Hellinger Distance Trees for Imbalanced Streams

Lyon, R.J.; Brooke, J.M.; Knowles, Joshua; Stappers, B.W.

DOI:
[10.1109/ICPR.2014.344](https://doi.org/10.1109/ICPR.2014.344)

Document Version
Peer reviewed version

Citation for published version (Harvard):
Lyon, RJ, Brooke, JM, Knowles, J & Stappers, BW 2014, Hellinger Distance Trees for Imbalanced Streams. in *Proceedings 22nd International Conference on Pattern Recognition (ICPR), 2014*. Institute of Electrical and Electronics Engineers (IEEE), pp. 1969-1974, 22nd International Conference on Pattern Recognition, ICPR 2014, Stockholm, Sweden, 24/08/14. <https://doi.org/10.1109/ICPR.2014.344>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

(c) 2014 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other users, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works for resale or redistribution to servers or lists, or reuse of any copyrighted components of this work in other works.

Checked Jan 2016

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Hellinger Distance Trees for Imbalanced Streams

R. J. Lyon, J. M. Brooke, J. D. Knowles
School of Computer Science
University of Manchester
Manchester, UK

B. W. Stappers
Jodrell Bank Centre for Astrophysics
University of Manchester
Manchester, UK

Abstract—Classifiers trained on data sets possessing an imbalanced class distribution are known to exhibit poor generalisation performance. This is known as the imbalanced learning problem. The problem becomes particularly acute when we consider incremental classifiers operating on imbalanced data streams, especially when the learning objective is rare class identification. As accuracy may provide a misleading impression of performance on imbalanced data, existing stream classifiers based on accuracy can suffer poor minority class performance on imbalanced streams, with the result being low minority class recall rates. In this paper we address this deficiency by proposing the use of the Hellinger distance measure, as a very fast decision tree split criterion. We demonstrate that by using Hellinger a statistically significant improvement in recall rates on imbalanced data streams can be achieved, with an acceptable increase in the false positive rate.

I. INTRODUCTION

Imbalanced data sets are characterised by a skewed class distribution, which typically favours a single majority class, over one or more minority classes. Nearly all datasets are imbalanced to some degree, however the extent to which they are imbalanced varies greatly. Those data sets which exhibit the largest imbalances present significant problems for inductive learners, since algorithms trained on such data often exhibit poor generalisation capabilities on test data. The cause of such poor performance has been the subject of extensive study by the research community since the early 2000's [1], [2], [3]. The problem is now fairly well understood, and a number of useful approaches have been developed that generally seem to improve classifier performance on imbalanced datasets (see related work in Section VII). However these approaches are intended for use on static datasets. These are data sets of a known fixed size n . In situations where the size of the dataset is indeterminable as with data streams, these approaches become unsuitable, either due to their reliance on knowing n a priori, or their higher computational requirements. In this paper we combine the Hellinger distance used previously in some static imbalanced learners [4] with a stream classifier based on the Hoeffding bound. Together these can be used to produce a skew insensitive decision tree split criterion that substantially increases minority class recall rates. Our goal is to use this approach to extract rare class instances from a heavily imbalanced data stream arising from the pulsar astronomy domain.

II. MOTIVATION

The motivation for this work is provided by the development of the world's largest radio telescope, the Square Kilometre Array (SKA). The SKA will comprise thousands of individual 15 metre receiving dishes, which together form

an instrument capable of producing data at a rate of many TB/s [5]. When used to survey the sky for periodic signals produced by a rare type of stellar object, the radio pulsar, such a high data capture rate presents a number of significant computational challenges and barriers to discovery. Our principal dilemma centres on determining how to separate the small number of scientifically interesting signal detections made by the SKA, from the large volume of negative detections arising from radio frequency interference (RFI) and background noise. The data to be processed is an intrinsically imbalanced collection of overwhelmingly negative detections requiring separation, which maps neatly to a traditional machine learning classification problem. However since the SKA's high data capture rate makes it infeasible for all data to be stored permanently (due to financial restrictions), the classification of received data will likely have to be performed on arrival, in close to real-time. The continual arrival of new signals for processing is analogous to a stream of data flowing from the telescope, through a computational processing pipeline. Thus we view the identification of rare signals in SKA data as a stream classification problem.

III. PROBLEM DEFINITION

We are concerned with the identification of rare class instances within a data stream emerging from a radio telescope's data processing pipeline, which possesses a class distribution of approximately $+1:-10,000^1$. Each instance in the stream is a pulsar candidate, a summary description of each signal exhibiting specific characteristics of interest. Our aim is to build a binary classifier capable of separating candidates arising from noise or RFI (the negative majority class), from those generated by radio pulsars (the positive minority class). We define the input stream as $S = \{\dots, (x_i), \dots\}$, $i = 1, \dots, \infty$ which describes the candidates emerging from a pulsar search pipeline under a discrete time model. Each candidate in the stream $x_i \in X$ is defined as $x_i = \{(x_i^1), \dots, (x_i^m)\}$, where each $x_i^j \in \mathbb{R}$ for $j = 1, \dots, m$ is a single summary statistic that describes some aspect of the signal represented by the candidate x_i . We specify binary candidate labels $C = \{-1, 1\}$, where c_i is an individual label such that $c_i \in C$ (i.e. uninteresting = -1, interesting = 1). Our goal is to learn a function $f : X \mapsto C$ which maps each candidate to its correct label producing the set of labelled candidates. From these the set of positively labelled candidates can be obtained, which should be stored for expert inspection. Building a classifier to achieve this on a static data set is straightforward. We need only form a suitable training set and apply the classification model learned

¹Estimation based on data obtained during a past pulsar survey [6].

to a test set. In a streamed classification scenario there is no distinction between training and test sets, since all data resides in the stream. There are instead two general processing models used for learning.

- **Batch processing model:** at time step i a batch b of n unlabelled instances arrives, and is classified using some model trained on batches b_1 to b_{i-1} . At time $i+1$ labels arrive for batch b_i , along with a new batch of unlabelled instances b_{i+1} to be classified.
- **Incremental processing model:** a single data instance arrives at time step i defined as x_i , and is classified using some model trained on instances x_1 to x_{i-1} . At time $i+1$ a label arrives for x_i , along with a new unlabelled instance x_{i+1} to be classified.

We adopt the incremental processing model throughout this work since it neatly describes the arrival of each candidate. However the model must be adapted since the stream S as described previously is completely unlabelled, given it is impossible to know a priori the correct label for each candidate². To accommodate this we prepend labelled training examples³ to the start of the stream, ensuring a classification model is built from instances x_1 to x_{i-1} . We refer to this process here as ‘off-line pre-training’. Furthermore, true class labels cannot be obtained for every instance at $i+1$ due to the candidate volume and the cost of analysis. In reality only instances receiving positive predictions will be labelled following expert analysis. This is simulated in the experiments described in Section VI, by assigning a varying proportion of instances in the stream their correct label at $i+1$. Thus we now consider the stream $S = \{\dots, (x_i, c_i), \dots\}$ where c_i is a class label. For labelled instances in the stream originating from simulated expert analysis, c_i will assume the value of a class label $\in C$, whilst for those without labels $c_i = 0$ which our learning algorithms will ignore during training.

IV. IMBALANCED LEARNING ON DATA STREAMS

Training a classifier on an imbalanced dataset does not necessarily mean poor generalisation performance [7]. If the training data are discriminative enough to separate the different classes in data space, then the classifier will perform well regardless of the imbalance. A contrived example of this is shown in Fig.1. It demonstrates that the underlying data distribution alone is not the root cause of poor classification performance on imbalanced data [7]. Rather it is three characteristics often possessed by imbalanced data sets that make it difficult for a classifier to separate the minority and majority classes. These are *small sample size* [8], [7], *class inseparability* [9], [7] Fig.2(c), and *small disjuncts* in Fig.2(d). Ultimately these characteristics conspire to make it difficult for a classifier to construct an optimal decision boundary leading to sub-optimal classifier performance. Furthermore learning from non-static imbalanced data, is known to be a difficult and important problem yet to be overcome in data-mining research [10]. Whilst classifiers built upon imbalanced streams have strong accuracy overall, they possess poor performance on the

Fig. 1: Training distribution (a) and test distribution (b). Here a skewed class distribution has no effect on classifier performance since a decision boundary can be induced which allows all instances to be classified correctly. Based on diagrams used in [8].

Fig. 2: Class inseparability in (c) and small disjuncts in (d). Based on diagrams used in [8].

minority class [11], [12]. Consider as an example a two-class stream with an arbitrarily large imbalance i.e. $\lim_{+} \rightarrow 1$ and $\lim_{-} \rightarrow \infty$. At any moment such a stream will possess only a small number of minority class instances. Thus a learner retrained over a random batch will do so on an intrinsically small sample of minority class instances, whilst an incremental learner will rarely be presented with a minority class example to learn from. This is the streamed version of the small sample problem. As the inductive bias of many machine learning algorithms is geared towards forming the simplest hypothesis possible, its consequences are severe: algorithms will typically predict the majority class for all instances [13], particularly as it is easy to dismiss any positives seen as noise. In real world domains the small sample problem is worsened as streams are often either unlabelled or partially labelled. Any partial labelling of the stream will likely include only the majority class instances, and this becomes increasingly likely as the proportion of labelled instances decreases. Those minority class instances which would provide a more representative sample, are rarely, if ever labelled. Real world classification tasks are also non-trivial, i.e. there is some degree of class overlap. Even if these overlaps are caused by noise and are few in number, say $\frac{1}{1000}$ instances, if the minority class occurs infrequently enough in the stream i.e. $\frac{1}{10000}$, then in time this noise would drown out the minority concept leading to class inseparability. Indeed the noisy negative examples may quickly form new clusters in data space for which a classifier would need to develop small disjuncts to cover. In order to alleviate these problems we either need to provide a larger and more representative sample of minority class examples for learning, or alternatively, a new approach unaffected by class skew. In this work we have adopted the latter approach using the Hellinger distance measure as a skew insensitive split criterion for streamed decision trees. Note that we have not considered the effects of distributional drift (concept drift) here; we reserve such discussions for future work.

²Accurate labels can only be obtained via re-observation, which is financially costly and time consuming.

³There are some 2,000 pulsars known to science for which training candidates can be generated.

V. HELLINGER DISTANCE

The Hellinger distance is a symmetric and non-negative measure of distributional divergence, related to the Bhattacharyya coefficient (BC)⁴ and the Kullback-Leibler (KL) divergence⁵. Cieslak and Chawla proposed the use of the Hellinger distance as a decision tree split criterion for imbalanced data in [4], [14]. They were able to utilise the distance in this way by considering two distributions P and N , to be the normalised frequencies of feature values in a binary classification scenario. In which case when $P = N$ there is no difference between the distributions therefore the distance is zero, whilst if the two are completely disjoint then the distance will be one. Cieslak and Chawla then use the notion of ‘affinity’ between P and N as a decision tree split criterion. The aim is to split tree nodes on those features with minimal affinity i.e. maximal Hellinger distance. This approach is appealing since it enables the splitting of features based on how well they discriminate between the examples seen so far in the stream, rather than on the feature which describes the largest possible number of instances seen so far (as with information gain). Intuitively Hellinger is skew insensitive, since an abundance of examples of one class will only serve to make its sample distribution more representative of its real distribution. In which case if a feature is a good class discriminator, then irrespective of the balance it will remain as such. Theoretical evidence backing up this intuitive interpretation is supplied in [4], [14]. The Hellinger distance with respect to a single continuous feature for the two class case given in [4], where the feature is discretized into b bins is defined as follows,

$$d_H(X_+, X_-) = \sqrt{\sum_{j=1}^b \left(\sqrt{\frac{|X_{+j}|}{|X_+|}} - \sqrt{\frac{|X_{-j}|}{|X_-|}} \right)^2}, \quad (1)$$

where X_+ is the positive class count, X_- the negative class count, and X_{+j} the count of positives in bin j (defined similarly for X_{-j}). The problem with this formulation is that it requires knowledge of the normalised frequencies of values for each feature. This requires $O(lfb)$ memory for l leaves in the tree, f features, b bins and c classes. For some stream environments this could be too costly. However the Hellinger distance between two normal distributions P and N can be computed more straightforwardly, provided the mean and standard deviation of the feature is known. Given P with mean μ_1 , variance σ_1^2 and standard deviation σ_1 with N defined similarly; we can calculate the Hellinger distance as follows,

$$d_H(P, N) = \sqrt{1 - \sqrt{\frac{2\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2}} e^{-\frac{1}{4} \frac{(\mu_1 - \mu_2)^2}{\sigma_1^2 + \sigma_2^2}}}. \quad (2)$$

The advantage of Eq. 2 is that it requires we keep only simple statistics which describe the distribution of each feature requiring $O(lf \cdot 2c)$ memory. Thus we propose a new algorithm-level approach (as opposed to other approaches described in Section VII) which uses the Hellinger distance defined between two normal distributions, as a splitting criterion in

Hoeffding trees (the Very Fast Decision Tree developed in [15]). We call this the Gaussian Hellinger VFDT. Since this approach relies upon measuring distances between normal distributions, it will require a number of examples to be seen before a reasonable normal model can be computed for each feature. Consequently it does not initially possess the any-time property⁶. Furthermore we are also assuming that features will be normally distributed. Pseudo code for our implementation of the GH-VFDT is presented in Algorithm 1, based on the VFDT algorithm given in [15]. Note that in our implementation of Eq. 2 continuous attributes were discretized into 10 bins. By way of comparison we also developed a streamed version

Algorithm 1 Gaussian Hellinger Very Fast Decision Tree (GH-VFDT)

Require: An input dataset $S = \{\dots, (x_i, c_i), \dots\}$, such that each x_i is an instance and c_i is its true class label. Also requires δ the confidence desired, and τ a parameter which prevents ties.

```

1: procedure GH-VFDT( $S, \delta, \tau$ )
2:   Let  $DT$  be a decision tree with leaf  $l_1$ 
3:   for  $i \leftarrow 1$  to  $|S|$  do ▷ For each stream instance.
4:      $l \leftarrow \text{sort}(x_i, c_i)$  ▷ Sort instance  $x_i$  to leaf  $l$ .
5:      $k \leftarrow c_i$  ▷ Get class.
6:     for  $j \leftarrow 1$  to  $|x_i|$  do ▷ For each feature.
7:       update  $\mu_{ijk}(l)$ 
8:       update  $\sigma_{ijk}(l)$ 
9:   Label  $l$  with majority class of instances seen at  $l$ 
10:  if all examples seen at  $l$  don't belong to same class then
11:     $F_a \leftarrow \text{null}$  ▷ Best feature.
12:     $F_b \leftarrow \text{null}$  ▷ 2nd best feature.
13:    for  $j \leftarrow 1$  to  $|x_i|$  do ▷ For each feature.
14:       $\text{dist} \leftarrow d_H(x_i)$  ▷ From equation 2.
15:       $F_a, F_b \leftarrow \text{getBest}(\text{dist}, x_i)$ 
16:     $\epsilon = \sqrt{\frac{R^2 \ln(1/\delta)}{2n}}$  ▷ Hoeffding bound.
17:    if  $d_H(F_a) - d_H(F_b) > \epsilon$  or  $\epsilon < \tau$  then
18:      Replace  $l$  with new leaf that splits on  $F_a$ 
19:      for each branch of split do
20:        Add new leaf  $l_m$ 
21:        for  $k \leftarrow 1$  to  $|C|$  do ▷ For each class.
22:          for  $j \leftarrow 1$  to  $|x_i|$  do ▷ For each feature.
23:             $\mu_{ijk}(l_m) \leftarrow 0$ 
24:             $\sigma_{ijk}(l_m) \leftarrow 0$ 
25:  return  $DT$ 

```

of the HDTree [4] which we call HD-VFDT. This is the same algorithm as GH-VFDT, except Eq.2 is replaced by Eq.1.

VI. EXPERIMENTS

A. Static Tests

Static experiments attempted to reproduce the results for the static domain as described in [4]. We applied 5 × 2-fold cross validation (CV)⁷ in conjunction with the C4.4⁸, C4.5, and the static HDTree [4] algorithm which uses Eq.1; to four data sets used in [4] (indicated in Table I with an asterisk). To evaluate our results we used metrics that give a representative impression of classifier performance on the minority class (e.g. the G-Mean = $\sqrt{\frac{TP}{TP+FN} \times \frac{TN}{TN+FP}}$). The area under the

⁶Although this can be mitigated by a suitable period of ‘pre-training’.

⁷Data split into two chunks, the first used for training the second for testing. This is then reversed (train on second, test on first) and repeated five times. 10-fold CV not used as it can give an elevated type 1 error rate [4], [16].

⁸An unpruned decision tree using Laplace smoothing at the leaves.

⁴Though Hellinger obeys the triangle inequality whereas the BC does not.

⁵The squared Hellinger distance lower bounds the KL divergence.

Dataset	Instances	Attributes / Type	$\sim$ Balance	Distribution	
				+	-
Pulsar	11,219,848	22 / Continuous	+1 : -7000	1,611	11,218,237
Skin	245,057	3 / Discrete	+1 : -4	50,859	194,198
MiniBoone	130,065	50 / Continuous	+1 : -3	36,499	93,565
Letter*	20,000	16 / Continuous	+1 : -4	3,878	16,122
Magic	19,020	10 / Continuous	+1 : -2	6,688	12,332
Pen*	10,992	16 / Continuous	+1 : -9	1,055	9,937
Statlog Landsat*	6,435	36 / Continuous	+1 : -9	626	5,809
Statlog Image*	2,310	19 / Continuous	+1 : -6	330	1,980

TABLE I: Characteristics of the data sets used. Asterisks indicate data sets used in [4].

ROC curve (AUC), F-score and recall rate were also used. This test procedure was then repeated on larger datasets not used in [4]. The purpose of running these experiments was to determine a baseline level of performance for static classifiers, against which stream performance could be compared.

B. Stream Tests

Data stream experiments were performed on a VFDT, the GH-VFDT presented in Algorithm 1, and a streamed implementation of the HDTree called HD-VFDT also developed as part of this work. In this streamed scenario, the three largest data sets in Table I were shuffled (except for the temporally ordered pulsar data) and then randomly sampled in order to generate: 1) training sets containing 200 positive and 1000 negative instances to be used for pre-training, and 2) disjoint test sets possessing different levels of class imbalance and a varied proportion of labelled instances, simulating different levels of expert feedback. Each sampling permutation was treated as a single stream of data. For each algorithm, tests were repeated ten times for a given balance and labelling, allowing results to be averaged. An incremental test-then-train approach was adopted throughout the stream experiments, i.e. test on an instance, then train if the label is available. Results were evaluated predominately using the G-Mean. The popular AUC metric was not computed given that it is computationally expensive to do so for millions of instances⁹. To rigorously validate our stream results, one-factor analysis of variance (ANOVA) tests were performed, where the algorithm used was the factor. Tukey’s Honestly Significant Difference (HSD) test was then applied to determine which results were statistically significant using $\alpha = 0.01$.

C. Data

In total eight datasets were used for this work (see Table I). The largest dataset consisted of pulsar candidates obtained during the HTRU survey [19]¹⁰. The remaining datasets were all obtained from the UCI machine learning repository. These include the Skin Segmentation, MiniBoone and the Magic Gamma Telescope datasets, describing binary classification problems. The Letter, Pen, Statlog Landsat and Statlog Image datasets on the other hand, describe multi-class problems. These were converted to binary datasets by labelling the smallest class within them as the minority class, and labelling all other instances as belonging to the majority class. For the Letter data set this meant grouping all vowel instances (without

⁹Whilst methods have been developed to provide AUC approximations for large data sets [17], we do not make use of these here. Particularly as work in [18] suggests the AUC is an unsuitable metric for use upon imbalanced data.

¹⁰This data is currently not publicly accessible.

Dataset	Algorithm	AUC	G-Mean	F-Score	Recall
Pulsar	C4.4	.976 $\pm$.03	.796 $\pm$.03	.628 $\pm$.09	.634 $\pm$.14
	C4.5	.867 $\pm$.09	.790 $\pm$.07	.635 $\pm$.06	.624 $\pm$.11
	HDTree	.992 $\pm$.01	.734 $\pm$.04	.555 $\pm$.05	.540 $\pm$.06
Skin	C4.4	.999 $\pm$.01	.999 $\pm$.01	.998 $\pm$.01	.999 $\pm$.01
	C4.5	.999 $\pm$.01	.999 $\pm$.01	.998 $\pm$.01	.999 $\pm$.01
	HDTree	.999 $\pm$.01	.999 $\pm$.01	.998 $\pm$.01	.998 $\pm$.01
MiniBoone	C4.4	.999 $\pm$.01	.999 $\pm$.01	.999 $\pm$.01	.999 $\pm$.01
	C4.5	.999 $\pm$.01	.999 $\pm$.01	.999 $\pm$.01	.999 $\pm$.01
	HDTree	.999 $\pm$.01	.999 $\pm$.01	.999 $\pm$.01	.999 $\pm$.01
Letter	C4.4	.965 $\pm$.01	.913 $\pm$.03	.862 $\pm$.03	.861 $\pm$.05
	C4.5	.933 $\pm$.02	.906 $\pm$.04	.864 $\pm$.05	.844 $\pm$.07
	HDTree	.967 $\pm$.01	.907 $\pm$.02	.860 $\pm$.03	.848 $\pm$.05
Magic	C4.4	.888 $\pm$.02	.806 $\pm$.02	.883 $\pm$.01	.913 $\pm$.03
	C4.5	.856 $\pm$.03	.809 $\pm$.02	.884 $\pm$.01	.914 $\pm$.03
	HDTree	.884 $\pm$.01	.793 $\pm$.02	.853 $\pm$.01	.848 $\pm$.02
Pen	C4.4	.986 $\pm$.01	.965 $\pm$.02	.945 $\pm$.02	.936 $\pm$.05
	C4.5	.968 $\pm$.02	.962 $\pm$.02	.945 $\pm$.03	.929 $\pm$.03
	HDTree	.988 $\pm$.01	.965 $\pm$.02	.948 $\pm$.03	.936 $\pm$.05
Statlog Landsat	C4.4	.912 $\pm$.03	.716 $\pm$.04	.537 $\pm$.07	.540 $\pm$.07
	C4.5	.731 $\pm$.14	.710 $\pm$.07	.542 $\pm$.10	.528 $\pm$.10
	HDTree	.916 $\pm$.02	.708 $\pm$.04	.541 $\pm$.07	.525 $\pm$.05
Statlog Image	C4.4	.987 $\pm$.03	.982 $\pm$.02	.971 $\pm$.03	.969 $\pm$.05
	C4.5	.981 $\pm$.02	.980 $\pm$.02	.970 $\pm$.03	.964 $\pm$.03
	HDTree	.992 $\pm$.01	.987 $\pm$.02	.978 $\pm$.03	.978 $\pm$.04

TABLE II: Results obtained on static datasets using 5×2 CV.

‘y’) to form the minority class, whilst for the Pen data set those instances with the class label ‘5’ formed the minority class.

D. Results

On static data sets we find that the HDTree algorithm developed in [4] achieves the best AUC values on 5 of the datasets used, as can be seen in Table II. On those that remain the HDTree has a joint best AUC value on two datasets, and is outperformed only once on the Magic dataset by the C4.4 algorithm. The AUC results obtained for the Letter, Pen, Landsat and Image datasets are similar to those reported in [4]. Thus our results generally support those reported in [4] for the HDTree. Though here we can see that possessing the best AUC value does not necessarily translate to the best imbalanced performance. For instance, whilst the AUC for the HDTree is higher on the Letter, Pulsar and Landsat data sets, the G-Mean and the recall rate is lower. Whilst the HDTree has a better AUC value, the other algorithms, particularly C4.4, must be operating at points on the ROC curve which provide better results.

In the tables that follow the results for the streamed test scenarios are given. In these tables the significance of each result is indicated using colour coding. Different colours indicate statistically different results at the level $\alpha = 0.01$. Here green indicates the best performing and red the worst i.e. green > yellow > red. Table III shows the results obtained on data streams with imbalances ranging from +1:-10 to +1:-100 with varied levels of labelling. For the least imbalanced streams (imbalance +1:-10), we see that for pulsar data there was no significant difference between the algorithms at the $\alpha = 0.01$ level. However there were significant differences observed during testing on the Skin and MiniBoone datasets. Here the VFDT consistently outperforms the two Hellinger based approaches. The performance of the GH-VFDT in particular worsens on Skin data as the labelling increases. The HD-VFDT is also affected since its performance drops when labelling reaches 75%. Crucially however, as the imbalance worsens to +1:-100 this situation reverses. The two Hellinger based approaches maintain their high G-Mean values, whilst

Dataset	Balance	+1 : -10				Rank	Balance	+1 : -100				Rank
	Labelling (%)	10	50	75	100		Labelling (%)	10	50	75	100	
Pulsar	HTree	.858/.926	.860/.904	.858/.901	.860/.905	2.5	HTree	.673/.876	.752/.854	.753/.852	.767/.859	2.75
	HD-VFDT	.860/.922	.855/.923	.855/.933	.850/.920	1.25	HD-VFDT	.457/.835	.472/.929	.529/.919	.493/.931	1.5
	GH-VFDT	.858/.918	.852/.915	.851/.909	.850/.917	2.25	GH-VFDT	.518/.903	.483/.920	.536/.916	.552/.916	1.75
Skin	HTree	.843/.911	.894/.958	.920/.969	.930/.973	1	HTree	.472/.671	.486/.615	.541/.655	.574/.693	3
	HD-VFDT	.737/.893	.746/.898	.726/.877	.727/.888	2.75	HD-VFDT	.277/.887	.292/.885	.319/.903	.263/.892	2
	GH-VFDT	.815/.911	.824/.904	.835/.913	.816/.914	1.75	GH-VFDT	.399/.898	.478/.915	.441/.904	.430/.900	1
MiniBoone	HTree	.954/.958	.993/.993	.992/.993	.992/.992	1	HTree	.918/.932	.964/.972	.972/.978	.980/.982	1
	HD-VFDT	.889/.913	.859/.786	.882/.912	.859/.818	2	HD-VFDT	.874/.866	.791/.857	.758/.854	.876/.949	2
	GH-VFDT	.568/.398	.666/.450	.661/.497	.760/.665	3	GH-VFDT	.742/.597	.547/.395	.623/.498	.511/.459	3

TABLE III: F-score/G-Mean results for the *HoeffdingTree* (HTree), *Hellinger Distance Tree* (HD-VFDT) [4] and *GH-VFDT* classifiers, trained before classifying a stream. Imbalances ranging from +1:-10 to +1:-100, significance tests used $\alpha = 0.01$.

Dataset	Balance	+1 : -1,000				Rank	Balance	+1 : -10,000				Rank
	Labelling (%)	10	50	75	100		Labelling (%)	10	50	75	100	
Pulsar	HTree	.372/.746	.576/.785	.581/.795	.610/.781	3	HTree	.194/.544	.114/.289	.159/.334	.170/.331	3
	HD-VFDT	.106/.924	.107/.929	.106/.924	.100/.934	1	HD-VFDT	.014/.921	.013/.928	.011/.938	.016/.919	1
	GH-VFDT	.137/.910	.131/.915	.128/.905	.125/.928	2	GH-VFDT	.028/.904	.016/.920	.020/.917	.026/.886	2
Skin	HTree	.137/.532	.106/.302	.069/.225	.087/.239	3	HTree	.018/.511	.017/.219	.011/.095	.021/.141	3
	HD-VFDT	.043/.802	.047/.891	.051/.893	.046/.882	2	HD-VFDT	.004/.871	.005/.914	.005/.908	.006/.869	1.75
	GH-VFDT	.082/.911	.082/.921	.080/.922	.080/.908	1	GH-VFDT	.008/.895	.009/.892	.009/.917	.009/.916	1.25
MiniBoone	HTree	.576/.617	.918/.951	.812/.862	.907/.943	1.5	HTree	.604/.672	.586/.667	.651/.830	.649/.707	2
	HD-VFDT	.284/.842	.650/.837	.206/.776	.455/.944	1.5	HD-VFDT	.021/.844	.039/.792	.048/.875	.328/.853	1
	GH-VFDT	.293/.469	.709/.656	.332/.200	.331/.300	3	GH-VFDT	.667/.500	.260/.447	.603/.494	.302/.489	3

TABLE IV: F-score/G-Mean results for the *HoeffdingTree* (HTree), *Hellinger Distance Tree* (HD-VFDT) [4] and *GH-VFDT* classifiers, trained before classifying a stream. Imbalances ranging from +1:-1,000 to +1:-10,000, significance tests used $\alpha = 0.01$.

for the VFDT this metric drops. The Hellinger algorithms now perform significantly better than the VFDT on both the Pulsar and Skin datasets at $\alpha = 0.01$. This trend continues as the imbalance worsens further as shown in Table IV. Here we see that the G-Mean of the VFDT has effectively dropped from .905 for the least imbalanced pulsar stream with 100% labelling, to .331 for the most imbalanced with the same labelling. Both Hellinger approaches by contrast have maintained consistent G-Mean values regardless of the imbalance. This is reflected in the minority class recall rates for the three classifiers. The HD-VFDT and GH-VFDT return more than double the true positives of the VFDT on pulsar data (see Table V). The MiniBoone results are an exception to this trend, since all three algorithms perform similarly here throughout.

Balance	+1 : -10,000			
Labelling (%)	10	50	75	100
HTree	.307/.001	.087/.001	.116/.001	.121/.001
HD-VFDT	.863/.017	.876/.016	.895/.017	.860/.016
GH-VFDT	.829/.014	.861/.017	.855/.013	.801/.012

TABLE V: Recall/false positive rate on pulsar data with an imbalance of +1:-10,000 and 10% labelled data.

In summary these results indicate that for streams possessing class imbalances greater than +1:-10, the use of algorithms such as the GH-VFDT and HD-VFDT can result in statistically significant increases in the minority class recall rate over standard stream classifiers such as the VFDT. Of the Hellinger distance based approaches described in this paper, our stream implementation of the HDTree algorithm from [4] was the best performer. It typically obtained the highest G-Mean values, and significantly improved recall rates over the VFDT. Our algorithm GH-VFDT, also significantly improved recall rates over the VFDT. Whilst it did not achieve recall rates quite as high as the HD-VFDT, its false positive rate was in fact

lower. Crucially the difference between these algorithms on the most imbalanced streams was not significant at $\alpha = 0.01$ (on all but the MiniBoone data). The GH-VFDT and HD-VFDT also achieved G-Mean values and recall rates which were higher than those observed for each of the static classifiers described in Section VI-A. Thus we argue that the GH-VFDT is a viable classifier for imbalanced streams particularly when the memory required by the HD-VFDT hinders its use.

VII. RELATED WORK

The problems associated with learning from an imbalanced class distribution, are typically tackled using one of three approaches. *Data level* approaches seek to modify data sets in order to rebalance the class distribution directly. *Algorithm level* approaches focus on modifying algorithms rather than datasets, to achieve improved performance on the minority class examples. *Cost sensitive* approaches on the other hand utilise different cost matrices describing the costs associated with misclassifying a particular data instance, reflecting the fact that it is usually more costly to misclassify rare class instances. There are also many hybrids of these approaches, particularly of those that mix data and algorithm level approaches in some way. However practical efforts at solving the imbalanced learning problem have focused on developing new data level sampling techniques. Some representative approaches include One-sided selection (OSS) [20], Wilsons Editing (WE) [21], [22], the Synthetic Minority Over-sampling Technique (SMOTE) [23], Borderline SMOTE [24], and ADASYN [25].

More recently some have begun investigating how these perform on streaming data. Nguyen et al. [26] for instance compared over-sampling and under-sampling techniques in the context of streamed data. Their results suggest that whilst under-sampling performs better than over-sampling at smaller training set sizes, performance converges as the training set

becomes larger (i.e. as more streamed data is used to train a classifier). Based upon these results the authors propose a new approach, multiple random under-sampling (MRUS), which generates m random undersamplings of the stream, each of which is used to train a separate classifier forming an ensemble. In [27] Chen et al. developed their recursive weighted ensemble approach (REA) for classifying non-stationary imbalanced data streams. REA adaptively pushes minority class examples into the current data chunk to explicitly balance the class distribution. The Stream ensemble framework developed in [28] attempts to mitigate similar problems by combining an ensemble classifier with sampling approaches. Each classifier in the ensemble is trained on a data set containing all the positive instances seen in the stream, and a unique subset of negative instances undersampled from the stream. The approach is designed to operate under a batch model.

VIII. CONCLUSION

In this paper we have presented a new classification algorithm for imbalanced data streams called GH-VFDT. Through an empirical investigation we have demonstrated that the algorithm can effectively improve minority class recall rates on imbalanced data, with similar levels of performance to the algorithm described in [4], whilst returning fewer false positives. A natural extension to this work would be to modify the decision tree algorithm so that rather than predicting the majority class at a tree leaf, we predict the class that possesses the maximal summed affinity to the features possessed by an instance. Furthermore, we would like to change how split points are calculated at nodes, such that the best split is obtained by choosing the point with the lowest Hellinger distance variance.

ACKNOWLEDGMENT

This work was supported by grant EP/I028099/1 for the University of Manchester Centre for Doctoral Training in Computer Science, from the UK Engineering and Physical Sciences Research Council (EPSRC). Experiments were carried out upon data obtained by the Parkes Observatory, funded by the Commonwealth of Australia and managed by the CSIRO.

REFERENCES

- [1] Proceedings of the AAAI 2000 Workshop on Learning from Imbalanced Data Sets, "Learning from Imbalanced Data Sets", AAAI Tech Report WS-00-05, Nathalie Japkowicz ed., AAAI Press, 2000.
- [2] Proceedings of the ICML 2003 Workshop on Learning from Imbalanced Data Sets II, World Wide Web accessed (12/12/2013), <http://www.site.uottawa.ca/~nat/Workshop2003/workshop2003.html>, Nitesh Chawla, Nathalie Japkowicz, Alek Kolcz eds., 2003.
- [3] N. V. Chawla, N. Japkowicz, and A. Kotcz, "Editorial: Special Issue on Learning from Imbalanced Data Sets", *SIGKDD Explor. Newsl.*, vol. 6, no. 1, pp. 1–6, 2004.
- [4] D. A. Cieslak and N. V. Chawla, "Learning Decision Trees for Unbalanced Data", in *Machine Learning and Knowledge Discovery in Databases* (W. Daelemans, B. Goethals, and K. Morik, eds.), vol. 5211 of *LNCS*, pp. 241–256, 2008.
- [5] R. Smits, M. Kramer, B. Stappers, et al., "Pulsar Searches and Timing with The Square Kilometre Array", *A & A*, vol. 493, no. 3, 2009.
- [6] D. R. Lorimer, A. J. Faulkner, A. Lyne, et al., "The Parkes Multibeam Pulsar Survey - VI. Discovery and Timing of 142 Pulsars and a Galactic Population Analysis", *MNRAS*, vol. 372, pp. 777–800, 2006.
- [7] M. Galar, A. Fernandez, E. Barrenechea, et al., "A Review on Ensembles for the Class Imbalance Problem: Bagging-, Boosting-, and Hybrid-Based Approaches", *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, vol. 42, no. 4, pp. 463–484, 2012.
- [8] H. He and E. Garcia, "Learning from Imbalanced Data", *IEEE Transactions on Knowledge and Data Engineering*, vol. 21, no. 9, 2009.
- [9] N. Japkowicz and S. Stephen, "The Class Imbalance Problem: A Systematic Study", *Intelligent Data Analysis*, vol. 6, pp. 429–449, 2002.
- [10] X. Wu and Q. Yang, "10 Challenging Problems In Data Mining Research", *International Journal of Information Technology & Decision Making*, vol. 05, no. 04, pp. 597–604, 2006.
- [11] M. Wasikowski and X. wen Chen, "Combating the Small Sample Class Imbalance Problem Using Feature Selection", *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, 2010.
- [12] R. J. Lyon, J. D. Knowles, J. M. Brooke, and B. W. Stappers, "A Study on Classification in Imbalanced and Partially-Labelled Data Streams", in *IEEE Systems, Man, and Cybernetics*, pp. 1506–1511, 2013.
- [13] R. Akbani, S. Kwek, and N. Japkowicz, "Applying Support Vector Machines to Imbalanced Datasets", in *Machine Learning: ECML 2004*, vol. 3201 of *LNCS*, pp. 39–50, 2004.
- [14] D. A. Cieslak, T. R. Hoens, N. V. Chawla, et al., "Hellinger Distance Decision Trees Are Robust and Skew-insensitive", *Data Mining and Knowledge Discovery*, vol. 24, no. 1, pp. 136–158, 2012.
- [15] G. Hulten, L. Spencer, and P. Domingos, "Mining time-changing data streams", in *Proceedings of the seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2001.
- [16] T. Dietterich, "Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms", *Neural Computation*, vol. 10, no. 7, pp. 1895–1923, 1998.
- [17] R. R. Bouckaert, "Efficient AUC Learning Curve Calculation", in *Advances in AI*, vol. 4304 of *LNCS*, pp. 181–191, 2006.
- [18] D. J. Hand, "Measuring Classifier Performance: A Coherent Alternative to The Area Under The ROC Curve", *Machine Learning*, vol. 77, no. 1, pp. 103–123, 2009.
- [19] M. J. Keith, A. Jameson, W. van Straten, et al., "The High Time Resolution Universe Pulsar Survey - I. System configuration and initial discoveries", *MNRAS*, vol. 409, no. 2, pp. 619–627, 2010.
- [20] M. Kubat and S. Matwin, "Addressing the Curse of Imbalanced Training Sets: One-Sided Selection", in *Proceedings of the 14th International Conference on Machine Learning*, pp. 179–186, 1997.
- [21] D. L. Wilson, "Asymptotic Properties of Nearest Neighbor Rules Using Edited Data", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 2, pp. 408–421, 1972.
- [22] R. Barandela, M. Valdovinos, R. J. Sánchez, and F. Ferri, "The Imbalanced Training Sample Problem: Under or over Sampling?", in *Structural, Syntactic, and Statistical Pattern Recognition*, vol. 3138 of *LNCS*, pp. 806–814, 2004.
- [23] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique", *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002.
- [24] H. Han, W.-Y. Wang, and B.-H. Mao, "Borderline-SMOTE: A New Over-Sampling Method in Imbalanced Data Sets Learning", in *Advances in Intelligent Computing*, vol. 3644 of *LNCS*, pp. 878–887, 2005.
- [25] H. He, Y. Bai, E. Garcia, and S. Li, "ADASYN: Adaptive Synthetic Sampling Approach for Imbalanced Learning", in *IEEE International Joint Conference on Neural Networks*, pp. 1322–1328, 2008.
- [26] H. Nguyen, E. Cooper, and K. Kamei, "A Comparative Study on Sampling Techniques for Handling Class Imbalance in Streaming Data", in *Joint 6th International Conference on Soft Computing and Intelligent Systems and 13th International Symposium on Advanced Intelligent Systems*, pp. 1762–1767, 2012.
- [27] S. Chen and H. He, "Towards Incremental Learning of Nonstationary Imbalanced Data Stream: A Multiple Selectively Recursive Approach", *Evolving Systems*, vol. 2, pp. 35–50, 2011.
- [28] J. Gao, B. Ding, W. Fan, J. Han, and P. Yu, "Classifying Data Streams with Skewed Class Distributions and Concept Drifts", *Internet Computing*, vol. 12, no. 6, pp. 37–49, 2008.