Wechsler 2009), cancer detection (Liu 2004), and sentiment analysis (Kranjc et al. 2015). In astronomy, AL has already been successfully applied in multiple tasks: determination of stellar population parameters (Solorio et al. 2005), classification of variable stars (Richards et al. 2012b), optimization of telescope choice (Xia, Protopapas & Doshi-Velez 2016), static supernova photometric classification (Gupta et al. 2016), and photometric redshift

Figure 1. Comparison between simulated peak magnitudes in the SNPCC spectroscopic (red – training) and photometric (blue – target) samples. Violin plots show both distributions in each of the DES filters.

estimation (Vilalta et al. 2017). There are also reports based on similar ideas by Masters et al. (2015) and Hoyle et al. (2016).

In this work, we show how AL enables the construction of optimal training data sets for SNe photometric classification, providing observers with a spectroscopic follow-up protocol on a night-by-night basis. The framework respects the time evolution of the survey providing a decision process that can be implemented from the first observational nights, avoiding the necessity of adapting legacy data and the consequent translation between different photometric systems. The methodology herein employed virtually allows any machine learning technique to outperform itself by simply labelling the data (taking the spectrum) in a way that provides maximum information gain to the classifier. As a case study, we focus on the problem of binary classification, i.e. type Ia versus non-Ia, but the overall structure can be easily generalized for multiclassification tasks.

This paper is organized as follows. Section 2 describes the SNPCC data set, emphasizing the differences between spectroscopic and photometric samples. Section 3 details on the feature extraction, classifier, and metrics used throughout the paper. Section 4 explains the AL algorithm, details the configurations we chose to approach the SN photometric classification problem, and presents results for the idealized static, full light-curve scenario. Section 5 explains how our approach deals with the real-time evolution of the survey and its consequent results. Section 6 presents our proposal for real-time multiple same-night spectroscopic targeting (batch-mode AL). Having established a baseline data driven spectroscopic strategy, in Section 7 we estimate the amount of telescope time necessary to observe the AL queried sample. Lastly, Section 8 shows our conclusions.

2 DATA

In what follows, we use the data released after the SNPCC. This is a simulated data set constructed to mimic DES observations. The sample contains 20 216 supernovae (SNe) observed in four DES filters, $\{g, r, i, z\}$, among which a subset of 1103 are identified as belonging to the *spectroscopic* sample. This subset was constructed considering observations through a 4m (8m) telescope and limiting r -band magnitude of 21 (23.5) (Kessler et al. 2010). Thus, it resembles closely biases foreseen in a realistic spectroscopic sample when compared to the photometric one. Among them, we highlight the predominance of brighter objects (Fig. 1) with higher signal-to-noise ratio (SNR; Fig. 2), and the predominance of SNe Ia over other SN types (Fig. 3). Hereafter, the spectroscopic sample will be

⁸<http://www.astronomy.ohio-state.edu/assassin/index.shtml>

⁹<https://www.ptf.caltech.edu/iptf>

Figure 2. Distribution of mean SNR in the SNPCC spectroscopic (red – training) and photometric (blue – target) samples.

Figure 3. Populations of different supernova types in the SNPCC spectroscopic (training) and photometric (target) samples. The spectroscopic sample holds 599 (51 per cent) Ia, 144 (13 per cent) Ibc, and 400 (36 per cent) II, while the photometric sample comprises 4326 (22 per cent) Ia, 2535 (13 per cent) Ibc, and 12 442 (65 per cent) II.

designated *SNPCC spec* and the remaining objects will be addressed as *SNPCC photo*.

3 TRADITIONAL ANALYSIS

3.1 Feature extraction

For each supernova, we observe its light curve, i.e. the evolution of brightness (flux) as a function of time, in four DES filters $\{g, r, i, z\}$. For most machine learning applications, this information needs to be homogenized before it can be used as input to a learning algorithm.¹⁰ There are many ways in which this feature extraction step can be performed: via a proposed analytical parametrization (Bazin et al. 2009; Newling et al. 2011), comparisons with theoretical

¹⁰Exceptions include algorithms able to deal with a high degree of missing data (e.g. Charnock & Moss 2017; Naul et al. 2018).

Figure 4. Example of light curve fitted with the parametric function of Bazin et al. (2009) – equation (1). The plot shows measurements for a typical well-sampled type Ia at redshift $z \sim 0.20$, in each one of the four DES filters (dots and error bars) as well as its best-fitted results (lines).

and/or well-observed templates (Sako et al. 2008) or dimensionality reduction techniques (Richards et al. 2012a; Ishida & de Souza 2013). Literature has many examples showing that, for the same classification algorithm, the choice of the feature extraction method can significantly impact classification results (see Lochner et al. 2016, and references therein).

In what follows, we use the parametrization proposed by Bazin et al. (2009),

$$f(t) = A \frac{e^{-(t-t_0)/\tau_f}}{1 + e^{(t-t_0)/\tau_r}} + B, \quad (1)$$

where A, B, t_0, τ_f , and τ_r are parameters to be determined. We fit each filter independently in flux space with a Levenberg–Marquardt least-square minimization (Madsen, Nielsen & Tingleff 2004). Fig. 4 shows an example of flux measurements, corresponding errors and best-fitting results in all four filters for a typical, well-sampled, SN Ia from SNPCC data.

Although not optimal for such a diverse light-curve sample, the parametrization given by equation (1) was chosen for being a fast feature extraction method. Moreover, as any parametric function, it returns the same number of parameters independently of the number of observed epochs, which is crucial for dealing with an inhomogeneous time-series that changes on a daily basis (the importance of this homogeneous representation is further detailed in Section 5). We stress that a more flexible feature extraction procedure still holding the characteristics described above would only improve the results presented here.

3.2 Classifier

Once the data have been homogenized, we need a supervised learning model to harvest the information stored in the spectroscopic sample. Analogous to the feature extraction case, the choice of classifier also impacts the final classification results for a given static data set (Lochner et al. 2016). In order to isolate the impact of AL in improving a given configuration of feature extraction and machine learning pipeline, we chose to restrict our analysis to a single classifier. A complete study on how different classifiers respond to the update in training provided by AL is out of the scope of this work, but is a crucial question to be answered in subsequent studies.

All the results we present below were obtained with a *random forest* algorithm (Breiman 2001).

Random forest is a popular machine learning algorithm known to achieve accurate results with minimal parameter tuning. It is an ensemble technique made up of multiple decision trees (Breiman et al. 1984), constructed over different sub-samples of the original data. Final results are obtained by averaging over all trees (for further details, see appendices A and B of Richards et al. 2012a). The method has been successfully used for SN photometric classification (Richards et al. 2012a; Lochner et al. 2016; Revsbech et al. 2018). In what follows, we used the SCIKIT-LEARN¹¹ implementation of the algorithm with 1000 trees. In this context, the probability of being an SN Ia, p_{Ia} , is given by the percentage of trees in the ensemble voting for an SN Ia classification.¹²

3.3 Metrics

The choice of a metric to quantify classification success goes beyond the use of classical accuracy (equation 2) – especially when the populations are unbalanced (Fig. 3). In order to optimize information extraction, this choice must take into account the scientific question at hand.

In the traditional SN case, the goal is to improve the quality of the final SN Ia sample for further cosmological use. In this context, a false negative (an SN Ia wrongly classified as non-Ia) will be excluded from further analysis posing no damage on subsequent scientific results. On the other hand, a false positive (non-Ia wrongly classified as a Ia) will be mistaken by a standard candle, biasing the cosmological analysis. Thus, purity (equation 4) of the photometrically classified SN Ia set is of paramount importance. At the same time, we wish to identify as many SNe Ia as possible (high efficiency – equation 3) in order to guarantee optimal exploitation of our resources. Taking such constraints into consideration, Kessler et al. (2010) proposed the use of a figure of merit that penalizes classifiers for false positives (equation 5). Throughout our analysis, classification results will be reported according to these four metrics:

$$\text{accuracy} = \frac{N_{\text{sc}}}{N_{\text{tot}}}, \quad (2)$$

$$\text{efficiency} = \frac{N_{\text{sc,Ia}}}{N_{\text{tot,Ia}}}, \quad (3)$$

$$\text{purity} = \frac{N_{\text{sc,Ia}}}{N_{\text{sc,Ia}} + N_{\text{wc,nIa}}}, \quad \text{and} \quad (4)$$

$$\text{figure of merit} = \frac{N_{\text{sc,Ia}}}{N_{\text{tot,Ia}}} \times \frac{N_{\text{sc,Ia}}}{N_{\text{sc,Ia}} + W N_{\text{sc,nIa}}}, \quad (5)$$

where N_{sc} is the total number of successful classifications, N_{tot} is the total number of objects in the target sample, $N_{\text{sc,Ia}}$ is the number of successfully classified SNe Ia (true positives), $N_{\text{tot,Ia}}$ is the total number of SNe Ia in the target sample, $N_{\text{wc,nIa}}$ is the number of non-Ia SNe wrongly classified as SNe Ia (false positives), and W is a factor that penalizes the occurrence of false positives. Following Kessler et al. (2010), we always use $W = 3$.

¹¹<http://scikit-learn.org/>

¹²In this work, we are concerned only with Ia $\times$ non-Ia classification. The analysis of classification performance using other SN types will be the subject of a subsequent investigation.

In the AL framework we propose, the metrics above are used to quantify the classification results in the target sample. They were calculated after the classifications were performed and had no part in the decision making algorithm (further details in Section 4 and in the Appendix).

4 ACTIVE LEARNING

We now turn to the key missing ingredient in our pipeline. The tools we have described thus far allow us to process (Section 3.1), classify (Section 3.2), and evaluate classification results (Section 3.3), given a pair of labelled and unlabelled light-curve data sets. The question now is starting from this initial configuration, how can we optimize the use of subsequent spectroscopic resources in order to maximize the potential of our classifier? Or in other words, how can we achieve high generalization performance by adding a minimum number of new spectroscopically observed objects to the training sample? We advocate the use of a dedicated recommendation system tuned to choose the most informative objects in a given sample – those which will be spectroscopically targeted.

AL is an area of machine learning designed to optimize learning results while minimizing the number of required labelled instances. At each iteration, the algorithm suggests which of the unlabelled objects would be most informative to the learning model if its label was available (Cohn et al. 1996; Balcan et al. 2009; Settles 2012). Once identified, a query is made – in other words, the algorithm is allowed to interact with an external agent in order to ask for the label of that object.¹³ The queried object – along with its label – is then added to the training sample and the model is re-trained. This process is repeated until convergence is achieved, or until labelling resources are exhausted. The complete algorithm is illustrated in Fig. 5.

Different flavours of AL propose different strategies to identify which objects should be queried. In what follows, we report results obtained using pool-based AL¹⁴ via uncertainty sampling. In this framework, we start with two sets of data: labelled and unlabelled. At first, the machine learning algorithm (classifier) is trained using all the available labelled data. Then, it is used to provide a class probability for all objects in the unlabelled set. The object whose classification is most uncertain is chosen to be queried. In a binary classification problem as the one investigated here, this corresponds to querying objects near the boundary between the two classes – where the classifier is less reliable. A detailed explanation of the uncertainty sampling technique (as well as *query by committee*) is given in the Appendix.

Finally, whenever one wishes to quantify the improvement in classification results due to AL, it is important to keep in mind that simply increasing the number of objects available for training changes the state of the model – independently of how the extra data were chosen. Thus, AL results should always be compared to the *passive learning* strategy, where at each iteration objects to be queried are randomly drawn from the unlabelled sample. Moreover, in the specific case of SN classification, we also want to investigate how the learning model would behave if the same number of ob-

¹³In the machine learning jargon, this agent is called an *oracle*. It can be a human, machine, or piece of software capable of providing labels. In our context, making a query to the oracle corresponds to spectroscopically determining the class of a given object.

¹⁴Our analysis used the LIBACT PYTHON package developed by Yang et al. (2017).

Figure 5. Schematic illustration of the AL work flow in the context of photometric light curves classification. Starting at the top left, the training set (spectroscopic sample - grey circles), is used to train a machine learning algorithm – resulting in a model which is then applied to the unlabelled data (photometric light curves – yellow pentagons). This initial model returns a classification for each data point of the unlabelled set (now represented as red squares and blue triangles). The AL algorithm is then used to choose a data point of the unlabelled data with highest potential to improve the classification model (identified by the grey arrow). The label of this point is then queried (a spectrum is taken). Once the true label of the queried point is known, it is added to the training set (converted into a grey circle), and the process is repeated.

jects were queried following the *canonical* spectroscopic targeting strategy, where objects to be queried are randomly drawn from a sample that follows closely the initial *SNPCC* labelled distribution. Diagnostic results presented below show outcomes from all three strategies: canonical, passive learning, and AL via uncertainty sampling.¹⁵

4.1 Static full light-curve analysis

We begin by applying the complete framework described in the previous subsections to static data. This is the traditional approach, where we consider that all light curves were completely observed at the start of the analysis. Although this is not a realistic scenario (one cannot query, or spectroscopically observe, an SN that has already faded away), it gives us an upper limit on estimated classification results. Section 5 considers more realistic constraints on available query data and light-curve evolution.

For each light curve and filter, all available data points were used to find the best-fitting parameters of equation (1) following the procedure described in Section 3. Best-fitting values for different filters were concatenated to compose a line in the data matrix. In

order to ensure the quality of fit, we considered only SNe with a minimum of five observed epochs in each filter; this reduced the size of our spectroscopic and photometric samples to 1094 and 20 193 objects, respectively.

4.1.1 Sub-samples

The iterative framework presented above corresponds to the AL strategy for choosing the next object to be queried. In this description, we have two samples: labelled and unlabelled. In case, we wish to quantify the performance of the ML algorithm after each iteration, the recently re-trained model must be used to predict the classes of objects in a third sample – one that did not take part in the AL algorithm. Classification metrics are then calculated, after each iteration, from predictions on this independent sample. In this scenario, we need three samples: training, query, and target. The *query sample* corresponds to the set of all objects *available for query* upon which the model evolves.¹⁶ On the other hand, the *target sample* corresponds to the independent one over which diagnostics are computed. In the traditional analysis, query and target sample follow the same underlying distribution in feature space; this separation helps avoid overfitting. This is the case for the static full light-curve analysis. In the results presented in this sub-section, *SNPCC photo* was randomly divided into query (80 per cent) and target (20 per cent) samples.

Finally, we quantified the evolution of the classification results when new objects are added to the training sample according to the canonical spectroscopic follow-up strategy, by constructing a *pseudo-training* sample. For each element of *SNPCC spec*, we searched for its nearest neighbour in *SNPCC photo*.¹⁷ This allowed us to construct a data set that follows very closely the distribution in the parameter space covered by the original *SNPCC spec*. Thus, randomly drawing elements to be queried from this *pseudo-training* sample is equivalent to feeding more data to the model according to the canonical spectroscopic follow-up strategy.

4.1.2 Results

In this section, we present classification results for the static full light-curve scenario according to three spectroscopic targeting strategies: canonical, passive learning, and AL via uncertainty sampling (Section 4). In all three cases, at each iteration one object was queried and added to the training sample (the one with highest uncertainty). We allowed a total of 1000 queries, almost doubling the original training set.

Fig. 6 shows how classification diagnostics evolve with the number of queries. The red inverse triangles describe results following the canonical strategy (random sampling from the pseudo-training sample), yellow circles show results from passive learning (random sampling from the query sample), and blue triangles represent results for AL via uncertainty sampling. We notice that the canonical spectroscopic targeting strategy does not add new information to the model – even if more labelled data are made available. Thus, there is almost no change in diagnostic results after 1000 queries.

¹⁶Not to be confused with the *set of queried objects*, which comprises the specific objects added to the training set (1 per iteration).

¹⁷This calculation was performed in a 16 dimensions parameter space: type, redshift, simulated peak magnitude, and mean SNR in all four filters. For all the numerical features, we used a standard Euclidean distance.

¹⁵The code used in this paper can be found in the COINtoolbox – <https://github.com/COINtoolbox/ActSNClass>.

Figure 6. Evolution of classification results as a function of the number of queries for the static full light-curve analysis.

On the other hand, the canonical strategy is very successful in identifying SN Ia (approximately 92 per cent efficiency); however, by prioritizing bright events, it fails to provide the model with enough information about other SN types. Consequently, its performance in other diagnostics is poor (~ 60 per cent accuracy, 36 per cent purity and a figure of merit of 0.15). At the same time, passive learning and AL via uncertainty sampling show very similar efficiency results up to 400 queries. Accuracy levels stabilize quickly (84 per cent/87 per cent after only 200 queries), followed closely by purity results (73 per cent after 600 queries). The biggest difference appears on efficiency levels. We can recognize an initial drop in efficiency up to 400 queries. This is expected, since both strategies prioritize the inclusion of non-Ia objects in the training sample: passive learning simply led by the higher percentages of non-Ia SNe in the target sample (Fig. 3), and AL by aiming at a more diverse information pool. This implies that high accuracy and purity levels are accompanied by a decrease in efficiency (from 92 per cent to 68 per cent at 200 queries). After a minimally diverse sample is gathered, passive learning continues to lose efficiency, stabilizing at 63 per cent after 700 queries, while AL is able to harvest further information to stabilize at 72 per cent after 800 queries. Thus, after 1000 new objects were added to the training sample, passive learning achieves a figure of merit of 0.32 (2.1 times higher than canonical), while AL via uncertainty sampling achieves a figure of merit of 0.39 (2.6 times higher than canonical).

Fig. 7 illustrates how the distribution of peak i -band magnitude in the set of queried elements evolves with the number of queries. For the sake of comparison, we also show the static distributions for the training (yellow) and target (blue) samples. As expected, the canonical strategy consistently follows the spectroscopic sample distribution. Meanwhile, passive learning completely ignores

the existence of the initial training – consequently, its initial queries overlap with regions already covered by the training sample, allocating a significant fraction of spectroscopic resources to obtain information already available in the training. The AL strategy, even in very early stages, takes into account the existence of the training sample, focusing its queries in the region not covered by training data (higher magnitudes). At 900 queries, the set of queried objects chosen by passive learning (red line, middle column) follows closely the distribution found in the target sample (blue), - but this does not translate into a better classification because the bias present in the original training was not yet overcome. On the other hand, the discrepancy in distributions between the target sample (blue region) and the set of objects queried by AL (red line, right-most column) at 900 queries is a consequence of the existence of the initial training.¹⁸ The fact that AL takes this into account is reflected in the classification results (Fig. 6).

These results provide evidence that AL algorithms are able to improve SN photometric classification results over canonical spectroscopic follow-up strategies, or even passive learning in a highly idealized environment.¹⁹ However, in order to have a more realistic description of an SN survey, we need to take into account the transient nature of the SNe and the evolving aspect of an observational survey.

Although we chose to illustrate non-representativeness between samples in terms of peak brightness in different bands (e.g. Figs 1, 7, and 12), these features are absent in the input data matrix. Our

¹⁸The reader should keep in mind that after 1000 queries the model is trained in a sample containing the complete SNPCC spectroscopic sample added to the set of queried objects.

¹⁹A result already pointed out by Gupta et al. (2016).

Figure 7. Simulated *i*-band peak magnitude distribution as a function of the number of queries for the static full light-curve scenario. The yellow (blue) region shows distribution for the training (target) samples, while the red curves denote the composition of sample queried by AL. Lines go through 10–900 queries (from top to bottom). Different columns correspond to different learning strategies: canonical, passive, and AL via uncertainty sampling (from left to right).

goal is to emphasize that the underlying astrophysical properties are tracked differently by the AL and passive learning strategies – even if these are not explicitly used.

5 REAL-TIME ANALYSIS

In this section, we present an approach to deal with the time evolving aspect of spectroscopic follow-ups in SN surveys. This is done through the daily update of

- (i) identification of objects allocated to query and target samples,
- (ii) feature extraction, and
- (iii) model training.

We begin considering the full SNPCC spectroscopic sample completely observed at the beginning of the survey; this allows us to have an initial learning model. Then, at each observation day d , a given SN is included in the analysis if, until that moment, it has at least five observed epochs in each filter. If this first criterion is fulfilled, the object is designated as part of the *query sample* if its *r*-band magnitude is lower than or equal to 24 ($m_r \leq 24$ at d) – otherwise, it is assigned to the target sample.²⁰ Fig. 9 shows how

the number of objects in the query (yellow circles) and target (blue triangles) samples evolve as a function of the number of observing days. Although the survey starts observing at day 1, we need to wait until day 20 in order to have at least 1 object with a minimum of five observed epochs in each filter. From then on, the query sample begins with 666 objects (at day 20) and shows a steady increase until it almost stabilizes ~ 2100 SNe (around day 60). On the other hand, the target sample shows a steep increase until $d \sim 80$ (hereafter, *build-up phase*) and continuous to grow from there until the end of the survey, although at a lower rate. This behaviour is expected, since, in this description, the *query sample* corresponds to the fraction of photometric objects satisfying the magnitude threshold ($m_r \leq 24$) at a specific time. Notice that as the survey evolves, an object whose detection happened in a very early phase will be assigned to the target sample during its rising period, but if its brightness increases enough to allow spectroscopic targeting it will move to the query sample – where it will remain for a few epochs. After its maximum passes, the SN will eventually return to the target sample as soon as it fades below the magnitude threshold – remaining there until the end of the survey. Thus, it is important to keep in mind that, despite its size being practically constant after the build-up phase, individual objects composing the query sample might not be the same for consecutive days.

The feature extraction process is also performed on a daily basis, considering only the epochs measured until that day. This clarifies why we consider an analytical parametrization a simple, and

²⁰We consider an object with *r*-band magnitude of 24 to have the minimum brightness necessary to allow spectroscopic observation with a 8-m class telescope.

Figure 8. Evolution of the classification results as a function of the survey duration for the time-domain AL, considering the SNPCC training set as completely given in the beginning of the survey.

efficient enough, feature extraction procedure. It reasonably fast and encompasses prior domain knowledge on light curve behaviour while returning the same number of parameters independently of the number of observed epochs. Moreover, it avoids the necessity of determining the time of maximum brightness or performing any type of epoch alignment (see e.g. Richards et al. 2012a; Ishida & de Souza 2013; Revsbech et al. 2018). Thus, we are able to update the feature extraction step as soon as a new epoch is observed and still construct a homogeneous and complete low-dimensionality data matrix. The only constraint is the number of observed epochs, which must be at least equal to the number of parameters in all filters.

Finally, at the end of each night, the model is trained using the features and labels available until that point. The AL algorithm is allowed to query only the objects belonging to the *query sample*. Once a query is made, the targeted object and corresponding label are added to the training sample, the model is re-trained and the result applied to the target sample (fig. 5). Given the time span of the SNPCC data, we are able to repeat this analysis for a total of 180 d.

Fig. 8 shows the evolution of classification results considering the complete SNPCC spectroscopic sample as a starting point. Here, we can clearly see the effect of the evolving sample sizes: accuracy and efficiency results oscillate, while purity and figure of merit remain indifferent to the learning strategy, during the build-up phase (grey region). Once this phase is over, results start to differ and the AL with uncertainty sampling clearly surpasses the other two, achieving

80 per cent accuracy, 55 per cent purity and a figure of merit of 0.23, while the passive learning only goes up to 72 per cent accuracy, 45 per cent purity and figure of merit of 0.18. The canonical strategy continues to output better efficiency, but its loss in purity does not allow it to overcome even passive learning in figure of merit levels.

5.1 No initial training

This leaves one open question: what should we do at the beginning of a given survey, when a training set with the same instrument characteristics (e.g. photometric system) is not yet available? Or even more drastically: if the algorithm is capable of building its own training sample, do we even need an initial training at all? The answer is no.

Fig. 10 shows how the classification results behave when the initial model is trained in one randomly selected object from the query sample, meaning we start with a random classifier. In this context, diagnostics are meaningless until around 100 d (a little after the build-up phase) when all samples involved are under construction. After this period, AL with uncertainty sampling starts to dominate purity and, consequently, figure of merit results. After 150 observation days (or after 130 objects were added to the training), the active and passive learning strategies achieve purity levels comparable to the one obtained in the unrealistic full light-curve scenario (~80 per cent). Thus, at the end of the survey, AL efficiency results (27 per cent) are 80 per cent higher than the one obtained by passive learning (15 per cent), which translates into an almost doubled fig-

Figure 9. Number of objects in the query (yellow circles) and target (blue triangles) samples as a function of the days of survey duration. The grey region highlights the initial build-up phase of the survey, where there is a steep increase in the number of objects in the target sample.

ure of merit (0.14 from AL and 0.08 from the canonical strategy). Compare these results with the *initial state* of the full light-curve analysis: Fig. 6 (accuracy 60 per cent, efficiency 92 per cent, purity of 35 per cent and 0.15 figure of merit) was obtained using complete light curves for all objects, all SNe in the original SNPCC spectroscopic sample surviving the minimum number of epochs cuts (1094

objects) and the same random forest classifier. Final results of the real-time AL analysis (Fig. 10) surpasses the full light-curve initial state accuracy results in 33 per cent, more than doubles purity and achieves comparable figure of merit results. All of these while respecting the time evolution of observed epochs of only 161 SNe in the training set, or 15 per cent the number of objects in the original SNPCC spectroscopic sample.

Accuracy levels of real time AL with (Fig. 8) and without (Fig. 10) the full initial training sample are comparable, while efficiency and figure of merit are higher for the former case. However, purity levels are 45 per cent higher without using the initial training. This is a natural consequence of the higher number of SNe Ia in the SNPCC spectroscopic sample (Fig. 9), which requires the algorithm to unlearn the preference for Ia classifications before it can achieve its full potential in purity results. Fig. 10 also shows that regarding purity, passive learning is able to achieve the same results as those obtained with uncertainty sampling while efficiency is severely compromised – exactly the opposite behaviour shown by the canonical strategy. This is a consequence of the populations targeted by each of these strategies. By prioritizing brighter objects, the canonical strategy introduces a bias in the learning model towards SN Ia classifications. On the other hand, by randomly sampling from the target, passive learning adds a larger number of non-Ia examples to the training, introducing an opposite bias, at least in the early stages of the survey.

In summary, given the intrinsic bias present in all canonically obtained samples, we advocate that the best strategy for a new survey is to construct its own training during the first running seasons.

Figure 10. Evolution of classification results as a function of survey duration in the real-time analysis, with a random initial training of one object.

Letting its own photometric sample guide the decisions of spectroscopic targeting. This is specially important if one has the final goal of supernova cosmology in mind, where the main objective is to maximize purity (minimize false positives) as well as many other scientific SN objectives.

6 BATCH-MODE AL

In this section, we take another step towards a more realistic description of a spectroscopic follow-up scenario. Instead of choosing one SN at a time, spectroscopic follow-up resources for large-scale surveys will probably allow a number of SNe to be spectroscopically observed per night. Thus, we need a strategy that allows us to extend the AL algorithm, optimizing our choice from one to a set (or a batch) of objects at each iteration. We focus on two methods derived from the notion of uncertainty sampling: *N-least certain* and *Semisupervised uncertainty sampling*.

The *N-least certain* batch query strategy uses the same machinery described in the sequential uncertainty sampling method but, instead of choosing a single unlabelled example, it selects the N objects with highest uncertainties, and queries all of them. This tactic carries a disadvantage, since a set of objects whose predictions exhibit similar uncertainties will probably also be similar among themselves (i.e. will be close to each other in the feature space). Thus, querying for a set of labels is not likely to lead to a model much different than the one obtained by adding only the most uncertain object to the training set. In dealing with a batch mode scenario, we should also require that the elements of the batch be as diverse as possible (maximizing their distance in the feature space).

Semisupervised uncertainty sampling (e.g. Hoi et al. 2008), in contrast, avoids the need to call the oracle at each individual iteration by using the uncertainty associated to each predicted label as a proxy for class assignment. The algorithm must be trained in the available initial sample in order to create the first batch. The object with the greatest classification uncertainty is then identified. Suppose this object has a probability p of being SN Ia. A pseudo-label is then drawn from a Bernoulli distribution, where success is interpreted as ‘Ia’ label (with probability p) and failure as ‘non-Ia’ (with probability $1 - p$). The object features and corresponding pseudo-label are temporarily added to the training sample and the model is re-trained. This is repeated until we reach the size of the batch (see algorithm 1). The benefit of using the model to produce pseudo-labels comes with the inevitable uncertainty attached to model predictions: they come unwarranted. However, the problem attached to the *N-least certain* strategy is here, to a certain degree, overcome. Similar unlabelled instances are less likely to be included in the same batch.

The optimum number of elements in each batch, N , is highly dependent on the particular combination of data set and classifier at hand. At each iteration, we are actually stretching the capabilities of the learning model in a feedback loop that cannot be expected to perform well for large batches. For the SNPCC data, our tests show that semisupervised learning outperforms the *N-least certain* strategy for $N \in [2, 8]$ with maximum results obtained with $N = 5$.

Fig. 11 shows classification results for canonical and passive learning (both at each iteration drawing 5 random elements from the pseudo-training and target sample, respectively), AL via *N-least certain* and semisupervised uncertainty sampling, when the initial training consists of five randomly drawn objects from the query sample and $N = 5$. We see that in this scenario semisupervised AL is able to achieve the same figure of merit (~ 0.22) as sequential uncertainty sampling when the entire initial training sample is

Data: Training set, T_{tr} , Unlabelled set, T_u and batch size, N

Result: List of data points to be queried

```

R ← []
Batch ← []
for i=1,...,N do
    Use  $T_{tr}$  to train learning model  $\mathcal{M}$ 
    for  $\mathbf{x} \in T_u$  do
        ID, class, Ia_prob =  $\mathcal{M}(\mathbf{x})$ 
        R ← {ID, class, Ia_prob,  $\mathbf{x}$ }
    end
     $r^* \leftarrow$  element in R with largest uncertainty
    Batch ←  $r^*[1]$ 
     $d \sim \mathcal{B}(p = r^*[3])$ 
    if  $d \equiv \text{True}$  then
        |  $y \leftarrow r^*[2]$ 
    else
        |  $y \leftarrow$  alternate class ( $\neq r^*[2]$ )
    end
    Add element to training set  $T_{tr} \leftarrow (r^*[4], y)$ 
end
return Batch

```

Algorithm 1: Semi-supervised uncertainty sampling algorithm to identify which elements must be included in the batch.

available (Fig. 8). However, it does so using only 63 per cent of the number of objects for training (or 800 SNe in the training after 180 d, against 1263 SNe in the full training case). Moreover, although efficiency results show a steady increase until the end of the survey, purity achieves saturation levels (~ 0.8 – the same as the final results obtained with the static full light-curve scenario, Fig. 6) after only 100 d (corresponding to a training set with 405 objects). A numerical description of the final classification results and corresponding training size is shown in Table 1.

From Fig. 11, we see that samples containing the same number of objects lead to different classification results. Moreover, considering that the query sample only contains objects with $m_r \leq 24$, we should not expect the set of objects queried by AL to be representative of the target sample, despite the improvement in classification results driven by AL. This is clearly shown in Fig. 12, where we compare distribution of maximum observed brightness in each filter for the SNPCC spectroscopic (red) and photometric (blue) samples with the set of objects queried by AL (dark grey). The latter provides a slight advantage in coverage when compared to the original spectroscopic sample, but it is still significantly different from the photometric distribution. A similar behaviour is found when we compare the populations of different SN types (Fig. 13) and redshift distribution (Fig. 14). These results confirm that, although a slight adjustment is necessary in order to optimize the allocation of spectroscopic time, a significant improvement in classification results may be achieved without a fully representative sample.

7 TELESCOPE ALLOCATION TIME

As a final remark, we must address the question of how much spectroscopic telescope time is required to obtain the labels queried by the AL algorithm – in comparison to the time necessary to get all labels from SNPCC spectroscopic sample (canonical strategy).

In the realistic case of a survey adoption of the framework proposed here, a term taking into account the telescope time needed for spectroscopy observations must be added to the cost function of the AL algorithm. This was not explicitly taken into account

Figure 11. Evolution of classification results as a function of survey duration for the batch-mode real-time analysis with $N = 5$ and a random initial training of five objects.

Table 1. Classification results for the AL by uncertainty sampling (UNC) and semisupervised batch mode (BATCH 5) strategies.

	Static, full LC initial training UNC	Time domain initial training UNC	Time domain initial training BATCH 5	Time domain BATCH 5
Training size	2093	1255	1093	810
Accuracy	0.89	0.80	0.85	0.83
Efficiency	0.73	0.78	0.69	0.44
Purity	0.78	0.55	0.69	0.76
Figure of merit	0.39	0.23	0.31	0.22

in this paper, but we considered a constraint on magnitudes for the set of SNe available for spectroscopic follow-up ($r_{\text{mag}} \leq 24$). We were able to estimate the integration time required for each object to achieve a given SNR by considering its magnitude and typical values for statistical noise of the sources.²¹ In the SNPCC spectroscopic sample, we considered the spectra taken at maximum brightness. For the set of AL queried objects, we used the magnitude at the epoch in which the object was queried. Considering an SNR of 10 (more than enough to enable classification) the ratio between the total spectroscopic time needed to get the labels for the SNPCC spectroscopic sample and the set of objects queried by

semisupervised AL is 0.9992. This indicates that the set of objects queried by AL would require less than 2.9 s more time than the SNPCC spectroscopic sample to be observed at each hour. Also, if a more realistic estimation had been performed considering instrumental overheads, the set of objects queried by AL would have significant advantage, as it contains 26 per cent less objects than the SNPCC spectroscopic sample. This gives us the first indication that AL-like approaches are feasible alternatives to minimize instrumental usage and, at the same time, optimize scientific outcome of photometrically classified samples.

For the specific case studied here, the high purity values achieved in early stages of the batch-mode AL, accompanied by the steady increase in efficiency (Fig. 11) renders our final SN Ia sample optimally suited for photometric classification in cosmological analysis

²¹Namely, counts in the sky, $\approx 13.8 \text{ e}^-/\text{s}/\text{pix}$ and read-out noise, $\approx 8 \text{ e}^-$ (e.g. Bolte 2015).

Figure 12. Distributions of maximum observed brightness, in all DES filters, for the set of objects queried by AL via batch-mode semisupervised uncertainty sampling with $N = 5$ (dark grey). This is compared to distributions from SNPCC spectroscopic (red – top) and SNPCC photometric (blue – bottom) samples.

Figure 13. Populations of different supernova types in the original SNPCC spectroscopic and photometric samples, and in time domain batch mode ($N = 5$) semisupervised AL query sample after 180 d of observations. The composition of the SNPCC samples are the same as shown in Fig. 3. The AL query sample holds 390 (48 per cent) Ia, 122 (15 per cent) Ibc, and 298 (37 per cent) II.

Figure 14. Redshift distribution of the original SNPCC spectroscopic (red – dotted) and photometric (blue – dashed) samples, superimposed to the redshift distribution of the AL query set for the time domain semisupervised batch mode AL strategy without the use of an initial training (dark blue – full). In each observation night, the algorithm queried for five SNe. The distribution shows redshift for the query sample after 180 observation nights.

– albeit being smaller in number of objects and requiring almost the same amount of spectroscopic resources to be secured.

8 CONCLUSIONS

Wide-field sky rolling surveys will detect an unprecedented number of astronomical transients every night. However, the usefulness of these photometric data for cosmological analysis is conditioned on our ability to perform automatic and reliable light-curve classification using a very limited number of spectroscopic observations for validation. Traditional attempts to address this issue via supervised learning methods focus on a learning model that postulates a static, fully observed pair of spectroscopic (training) and photometric (target) samples. Such studies are paramount to assess the requirements and performance of different classifiers. Nevertheless, they fail to address fundamental aspects of astronomical data acquisition, which renders the problem ill suited for textbook machine learning algorithms. The most crucial of these issues is the non-representativeness between spectroscopic and photometric samples.

This mismatch has its origins in a follow-up strategy designed to maximize the number of spectroscopically confirmed SNe Ia for cosmology, resulting in a highly biased spectroscopic set – and a sub-optimal training sample. Given such data configuration, not even the most suitable classifier can be expected to achieve its full potential. In this work, we advocate that any attempt to improve SN photometric classification must include a detailed strategy for constructing a more representative training set, without ignoring the constraints intrinsic to the observational process.

Our proposed framework updates on a daily basis crucial steps of the SN photometric classification pipeline and uses AL to optimize the scientific outcome from machine learning algorithms. On each day, we consider only the set of available observed epochs and perform feature extraction via a parametric light-curve representation. The identification of SNe available for spectroscopic targeting (objects with $m_r \leq 24$ on that day) as a separate group from the full photometric sample is also updated daily. Finally, by using AL, we allow the algorithm itself to target those objects available for spectroscopic targeting that would maximally improve the learning model if added to the training set. Using the proposed semisupervised batch-mode AL, we designate an optimal set of new objects

to be spectroscopically observed on each night. Once the batch is identified, the model is re-trained and new spectroscopic targets are selected for the subsequent night. This method avoids the necessity of an initial training sample: it starts with a random classifier, allowing the algorithm to construct an optimal training sample from scratch, specifically adapted to the survey at hand. The framework was successfully applied to the simulated data released after the SNPCC (Kessler et al. 2010).

Our results show that the proposed framework is able to achieve high-purity levels, ~ 80 per cent, after only 100 observation days – which corresponds to a training set of only 400 objects. After 180 observation days, or 800 queries, we are able to reach a figure of merit of 0.22 (Fig. 11) – highly above the values we would have obtained by using the canonical strategy for the idealized full training and full light-curve analysis (which achieves 35 per cent purity and figure of merit of 0.14, red curve in Fig. 6).

After showing the classification improvements achieved with the AL algorithm, we examined the characteristics of the set of objects queried by AL. As expected, the $m_r \leq 24$ requirement does not allow this set to deviate strongly from the characteristics of the SNPCC spectroscopic sample (see Figs 12–14). However, these small differences translate into significant improvements in purity, and consequently, figure of merit results. This comes with no penalty whatsoever in the necessary observational time for spectroscopic follow-up. The ratio of integration times required to observe the complete set of queried objects and the SNPCC spectroscopic sample is close to unity. This is still additional evidence that AL is a viable strategy to optimize the distribution of spectroscopic resources.

In all our results, we observe that the canonical targeting strategy, which prioritizes the spectroscopic follow-up of bright events that resemble SNe Ia, has higher efficiency results (which means that this strategy is more successful in identifying a high number of SNe Ia in the target sample). However, as the diversity present in the training sample is very low, purity levels are always compromised. The canonical strategy is successful in targeting SNe Ia, but it is not optimal when the goal is to construct a training sample for machine learning classifiers.

In the present analysis, we restricted ourselves to the SNe photometric classification case of SNe Ia versus non-Ia. However, we stress that our methodology is easily adaptable to more general classifications that include other transients (e.g. Narayan et al. 2018). This exercise will be especially informative when applied to the data from the upcoming *Photometric LSST Astronomical Time-series Classification Challenge*²² (PLAsTiCC).

Finally, it is reasonable to expect that different combinations of classifiers and feature extraction methods will react differently to the iterative AL process. The same is valid for the AL algorithm itself. Moreover, there are other practical issues that we did not take into account, like the time delay necessary to treat the spectra and provide a classification or the fact that sometimes one spectra is not enough to get a classification. In summary, we recognize that each item in our pipeline can be refined, potentially leading to even more drastic improvements in classification results. The results we show here are the first evidence that an iterative learning process adapted to the specificities of astronomical observations can lead to significant optimization in the allocation of observational resources.

Adapting to the era of big data in astronomy will entail adapting machine learning techniques to the unique reality of astronomical

observations. This requires tackling fundamental issues that will always be present in astronomical data (e.g. the discrepancy between training and test samples and the time evolution of a transient survey). Astronomy has once again the opportunity to provide ground for developments in other research areas by providing unique data situations not commonly present in other scenarios, as long as a consistent interdisciplinary environment is available. We are convinced that the most exciting part of this endeavour is still to come.

ACKNOWLEDGEMENTS

This work was created during the 4th COIN Residence Program²³ (CRP#4), held in Clermont-Ferrand, France on August 2017, with support from Université Clermont-Auvergne and La Région Auvergne-Rhône-Alpes. This project is financially supported by CNRS as part of its MOMENTUM programme over the 2018–2020 period. EEOI thanks Michele Sasdelli for comments on the draft and Isobel Hook for useful discussions. AKM acknowledges the support from the Portuguese Fundação para a Ciência e a Tecnologia (FCT) through grants SFRH/BPD/74697/2010, from the Portuguese Strategic Programme UID/FIS/00099/2013 for CENTRA, the ESA contract AO/1-7836/14/NL/HB and Caltech Division of Physics, Mathematics and Astronomy for hosting a research leave during 2017–2018, when this paper was prepared. RSS thanks the support from NASA under the Astrophysics Theory Program Grant 14-ATP14-0007. RB acknowledges support from the National Science Foundation (NSF) award 1616974 and the NKFI NN 114560 grant of Hungary. BQ acknowledges financial support from CNPq-Brazil under the process number 205459/2014-5. AZV acknowledges financial support from CNPq. AM thanks partial support from NSF through grants AST-0909182, AST-1313422, AST-1413600, and AST-1518308. This work has made use of the computing facilities of the Laboratory of Astroinformatics (IAG/USP, NAT/Unicisul), whose purchase was made possible by the Brazilian agency FAPESP (grant 2009/54006-4) and the INCT-A. This work was partly supported by the Center for Advanced Computing and Data Systems (CACDS) and by the Texas Institute for Measurement, Evaluation, and Statistics (TIMES) at the University of Houston. This project has been supported by a Marie Skłodowska-Curie Innovative Training Network Fellowship of the European Commission’s Horizon 2020 Programme under contract number 675440 AMVA4NewPhysics.

The Cosmostatistics Initiative²⁴ (COIN) is a non-profit organization whose aim is to nourish the synergy between astrophysics, cosmology, statistics, and machine learning communities. This work benefited from the following collaborative platforms: Overleaf,²⁵ Github,²⁶ and Slack.²⁷

REFERENCES

- Balcan M. F., Beygelzimer A., Langford J., 2009, *J. Comput. Syst. Sci.*, 75, 78
- Bazin G. et al., 2009, *A&A*, 499, 653
- Betoule M. et al., 2014, *A&A*, 568, A22
- Bolte M., 2015, *Modern Observational Techniques*, <http://www.icolick.org/bolte/AY257/s.n.pdf>

²³<https://iaacoin.wixsite.com/crp2017>

²⁴<https://github.com/COINtoolbox>

²⁵<https://www.overleaf.com>

²⁶<https://github.com>

²⁷<https://slack.com/>

²²<https://plasticblog.wordpress.com/>

- Breiman L., 2001, *Mach. Learn.*, 45, 5
- Breiman L., Friedman J. H., Olshen R. A., Stone C. J., 1984, *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA
- Campbell H. et al., 2013, *ApJ*, 763, 88
- Charnock T., Moss A., 2017, *ApJ*, 837, L28
- Childress M. J. et al., 2017, *MNRAS*, 472, 273
- Cohn D. A., Ghahramani Z., Jordan M. I., 1996, *J. Artif. Intell. Res.*, 4, 129
- Conley A. et al., 2011, *ApJS*, 192, 1
- Cover T. M., Thomas J. A., 2006, *Elements of Information Theory* (Wiley Series in Telecommunications and Signal Processing). Wiley-Interscience, New York, NY
- Dai M., Kuhlmann S., Wang Y., Kovacs E., 2018, *MNRAS*, 477, 4142
- DeBarr D., Wechsler H., 2009, *Sixth Conference on Email and Anti-Spam*. Mountain View, California, p. 1
- Foley R. J., Mandel K., 2013, *ApJ*, 778, 167
- Gamow G., 1948, *Nature*, 162, 680
- Goobar A., Leibundgut B., 2011, *Ann. Rev. Nucl. Part. Sci.*, 61, 251
- Gupta K. D., Pampana R., Vilalta R., Ishida E. E. O., de Souza R. S., 2016, in Jin Y., Kollias S., eds, 2016 IEEE Symp. Ser., Computational Intelligence (SSCI), IEEE, Athens, Greece
- Hillebrandt W., Niemeyer J. C., 2000, *Ann. Rev. Astron. Astrophys.*, 38, 191
- Hlozek R. et al., 2012, *ApJ*, 752, 79
- Hoi S. C. H., Jin R., Zhu J., Lyu M. R., 2008, in Shapiro L., Ahuja N., Boyer K., Shah M., Syeda-Mahmood T., eds, 2008 IEEE Conf., Computer Vision and Pattern Recognition, IEEE, Anchorage, p. 1
- Hoyle B., Paech K., Rau M. M., Seitz S., Weller J., 2016, *MNRAS*, 458, 4498
- Ishida E. E. O., de Souza R. S., 2013, *MNRAS*, 430, 509
- Johnson B. D., Crotts A. P. S., 2006, *AJ*, 132, 756
- Jones D. O. et al., 2017, *ApJ*, 843, 6
- Karpenka N. V., Feroz F., Hobson M. P., 2013, *MNRAS*, 429, 1278
- Kessler R. et al., 2010, *PASP*, 122, 1415
- Kranjc J., Smailović J., Podpečan V., Grčar M., Žnidaršič M., Lavrač N., 2015, *Inf. Process. Manage.*, 51, 187
- Kuznetsova N. V., Connolly B. M., 2007, *ApJ*, 659, 530
- Liu Y., 2004, *J. Chem. Inf. Comput. Sci.*, 44, 1936
- Lochner M., McEwen J. D., Peiris H. V., Lahav O., Winter M. K., 2016, *ApJS*, 225, 31
- Madsen K., Nielsen H. B., Tingleff O., 2004, *Methods for Non-Linear Least Squares Problems*, 2nd edn., Technical University of Denmark, Lyngby, Denmark
- Masters D. et al., 2015, *ApJ*, 813, 53
- Möller A. et al., 2016, *J. Cosmol. Astropart. Phys.*, 12, 008
- Narayan G. et al., 2018, *ApJS*, 236, 9
- Naul B., Bloom J. S., Pérez F., van der Walt S., 2018, *Nature Astron.*, 2, 151
- Newling J. et al., 2011, *MNRAS*, 414, 1987
- Perlmutter S. et al., 1999, *ApJ*, 517, 565
- Perrett K. et al., 2010, *AJ*, 140, 518
- Phillips M. M., 1993, *ApJ*, 413, L105
- Planck Collaboration et al., 2016, *A&A*, 594, A1
- Poznanski D., Gal-Yam A., Maoz D., Filippenko A. V., Leonard D. C., Matheson T., 2002, *PASP*, 114, 833
- Poznanski D., Maoz D., Gal-Yam A., 2007, *AJ*, 134, 1285
- Revsbech E. A., Trotta R., van Dyk D. A., 2018, *MNRAS*, 473, 3969
- Richards J. W., Homrighausen D., Freeman P. E., Schafer C. M., Poznanski D., 2012a, *MNRAS*, 419, 1121
- Richards J. W. et al., 2012b, *ApJ*, 744, 192
- Riess A. G. et al., 1998, *AJ*, 116, 1009
- Rodney S. A., Tonry J. L., 2009, *ApJ*, 707, 1064
- Sako M. et al., 2008, *AJ*, 135, 348
- Settles B., 2012, *Active Learning*, Morgan & Claypool, San Rafael, CA
- Solorio T., Fuentes O., Terlevich R., Terlevich E., 2005, *MNRAS*, 363, 543
- Spergel D. N. et al., 2007, *ApJS*, 170, 377
- Sullivan M. et al., 2006, *AJ*, 131, 960
- Thompson C. A., Califf M. E., Mooney R. J., 1999, in Bratko I., Dzeroski S., eds, *Proc. 16th Int. Conf. on Machine Learning*, Bled, Slovenia, p. 406
- Tripp R., 1998, *A&A*, 331, 815
- Varughese M. M., von Sachs R., Stephanou M., Bassett B. A., 2015, *MNRAS*, 453, 2848
- Vilalta R., Ishida E. E. O., Beck R., Sutrisno R., de Souza R. S., Mahabal A., 2017, *IEEE Symp. Ser., Computational Intelligence (SSCI)*, IEEE, Honolulu
- Wang Y., Gjergo E., Kuhlmann S., 2015, *MNRAS*, 451, 1955
- Xia X., Protopapas P., Doshi-Velez F., 2016, in Venkatasubramanian S. C., Meira W., eds, 2016 SIAM Int. Conf. on Data Mining, Cost-Sensitive Batch Mode Active Learning: Designing Astronomical Observation by Optimizing Telescope Time and Telescope Choice, Soc. Ind. Appl. Math. Philadelphia, p. 477
- Yang Y.-Y., Lee S.-C., Chung Y.-A., Wu T.-E., Chen S.-A., Lin H.-T., 2017, preprint ([arXiv:1710.00379](https://arxiv.org/abs/1710.00379))

APPENDIX: ACTIVE LEARNING

We now give a more detailed description of the AL framework presented in Section 4. As mentioned above, AL is an area of study in machine learning that suggests which are the most informative samples to add to the training set to enhance the current classification model (Cohn et al. 1996; Balcan et al. 2009; Settles 2012). AL assumes two sample sets: T_{tr} and T_{u} , where T_{tr} is the standard training set made of pairs (feature vectors and class labels). The new set $T_{\text{u}} = \{(\mathbf{x}_i)\}_{i=1}^p$ is made of p unlabelled elements. After a model is built on T_{tr} , AL suggests which instances in T_{u} should be assigned a class label and incorporated into T_{tr} to build a new (refined) model. The task of finding the class label corresponding to an instance in T_{u} is called *sample query*. AL tries to minimize the cost of querying samples from T_{u} by focusing on objects with the highest potential to force a change in the current predictive model.

In our particular study, T_{tr} plays the role of spectroscopic data, where class labels are known. Photometric observations were divided into T_{u} (objects available for query, whose $r_{\text{mag}} \leq 24$) and T_{te} (the remaining non-labelled objects). Final accuracy is measured on test set T_{te} . The potential distributional discrepancy between spectroscopic and photometric samples is tackled by querying examples from the true target (photometric) distribution, whereas the spectroscopic sample simply serves to provide an initial model amenable to refinement. However, as we showed in Sections 5 and 6, the algorithm is effective even when this initial training set is non-informative.

There are many variants of AL. Some examples include *query synthesis* where the learner can query instances from any region of the input (i.e. feature) space $\mathcal{X}$; or *stream-based selective sampling*, where instances are sampled according to the underlying sample distribution. The most common approach to AL, and the one adopted in this study, is known as *pool-based learning* where we assume the existence of a data set T_{u} from which unlabelled instances can be queried. A key concept in active learning is that each query is associated with a cost; a trade-off then exists between improving the current model by adding more labelled instances, and minimizing the overall cost needed to acquire the corresponding class labels. We describe two instantiations of the pool-based sampling approach in the following subsections.

A1 Uncertainty sampling

One approach to (pool-based) AL is called *uncertainty sampling*. The main idea is to iteratively train a model $f(\mathbf{x}|\theta)$ on T_{tr} and choose the single instance $\mathbf{x}^*$ from T_{u} with the highest uncertainty on its class label. After querying the label y of $\mathbf{x}^*$, the algorithm incorporates $(\mathbf{x}^*, y)$ into training set T_{tr} , and a new model is induced. This

iterative process continues until we reach a user-defined maximum in the number of allowed queries.

A key component in uncertainty sampling is the criterion to select the instance $\mathbf{x}^*$ with highest uncertainty. Different criteria exist to measure the degree of uncertainty. In every case, it is commonly assumed that the learning algorithm can produce not only a class prediction from model $f(\mathbf{x}|\theta)$, but also a posterior probability $P(y|\mathbf{x}, \theta)$, which can be used to quantify the confidence in the prediction. As an illustration, one possible criterion is known as *maximum entropy*; here, $\mathbf{x}^*$ is selected based on the class with highest entropy:

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x}} - \sum_y P(y|\mathbf{x}, \theta) \log_2 P(y|\mathbf{x}, \theta), \quad (\text{A1})$$

where the right-hand side of the equation computes Shannon's entropy (Cover & Thomas 2006).

Another popular way to measure uncertainty is known as the *least-confident* approach, where instance $\mathbf{x}^*$ is selected if it minimizes the class posterior probability:

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} P(y|\mathbf{x}, \theta). \quad (\text{A2})$$

One more approach is to select $\mathbf{x}^*$ by looking for the *minimum margin*, defined as follows:

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} P(y'|\mathbf{x}, \theta) - P(y''|\mathbf{x}, \theta) \quad (\text{A3})$$

where y' and y'' are the first and second most likely predictions made by $f(\mathbf{x}|\theta)$.

For our results, we implemented pool-based least confident AL by using the fraction of votes among the trees when random forest acts as the classifier. The result is an estimation of $P(y'|\mathbf{x}, \theta)$, where y' is the predicted class, $\mathbf{x}$ are the set of best-fitting parameters of equation (1) in all four filters, and θ are the parameters of the random forest classifier.

A2 Query by committee

Within pool-based learning, another popular approach to AL is known as *query by committee* (QBC). The main idea is as follows. Instead of looking for the single instance with the lowest confidence on its prediction, we train a committee or ensemble of different models $\mathcal{C} = \{\theta_i\}$ and look for the single instance where most of the models *disagree* on their prediction. Intuitively, querying such an instance is critical in finding a better location for the final decision boundary. The degree of uncertainty is quantified here in terms of the degree of model disagreement. Specifically, let us refer to every different model according to its parameter set θ . In the previous section, we referred to $P(y|\mathbf{x}, \theta)$ as the class posterior probability of a model parametrized by θ . Under a committee of different models $\mathcal{C}$, we are now able to compute the average posterior probability of $\mathbf{x}$ (across models in $\mathcal{C}$):

$$\bar{P}_{\mathcal{C}}(y|\mathbf{x}) = \frac{1}{|\mathcal{C}|} \sum_{\theta} P(y|\mathbf{x}, \theta). \quad (\text{A4})$$

This average probability says something about the uncertainty of the prediction of $\mathbf{x}$, but this time in terms of an ensemble of models. The closer the value to 0.5,²⁸ the higher the uncertainty of the prediction. We can now look for the single instance that

maximizes the entropy of such average probability:

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x}} - \sum_y \bar{P}_{\mathcal{C}}(y|\mathbf{x}) \log \bar{P}_{\mathcal{C}}(y|\mathbf{x}). \quad (\text{A5})$$

The formula is known as *vote entropy* and is similar to the idea behind uncertainty sampling, but here the class posterior probability is computed through a consensus of committee models.

In our investigations, we applied three different flavours of QBC. In each case, we used vote entropy as the measure of disagreement. Specifically, let y represents the potential label of a data point $\mathbf{x}$, let $V(y)$ represents the number of votes received by that label, and let $|\mathcal{C}|$ be the number of models in the committee. Our (hard-vote) approximation of equation (A5) is defined as follows:

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x}} - \sum_y \frac{V(y)}{|\mathcal{C}|} \log_2 \frac{V(y)}{|\mathcal{C}|}. \quad (\text{A6})$$

For the actual software implementation of the different committee models, we used the publicly available *scikit-learn*²⁹ PYTHON library, with the default input parameters, unless explicitly specified otherwise.

In the following sections, we detail the two QBC flavours that we applied in our experiments.

A2.1 QBC 1

In our first implementation of the QBC approach, the committee consisted of a selection of the following five machine learning classifiers:

- (i) logistic regression,
- (ii) a random forest made up of 100 trees,
- (iii) a support vector machine classifier,
- (iv) a decision tree, and
- (v) and a k -nearest neighbour classifier with $k = 19$.

The models used by these classifiers are significantly different from one another; therefore, it is reasonable to expect their decision boundaries to differ as well. The rationale behind this approach is that the algorithm would query objects that are not only difficult to classify in a single model, but are challenging across all models, therefore are universally worthy of attention when targeting follow-up observations.

A2.2 QBC 2

In this particular implementation of QBC, we used a committee made up of 15 random forests, each comprising 100 trees. This choice ensures that all committee members have about the same classification performance, as they are indeed different randomized examples of the same model. However, the disadvantage compared to previous selections is that the population of queried objects may be biased by any potential inability of this specific model to fail to generalize well on a given region of the feature space.

A2.3 Results

Fig. A1 shows how QBC results evolve with the number of queries for the static full training and full light-curve scenario. The figure also shows results for the canonical, passive learning, and for AL

²⁸This is valid when considering binary classifications only.

²⁹<http://scikit-learn.org>

Figure A1. Evolution of classification results as a function of the number of queries for the full initial training and full light-curve analysis.

by uncertainty sampling in order to facilitate comparison (these are the same results shown in Fig. 6). We can see that QBC1 results are slightly better than passive learning for very early times – when the impact of non-representativeness is greater. However, the diversity within the committee soon becomes insufficient to tackle the small differences between objects – and consequently, results are very similar to the passive learning strategy. This can be explained by the fact that, although the elements of the committee are different from each other, they are very simple instances of their class. Increasing the complexity of the elements might provide improved results.

QBC2 results follow closely those obtained with the uncertainty sampling strategy concerning SNe Ia classification. This was expected, since both approaches utilize a random forest classifier, even though the decision process is somewhat different. In fact, the total number of trees used in each decision is 1500 for QBC2, and only 1000 for uncertainty sampling, which might explain the slight

but consistent advantage in accuracy (the total number of correct classifications) for QBC2. A detailed study of the results of active learning algorithms focusing on other SN types is outside the scope of this work, but will certainly be the subject of subsequent analysis.

Finally, although the tests we performed were not conclusive regarding the QBC designs, we advise further investigation of this type of algorithm – especially if new data become available (e.g. PLAsTiCC). This might provide interesting results in the presence of more complex committees and in investigations of other SN types.

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.