

RESEARCH

Open Access



Using discretization for extending the set of predictive features

Avi Rosenfeld^{1*}, Ron Illuz¹, Dovid Gottesman¹ and Mark Last²

Abstract

To date, attribute discretization is typically performed by replacing the original set of continuous features with a transposed set of discrete ones. This paper provides support for a new idea that discretized features should often be used in addition to existing features and as such, datasets should be extended, and not replaced, by discretization. We also claim that discretization algorithms should be developed with the explicit purpose of enriching a non-discretized dataset with discretized values. We present such an algorithm, D-MIAT, a supervised algorithm that discretizes data based on minority interesting attribute thresholds. D-MIAT only generates new features when strong indications exist for one of the target values needing to be learned and thus is intended to be used in addition to the original data. We present extensive empirical results demonstrating the success of using D-MIAT on 28 benchmark datasets. We also demonstrate that 10 other discretization algorithms can also be used to generate features that yield improved performance when used in combination with the original non-discretized data. Our results show that the best predictive performance is attained using a combination of the original dataset with added features from a “standard” supervised discretization algorithm and D-MIAT.

1 Introduction

Discretization is a data preprocessing technique that transforms continuous attributes into discrete ones. This is accomplished by dividing each numeric attribute A into m discrete intervals where $D = \{[d_0, d_1], (d_1, d_2], \dots, (d_{m-1}, d_m]\}$ where d_0 is the minimal value, d_m is the maximal value and $d_i < d_{i+1}$ for $i = 0, 1, \dots, m-1$. The resulting values within D constitute a discretization scheme for attribute A and $P = \{d_1, d_2, \dots, d_{m-1}\}$ is A 's set of cut points. Traditionally, discretization has been used in place of the original values such that after preprocessing the data, D is used instead of A [1].

Two basic types of discretization exist, supervised and unsupervised. Unsupervised discretization divides each A into a fixed number of intervals within D , typically through equal-width (EW) or equal-frequency (EF) heuristics [2]. Supervised discretization further considers the target class, C , in creating D . One popular supervised discretization algorithm is based on information entropy

maximization (IEM) whereby the set of cut points is created to minimize the entropy within D [3]. An interesting by-product of IEM, and other supervised discretization methods, is that if no cut points are found according to the selection criteria, then only one bin is created for that variable, effectively eliminating A from the dataset as D is the null set. In this way, supervised discretization can also function as a type of feature selection [4]. A large number of supervised discretization algorithms have been developed in addition to IEM, including Ameva, CAIM, CACC, Modified-Chi2, HDD, IDD, and ur-CAIM [5–11]. We refer the reader to a recent survey [1] for a detailed comparison of these and other algorithms.

While discretization was originally performed as a necessary preprocessing step for certain supervised machine-learning algorithms which require discrete feature spaces [3], several additional advantages have since been noted. First, discretization can at times improve the performance of some classification algorithms that do not require discretization. Consequently, multiple studies have used D instead of A to achieve improved prediction performance [3, 12–14]. Most recently, this phenomenon has particularly been noted within medical and bioinformatic datasets [12–15]. It has been hypothesized that the source for this improvement is due to the attribute selection

*Correspondence: rosenfa@jct.ac.il

¹Jerusalem College of Technology, Jerusalem, Israel

Full list of author information is available at the end of the article

component evident within supervised discretization algorithms [13]. A second advantage of discretization is that it contributes to the interpretability of the machine-learning results, as people can better understand the connection between different ranges of values and their impact on the learned target [3, 13].

This paper's first claim is that discretization should not necessarily be used to replace a dataset's original values but instead to generate new features that can augment the existing dataset. As such, each discretized feature D should be used in *addition* to the original feature A . Using discretization in this fashion is to the best of our knowledge a completely novel idea. We claim that using datasets with both A and D can often improve the performance of a classification algorithm for a given dataset.

This claim is somewhat surprising and may not seem intuitive. It is widely established that reducing the number of features through feature selection, and by extension using discretization for feature selection, helps to improve prediction performance [16]. At the core of this claim is that the curse of dimensionality can be overcome by reducing the number of candidate predictive attributes. Counter to this claim, we posit that it is not the number of attributes that is problematic, but the lack of added value within those attributes. At times, the information gained from the discretized attributes is significant, to the point that its addition to the original attributes improves prediction accuracy.

This paper also contains a second claim that discretization algorithms should be explicitly developed for the purpose of augmenting the original data. As support for this point, we present D-MIAT, an algorithm that discretizes numeric data based on minority interesting attribute thresholds. We believe D-MIAT to be unique as to the best of our knowledge, it is the first discretization algorithm that explicitly extends a set of features through discretization. D-MIAT generates features where only the minority of an attribute's values strongly point to one of the target classes. We claim that at times it can be important to create discretized features with such indications. However, attribute selection approaches to date typically treat all values within a given attribute equally, and thus focus on the general importance of all values within a given attribute, or combinations of the full set of different attributes' values [17, 18]. Hence, these approaches would typically not focus on cut points based on strong indications within only a small subset of values. Once again, the potential success of D-MIAT may seem counterintuitive as it generates features in addition to the original dataset, something often believed to reduce performance [16].

To support these claims we studied the prediction accuracy within the 28 datasets of a previous discretization study [11]. We considered the performance of seven different classification algorithms: Naive Bayes, SVM, K-nn,

AdaBoost, C4.5, Random Forest, and Logistic regression, on five different types of datasets. First, we considered the accuracy of the original dataset (A) without any discretization. Second, we created a dataset with the original data combined with the features generated by D-MIAT and studied where this combination was more successful. Third, we compared the accuracy of the baseline datasets (A) to the discretized datasets (D) from the 10 algorithms previously considered [11], using the training and testing data from that paper. Somewhat surprisingly, we found that prediction accuracy on average *decreased* when only the discretized data was considered. Based on the understanding that discretized features can improve performance, we then created a fourth dataset which appended the original data to the features generated by each of the 10 canonical discretization algorithms, creating 10 new combined datasets based on A and D . Again, we noted that the combination of the discretized features with the original data has improved predictive performance. Fifth, we studied how combinations of features from different discretization algorithms can be created. Specifically, we created a dataset that combined the discretized values of D-MIAT and the three discretization algorithms with the best prediction performance. The combination of the three types of features, the original data, those of the "standard" discretization algorithms and D-MIAT, performed the best.

2 The D-MIAT Algorithm

Contrary to other discretization algorithms, the D-MIAT was explicitly developed to augment the original data with discretized features. It will only generate such features when strong indications exist for one of the target classes, even within a relatively small subsets of values within A . Our working hypothesis is that these generated features improve prediction accuracy by encapsulating features that classification algorithms could miss in the original data or using "standard" discretization methods.

This assumption is motivated by recent findings that at times values of important subsets of attributes can serve as either "triggers" or "indicators" for biological processes. Recent genomic (DNA) and transcriptomic (RNA) research has shown that some people may have a natural predisposition or immunity towards certain diseases [19]. Similarly, we posit that even small subsets of values pointing to one of the target classes can be significant, even within non-medical datasets. The success of D-MIAT is in finding these subsets and thus we shift our focus from studying all attribute values, as has been done to date by other discretization algorithms, to finding those important subsets of values within the range of a given attribute.

To make this general idea clearer, consider the following example. Assume that attribute A is a numeric value for how many cigarettes a person smokes in a given day. The dataset contains a total population of 1000 where only 10% smoked more than three cigarettes a day. Most of these heavier smokers develop cancer (say 95 of 100 people) while the cancer rates within the remaining 90% (900 people) are not elevated. Traditional discretization will analyze all values within the attribute equally and may thus ignore this relatively small, but evidently important, subset of this dataset. Methods such as IEM that discretize based on the overall effectiveness of the discretization criteria; here, entropy reduction will not find this attribute interesting as this subset is not necessarily large enough for a significant entropy reduction within the attribute. Accordingly, it will not discretize this attribute, effectively removing it from the dataset. Similarly, even discretization algorithms with selection measures not based on information gain will typically ignore the significance of this subset of minority attribute values due to its size, something that D-MIAT was explicitly designed to do.

Specifically, we define the D-MIAT algorithm as follows. We assume that n numeric attributes exist in the dataset, which are denoted $A_1 \dots A_n$. Each given attribute, A_j , has a continuous set of values, which can be sorted and then denoted as $val_1 \dots val_b$. There are c target values (class labels) within the target variable C with $c \geq 2$. A new discretized cut will be created if a subset of a minimum size $supp$ exists with a strong indication, $conf$, for one of the target values within C . As per D-MIAT's assumption, we only attempt to create up to two cuts within A_j between the minimum val_1 to d_1 and between the maximum value val_b and d_2 . In both cases, the size of each subset needs to contain at least $supp$ records.

We considered two types of "strong" indications as support, one type based on entropy and one based on lift. Most similar to IEM, we considered confidence thresholds based on entropy. Entropy can be defined as $\sum_{i=1}^C -p_i \log_2 p_i$ where p_i here is the relative size of C_i within C . This value can be used a threshold to quantify the support as the percentage of values within a given attribute that points to a given target C_i . For example, if the smallest 10% of values of A_1 point to the first target value C_1 , then the entropy within this subset is 0 and thus constitutes a strong indication. A second type of confidence is based on lift, which is typically defined as $\frac{p(x|c)}{p(x)}$ where $p(x|c)$ is the conditional probability for x given c divided by the general probability of x . We use lift in this content as a support measure for the relative strength of subset for a target value given its general probability. For example, let us again consider a subset with the lowest 10% of values within A_1 with respect to C_1 . The relative

strength of a subset can here be defined as the probability of $P([val_1 \dots d_1] | C_1) / P([val_1 \dots d_1])$. Assuming that the general probability of C_1 is 0.05 as it occurs 50 times out of 1000 records of which 40 of these occurrences are within the 100 lowest values of A_1 and only 10 additional times within the other 900 values, then its lift will be $(40/100)/(50/1000)$ or 8. Assuming D-MIAT's value for the $conf$ parameter is less than 8, then this subset would be considered significant and D-MIAT will create a discretized cut with that subset.

It is worth noting that the two parameters within D-MIAT, $supp$ and $conf$, are motivated from association rule learning [20]. In both cases, we attempt to find even relatively small subsets of values, however, above a minimum amount of support $supp$. Similarly, in both cases, we refer to a confidence threshold that can be defined as the absolute and relative number of instances within a subset corresponding to a target value. To the best of our knowledge, D-MIAT is the first discretization algorithm using support and confidence to find these significant subsets of attribute values.

Algorithm 1: D-MIAT Algorithm to Generate Features with Important Values

```

1: Procedure D-MIAT( $\{A\}$ )
2: Initialize  $\{D\}$  to be empty
3: for  $j = A_1$  to  $A_n$  do
4:   Sort( $A_j$ )
5:   for  $i = C_1$  to  $C_c$  do
6:     bound1 = BinarySearchLower( $A_j, C_i, conf$ )
7:     if Size[ $A_1, A_{bound1}$ ]  $\geq$   $supp$  then
8:       cut1 = Discretize( $A_1, A_{bound1}$ )
9:        $D = D + cut1$ 
10:    end if
11:    bound2 = BinarySearchHigher( $A_j, C_i, conf$ )
12:    if Size[ $A_{bound2}, A_b$ ]  $\geq$   $supp$  then
13:      cut2 = Discretize( $A_{bound2}, A_b$ )
14:       $D = D + cut2$ 
15:    end if
16:  end for
17: end for
18: return ( $\{A\} + \{D\}$ )

```

Based on these definitions, Algorithm 1 presents D-MIAT. Line 1 begins with creating a null set of discretized features D that will be extracted from the full set of attributes, A . Lines 2–4 loop through all attributes within the dataset (A), sort the continuous values within each attribute A_j , and consider each target variable C_i . We then consider two discretization ranges, one beginning with the smallest value of A_j and one ending with the largest value in A_j (lines 6 and 11). The algorithm uses a binary

search (BinarySearch) to find potential cut points based on the selection criteria, *conf*. Trivially, this step could end with bound1 being equal to the smallest value within A_j and bound2 being A_j 's largest value. Typically, the subset will be beyond these two points. Regardless, lines 7 and 12 check if the number of records within this interval is larger than the predefined support threshold, *supp*. If so, a new discretized variable is created. In our implementation, the new discretized variable will have values of 1 for all original values within $[A_1, A_{bound1}]$ and zero for $(A_{bound1}, b]$ or a value of 1 for $[A_{bound2}, A_b]$ and zero for $[A_1, A_{bound2})$. In lines 9 and 14 we add these new cuts to the generated features within D . Thus, this algorithm can potentially create two cuts for every value of A_j , but will typically create much fewer as per the stringent requirements typically assigned to *supp* and *conf* by association rule algorithms. Line 18 returns the new dataset combining the original attributes in A with the new discretized features from D-MIAT.

The motivation for D-MIAT's attempt to potentially create two cuts at the extremes, and not to search for support within any subset as per association rule theory, is based on medical research. One motivation for D-MIAT comes from the observation that the probability distribution of features relating to any medical classification problem need not be unimodal. In fact, as discussed in the previous smoker example, features are likely to be multi-modal with some modes having very low representation in the sampled dataset, but within those modes prevalence of one class may be significantly different from that for the remainder of the samples. Furthermore, we expect to find that the discretized cuts at the extreme values for A_j would have the highest amount of interpretability—one of the main general motivations behind discretization [3, 13].

Table 1 presents a small portion of one of the 28 datasets, glass, used in our experiments. In this dataset, the goal is to identify the type of glass with target values of 1–7. We present three attributes from within this

dataset: $A_1 = Ba, A_2 = Ca$ and $A_3 = Fe$ and due to limited space and the size of the dataset only a subset of values for these three attributes is shown. D-MIAT was run using a minimal support parameter of 10% (*supp*=0.1) of the dataset and a lift value of 1.5 as minimal confidence (*conf* \geq Lift(1.5)). Note that the values of Ba are sorted with the non-zero values being highlighted in the first column of Table 1. The algorithm found that for the 33 records of the training set where $A_1 > 0.09$, 23 records belonged to class 7 (see column 4). Class 7 occurs with approximately 0.1 probability within the training set and with approximately 0.7 probability for the cut $A_1 > 0.09$. D-MIAT then computed that the $Lift(P(7)|A_1 > 0.09) = P(7|A_1 > 0.09)/P(7) = 7$ which is much greater than 1.5. Thus, D-MIAT created a cut for the higher values of Ba based on line 11 of Algorithm 1, with the resulting cut, D_1 being shown in column 4 of Table 1. Note that the discretized cut is for all values within this range, including those with other target values, such as class 1. No cut was created based on the lower values of A_1 (line 6 of Algorithm 1) as the size of the resulting cut (line 7) was not greater than the required support. Conversely, for $A_2 = Ca$, a cut was created based on the value of bound1 in line 6 of Algorithm 1. Here, D-MIAT found strong indications for one of the target values—here target number 2—as it found that of the 80 instances in the training set with values less than 8.4, 43 corresponded to this target (see column 5). This probability (0.53) is significantly higher than the general probability of target 2 within the dataset (0.35) and thus the lift is greater than 1.5 (1.51).

For comparison, we present in columns 6–8 the discretized cuts from the IEM algorithm for these three attributes. The IEM algorithm is based on the use of the entropy minimization heuristic for discretization of a continuous value into 0, 2, or more intervals. For the Ba attribute it minimized the overall entropy with one cut at 0.4158, and thus created two intervals—either less than or greater than this value. For the Ca attribute, four

Table 1 A sample dataset with D-MIAT applied

| Ba | Ca | Fe | Lift=1.5 Ba >0.09: [23/33] from 7 | Lift=1.5 Ca <8.4: [43/80] from 2 | IEM-Ba | IEM-Ca | IEM-Fe | Target |
|---------|----------|--------|-----------------------------------|----------------------------------|--------------|-------------------|--------|--------|
| 0 | 6.07628 | 0 | 0 | 1 | (-inf-.4158] | (-inf-7.0171] | All | 2 |
| 0 | 7.35604 | 0.0612 | 0 | 1 | (-inf-.4158] | (7.0171-8.31906] | All | 2 |
| 0 | 7.59276 | 0 | 0 | 1 | (-inf-.4158] | (7.0171-8.31906] | All | 6 |
| 0 | 7.77568 | 0 | 0 | 1 | (-inf-.4158] | (7.0171-8.31906] | All | 1 |
| 0 | 7.82948 | 0 | 0 | 1 | (-inf-.4158] | (7.0171-8.31906] | All | 1 |
| 0 | 7.82948 | 0 | 0 | 1 | (-inf-.4158] | (7.0171-8.31906] | All | 2 |
| 0 | 7.86176 | 0 | 0 | 1 | (-inf-.4158] | (7.0171-8.31906] | All | 2 |
| 0 | 7.88328 | 0 | 0 | 1 | (-inf-.4158] | (7.0171-8.31906] | All | 2 |
| 0 | 7.89404 | 0 | 0 | 1 | (-inf-.4158] | (7.0171-8.31906] | All | 2 |
| 0 | 7.9048 | 0 | 0 | 1 | (-inf-.4158] | (7.0171-8.31906] | All | 2 |
| 0.1 | 8.67952 | 0 | 1 | 0 | (-inf-.4158] | (8.31906-9.90078] | All | 7 |
| 0.68985 | 8.66876 | 0 | 1 | 0 | (.4158-inf) | (8.31906-9.90078] | All | 1 |
| 0.75915 | 8.66876 | 0 | 1 | 0 | (.4158-inf) | (8.31906-9.90078] | All | 7 |
| 1.3797 | 10.66876 | 0 | 1 | 0 | (.4158-inf) | (9.90078-inf) | All | 7 |

intervals were created and for Fe no intervals were created (represented by the uniform value of “All” in this implementation). This example highlights both the similarities and differences between D-MIAT and other discretization algorithms. First, D-MIAT is a binary discretization method— each cut D-MIAT creates will only have two intervals, one where *supp* and *conf* are met and one where they are not met. In contrast, IEM, and other classic discretization algorithms maximize a score over all attribute values (such as entropy reduction for IEM). Thus, these algorithms often choose different numbers of cuts (such as IEM creating for four different interval cuts for Ca) and thresholds for these cuts (0.09 for D-MIAT within Ba and 0.4158 for IEM).

These different cut values impact the interpretability of the results. As D-MIAT focuses on a subset, it will focus one’s analysis on a range of values either below a given threshold (if the subset range starts at the minimum value) or above a given threshold (if the subset range ends at the maximum value). In Table 1, we see both examples. Please note that the first D-MIAT cut focuses one’s attention on the higher range of values ($Ba > 0.09$), where a strong indication exists for the target value of 7. The second cut focuses one’s attention on the smaller range of values ($Ca < 8.4$), where a strong indication exists for the target value of 2. Often these algorithms agree and will both find nothing of interest. Note that Fe had no discretized cuts formed by either D-MIAT or IEM as both algorithms’ conditions for creating cuts were not present for this attribute. As we explain below, all experiments were generated with the discretization cuts being generated only on the training data and then applied to the testing data.

3 Experimental results

We used the same 28 datasets within a recent discretization study [11]. These datasets were collected from the the KEEL [21] and UCI [22] machine learning repositories and represent a variety of complexity, number of classes, number of attributes, number of instances, and imbalance ratio (ratio of the size of the majority class to the minority class). Detailed information about the datasets can be found online¹. We downloaded the discretized versions of these datasets for the discretization algorithms they considered. This included the unsupervised algorithms of equal-width (EW) and equal-frequency (EF) [2], and the supervised algorithms of information entropy maximization (IEM) [3], class-attribute interdependence maximization (CAIM) [5], Ameva [6], Modified-Chi2 [7], Hypercube Division-based Discretization (HDD) [8], Class-Attribute Contingency Coefficient (CACC) [9], Interval Distance-Based Method for Discretization (IDD) [10] and ur-CAIM [11]. These datasets each contain 100 different folds whereby the discretized intervals were determined

by the training portion in the first 90% of the file and these intervals were then applied to the remaining 10% of the data used for testing. Thus, the 28 datasets contained 10 independently constructed training and testing components to effectively create 10-fold cross validation for a total of 280 training and testing pairs.

The general thesis of this work is that adding features is not inherently detrimental, so long as these features have some added value. Conversely, we claim that traditional use of discretization, where the continuous values are replaced with discretized ones, can be detrimental if important information is lost by removing the original continuous values. We expected to find that datasets enriched with D-MIAT provide more accurate results than those without it, and more generally removing the original features in favor of exclusively using the discretized ones would be less effective than using the combination of features in one extended set. We checked several research questions to study these issues:

1. Do the features generated by D-MIAT improve prediction performance when adding them to the original data?
2. Does the classic use of discretization of removing the continuous values help improve performance?
3. Is it advantageous to use the discretized features *in addition* to the original ones?
4. Should the features generated by D-MIAT be used in conjunction with other discretization algorithms?

In order to check these issues we then proceeded to create five different sets of data. The first set of data was composed of the 28 base datasets without any modification to their training and testing components (*A*) and used the data from a previous study [11]. The second set of data was the original data in addition to the discretized cuts (*D*) from D-MIAT. The third set of data consisted of the discretized versions of the 28 datasets with the 10 abovementioned algorithms (*D*) and were also generated based on a previous study [11]. The fourth set of data appended the original datasets (*A*) with each of the discretized datasets (*D*). Last, we created a fifth set of data by appending the original data to the features created by D-MIAT and several other best performing discretization algorithms. To facilitate replicating these experiments in the future, we have made a Matlab version of D-MIAT available at: <http://homedir.jct.ac.il/~rosenfa/D-MIAT.zip>². We ran D-MIAT on a personal computer with an Intel i7 processor and 32 GB of memory. Processing the full set of 280 files with D-MIAT took a total of approximately 15 min.

We then measured the predictive accuracy of 7 different classification algorithms: Naive Bayes, SVM using the RBF kernel, K-nn using the value of $k = 3$, AdaBoost,

C4.5, Random Forest, and Logistic regression on each of these datasets. The first 6 algorithms were implemented in Weka [23] with the same parameters used in previous work [11]. The last algorithm was added as it is a well-accepted deterministic classification algorithm that was not present in the previous study. The default parameters were used within Weka's Simple Logistic implementation of this algorithm.

As shown below, we found that the classic use of discretization often did not improve the average performance across these datasets. Instead, using discretization in addition to the original features typically yielded better performance, either through using D-MIAT in addition to the original data or through using discretized features from canonical algorithms in addition to the original data. The best performance was typically reached by the combined dataset of the original data enriched with D-MIAT and other discretization algorithms as posited by the research question #4. Thus, we present evidence that the most effective pipeline is using D-MIAT along with discretized features to extend a given set of features.

3.1 D-MIAT results

The goal of the D-MIAT algorithm is to supplement, not supplant, the original data. The decision as to whether additional features will be generated depends on the *supp* and *conf* parameters in Algorithm 1 and it is applied to every numeric attribute A_j within the dataset. Thus, D-MIAT could potentially generate two features for every attribute based on lines 7 and 12 of the Algorithm 1. In our experiments, we defined *supp* equal to 10% of the training data. Three confidence values, *conf*, were checked. The first value was $conf = Entropy(0)$, meaning the cut yielded a completely decisive indication for one of the target classes and thus yielded zero entropy as per the first type of confidence mentioned above. We also considered 2 types of lift confidence: $conf = Lift(1.5)$ and $conf = Lift(2.0)$. As per these confidence levels, we checked if the cut yielded a stronger indication for one of the target classes as measured per their lift values. Potentially, both of these lift thresholds could be satisfied. For example, assume a cut yielded a lift of 3, then both confidence thresholds would consider this cut significant and generate a new feature accordingly. We also considered the possibility that the cuts could be added cumulatively and thus overlapping cuts could be added based on combinations of these 3 different thresholds. Conversely, if all thresholds used were not met, no cuts were created for a given attribute.

To illustrate this point, Table 2 presents the number of continuous attributes in each of the 28 datasets and the total number of attributes D-MIAT generated within each of these datasets combined with the three thresholds for *conf* we considered. As D-MIAT generated its

cuts exclusively based on 10 different iterations within the training set, it was possible that the number of features that D-MIAT generated would vary across these iterations within the 28 datasets. For example, note that the first two datasets, abalone and arrhythmia, had no features generated by D-MIAT across any of the iterations and for any of the values of *conf* and *supp*. This shows that the D-MIAT can have a more stringent condition for generating cuts than all of the 10 canonical discretization algorithms previously studied that did generate features in all of these datasets. Also note that D-MIAT generated an average of 12.9 features within the glass dataset. Based on the training data, D-MIAT generated between 12 and 15 features, leading to the average being a fraction. Most times, as is the case here, the number of discretized features generated by D-MIAT is less than the number of the original features. However, a notable example is found in the iris dataset which contains only four continuous attributes, but on average, D-MIAT generated 20 features. In this case, many cuts were generated for each A_j as per each one of *conf* conditions in Algorithm 1.

We then checked if the features generated by D-MIAT improved predictive performance. Table 3 displays the average accuracy of the baseline data for each of the 15 files where D-MIAT generated cuts which are noted in Table 2. Please note that each of the D-MIAT *supp* thresholds—0 entropy in line 2, Lift of 1.5 in line 3, and lift of 2 in line 4—typically improved the prediction performance across all the classification algorithms. This demonstrates that the values chosen within this threshold are not extremely sensitive as all values chosen typically improve prediction accuracy. Note that all performance increases are highlighted in the table. A notable exception is the AdaBoost algorithm where no performance increase was noted (nor any decrease). Also, we note that certain algorithms, such as Naive Bayes, SVM, and logistic regressions consistently benefited, while the other algorithms did not. We also find that combination of the features created by the cuts from all three D-MIAT thresholds, the results of which are found in the fifth line of Table 3, yielded the best performance. As each of the cuts generated by D-MIAT is significant, it is typically advantageous to add all cuts in addition to the original data. For the remainder of the paper, we will use the results of D-MIAT using all cuts as we found this yielded the highest performance. Thus, we overall found support for this paper's thesis that the addition of features is not necessarily problematic, unless there is a lack of added value within those attributes, even if the thresholds within the discretized cuts are not necessarily optimal.

3.2 Using discretization in addition to the original data

We then proceeded to check research questions 2 and 3. We assumed that research question #2 would also be

Table 2 The total number of new features generated by three D-MIAT thresholds in the each of the 28 datasets

| Dataset | Attributes | Average features from D-MIAT | Range | Dataset | Attributes | Average features from D-MIAT | Range |
|----------------|------------|------------------------------|-------|-------------------|------------|------------------------------|---------|
| Abalone | 8 | 0 | - | Penbased | 16 | 0 | - |
| Arrhythmia | 279 | 0 | - | Pendigits | 16 | 0 | - |
| Glass | 9 | 12.9 | 12–15 | Pima | 8 | 3.1 | 3–4 |
| Heart | 13 | 6.8 | 6–7 | Satimage | 36 | 128.2 | 127–130 |
| Ionosphere | 33 | 26.1 | 22–31 | Segment | 19 | 97.6 | 96–10 |
| Iris | 4 | 20 | 19–21 | Sonar | 60 | 18.2 | 15–20 |
| Jm1 | 21 | 0 | - | Spambase | 57 | 32.9 | 31–35 |
| Madelon | 500 | 0 | - | Spectrometer | 102 | 0 | - |
| Mc1 | 38 | 12 | 12 | Texture | 40 | 0 | - |
| Mfeat-factors | 216 | 0 | - | Thyroid | 21 | 1 | 1 |
| Mfeat-fourier | 76 | 0 | - | Vowel | 13 | 0 | - |
| Mfeat-karhunen | 64 | 0 | - | Waveform | 40 | 39.5 | 37–42 |
| Mfeat-zernike | 47 | 0 | - | Winequality-red | 11 | 3.9 | 3–4 |
| Pc2 | 36 | 17.2 | 17–19 | Winequality-white | 4 | 2.7 | 2–4 |

found to be correct and that using the discretized features alone would typically be useful, as has been previously been found [3, 12–14]. To our surprise, we did not find this to often be the case within the 28 datasets we considered. Table 4, displays the average accuracy results of this experiment. The baseline here is again the average performance without discretization (A) to which we compared the 10 canonical discretization algorithm across the same 7 classification algorithms. Once again, each colored cell represents an improvement of a given discretization algorithm versus the baseline and note the lack of color within much of Table 2. While some algorithms, particularly Naive Bayes and SVM, showed often large improvements versus the non-discretized data, the majority of classification algorithms (K-nn, AdaBoost, C4.5, Random Forest and Logistic Regression) typically did not show any improvement across the discretization algorithms in these datasets. This point is illustrated in the last column of this table, which shows that for most discretization algorithms, a *decrease* in accuracy is on average found through exclusively using discretized features. Thus, we did find that for most classification algorithms in the datasets we considered, using discretization alone was *less* effective than using the original features.

We found significantly more empirical support for this paper's third research question, that using discretization in combination with the original features is more effective than using either the original set or those generated from discretization alone. The results of this experiment are found in Table 5 where we compare the results of the original dataset and the combination of the original data with the features from the discretization algorithms. We again find that the Naive Bayes and SVM algorithms benefit from this combined dataset the most. When considering these two learning algorithms, improvements were now noted through combining all discretization algorithms. We note a strong improvement in the results in Table 5 versus those in Table 4 within the AdaBoost and Logistic Regression algorithms. In contrast to the results in Table 4, the average performance across all algorithms (found in the last columns of the table) now shows an improvement. Also note that almost without exception, the combined features outperformed the corresponding discretized set. For example, the average accuracy for the Ameva algorithm with C4.5 is 70.46, but this jumps to 78.39 for Ameva combined with the original data. Even when improvements were not noted, we found that performance was typically not negatively affected by the

Table 3 Comparing the accuracy of the datasets without D-MIAT and with 4 variations of D-MIAT with parameter values of zero entropy (MIAT 0), Lift of 1.5 (MIAT 1.5), Lift of 2 (MIAT 2) and all three discretized values (MIAT 0 1.5 2)

| Dataset | Naïve Bayes | SVM | KNN | AdaBoost | C4.5 | RandomForest | Logistic | Average |
|-------------------------|-------------|-------|-------|----------|-------|--------------|----------|---------|
| Baseline | 77.23 | 74.91 | 83.33 | 71.96 | 83.09 | 86.17 | 83.31 | 80 |
| Baseline + MIAT 0 | 77.5 | 75.7 | 83.83 | 71.96 | 83.03 | 86.28 | 83.64 | 80.28 |
| Baseline + MIAT 1.5 | 77.49 | 79.72 | 82.46 | 71.96 | 83.36 | 85.89 | 83.54 | 80.63 |
| Baseline + MIAT 2 | 77.63 | 76.43 | 82.84 | 71.96 | 83.1 | 86.58 | 84.14 | 80.39 |
| Baseline + MIAT 0 1.5 2 | 78.36 | 82.02 | 83.56 | 71.96 | 83.22 | 85.82 | 84.08 | 81.29 |

Table 4 Comparing the accuracy results from seven different classification algorithms within the original, baseline data (A) and the discretized data (D) formed from the Ameva, CACC, CAIM, Chi2, EF, EW, HDD, IDD, IEM, and ur-CAIM algorithms

| Dataset | Naïve Bayes | SVM | KNN | AdaBoost | C4.5 | RandomForest | Logistic | Average |
|----------|-------------|-------|-------|----------|-------|--------------|----------|---------|
| Baseline | 74.06 | 71.43 | 80.45 | 52.57 | 79.16 | 81.95 | 80.57 | 74.31 |
| Ameva | 71.39 | 74.24 | 74.28 | 50.8 | 70.76 | 73.37 | 74.89 | 69.96 |
| CACC | 69.73 | 72.91 | 73.23 | 50.99 | 71.29 | 66.51 | 69.75 | 67.77 |
| CAIM | 76.68 | 79.13 | 78.8 | 52.26 | 75.4 | 78.42 | 80.23 | 74.42 |
| CHI2 | 76.18 | 79.34 | 79.16 | 51.14 | 74.22 | 77.65 | 79.85 | 73.93 |
| EF | 73.56 | 76.55 | 74.39 | 48.27 | 69.39 | 72.43 | 76.51 | 70.16 |
| EW | 74.95 | 75.19 | 73.34 | 48.78 | 70.65 | 72.5 | 76.12 | 70.22 |
| HDD | 68.82 | 65.52 | 65.4 | 47.88 | 59.86 | 59.72 | 65.12 | 61.76 |
| IDD | 74.83 | 77.75 | 77.4 | 51.96 | 71.27 | 75.27 | 78.67 | 72.45 |
| IEM | 76.17 | 79.95 | 79.92 | 51.52 | 77.21 | 79.89 | 80.75 | 75.06 |
| urCAIM | 76.29 | 79.17 | 79.63 | 52.62 | 78 | 76.69 | 80.16 | 74.65 |

addition of these features as one would assume due to the curse of dimensionality.

Despite the general support being found for the research question #3 in that using discretization to augment the original data is better than using the discretized data alone, we still note that some algorithms, particularly the K-nn, C4.5 and Random Forest learning algorithms, often did not benefit from any form of discretization. We further explore these results differences, and generally the impact of discretization on all algorithms, in Section 3.4.

3.3 Using D-MIAT and other discretization algorithms as additional features

Given the finding that adding features from both D-MIAT and canonical discretization algorithms helps improve performance, we hypothesized that the combination of these features would be effective to achieve the best performance as per research question #4. To evaluate this issue, we combined the D-MIAT features with those of the best performing discretization algorithms in these

datasets– CAIM, IEM, and urCAIM. We then considered the performance of this combination within the 15 datasets where D-MIAT generated some features and on average within all 28 datasets. The results of this experiment are found in Table 6.

As can be seen from the results from the 15 datasets (top portion of Table 6), the combination of discretization algorithms almost always outperforms the original data and improvements in predictive accuracy are noted in all three combinations in the Naive Bayes, SVM, K-nn, AdaBoost, C4.5, and logistic regression algorithms. The one exception seems to be the Random Forest algorithm where large performance improvements are not noted. For comparison, we also present the improvements across all 28 datasets in the bottom portion of Table 6, which include 13 datasets where there are no features generated by D-MIAT. As expected, the combination of D-MIAT was somewhat less significant once again demonstrating the benefit of adding the features from D-MIAT. Overall, and on average across the algorithms we considered, we

Table 5 Comparing the accuracy results from seven different classification algorithms within the original (A) and discretized data appended to the original data (A + D)

| Dataset | Naïve Bayes | SVM | KNN | AdaBoost | C4.5 | RandomForest | Logistic | Average |
|-----------------|-------------|-------|-------|----------|-------|--------------|----------|---------|
| Baseline | 74.06 | 71.43 | 80.45 | 52.57 | 79.16 | 81.95 | 80.57 | 74.31 |
| Appended Ameva | 75.53 | 77.01 | 79.51 | 52.74 | 78.39 | 79.91 | 81.27 | 74.91 |
| Appended CACC | 74.95 | 75.96 | 79.71 | 52.54 | 78.5 | 75.55 | 80.55 | 73.97 |
| Appended CAIM | 76 | 79.64 | 79.2 | 52.66 | 78.53 | 79 | 81.72 | 75.25 |
| Appended CHI2 | 76 | 80.25 | 80.41 | 52.73 | 76.98 | 78.67 | 81.59 | 75.23 |
| Appended EF | 74.78 | 77.95 | 77.16 | 52.59 | 76.77 | 74.09 | 81.14 | 73.5 |
| Appended EW | 75.07 | 77.14 | 75.94 | 52.75 | 77.53 | 74.4 | 81.44 | 73.47 |
| Appended HDD | 75.36 | 75.16 | 75.78 | 52.7 | 72.61 | 61.93 | 80.53 | 70.58 |
| Appended IDD | 75.5 | 78.33 | 78.82 | 52.71 | 77.79 | 76.6 | 81.26 | 74.43 |
| Appended IEM | 75.75 | 80.3 | 80.91 | 52.57 | 78.94 | 81.1 | 82.1 | 75.95 |
| Appended urCAIM | 75.63 | 79.72 | 80.07 | 52.79 | 78.62 | 80 | 81.58 | 75.49 |

Table 6 Combining D-MIAT with discretization features yields the highest performance when considering the 15 datasets where D-MIAT added features, and on average across all 28 datasets

| 15 Datasets with MIAT | NaiveBayes | SVM | KNN | AdaBoost | C4.5 | RandomForest | Logistic | Average |
|------------------------|------------|-------|-------|----------|-------|--------------|----------|---------|
| Baseline | 77.23 | 74.91 | 83.33 | 71.96 | 83.09 | 86.17 | 83.31 | 80.66 |
| Appended CAIM + MIAT | 80.03 | 84.66 | 84.23 | 72.14 | 83.58 | 85.93 | 85.08 | 82.24 |
| Appended IEM + MIAT | 80.03 | 84.06 | 84.59 | 72.12 | 83.59 | 86.22 | 84.4 | 82.14 |
| Appended urCAIM + MIAT | 79.98 | 84.49 | 84.41 | 72.29 | 84.29 | 86.16 | 84.87 | 82.36 |

| All 28 Datasets | Naive Bayes | SVM | KNN | AdaBoost | C4.5 | RandomForest | Logistic | Average |
|--------------------------|-------------|-------|-------|----------|-------|--------------|----------|---------|
| Baseline | 74.06 | 71.43 | 80.45 | 52.57 | 79.16 | 81.95 | 80.57 | 74.31 |
| Baseline + CAIM + MIAT | 76.27 | 79.93 | 79.81 | 52.66 | 78.83 | 79.29 | 81.85 | 75.52 |
| Baseline + IEM + MIAT | 76.21 | 80.45 | 81.23 | 52.57 | 79.08 | 81.31 | 81.89 | 76.11 |
| Baseline + urCAIM + MIAT | 75.95 | 80.16 | 80.53 | 52.79 | 79.25 | 80.12 | 81.52 | 75.76 |

found that this combination was the most successful, with prediction accuracies typically improving by over 1%. Thus, we found the thesis #4 to typically be correct and using D-MIAT in conjunction with existing discretization algorithms is recommended to enrich the set of features to be considered.

3.4 Discussion and future work

We were somewhat surprised that using discretization alone was not as successful as it has been previously found with such classification algorithms as Random Forests and C4.5 [3, 12–14]. We believe that differences in the datasets being analyzed is likely responsible for these gaps. As such, we believe that an important open question is to predict when discretization will be successful, given the machine learning algorithm and the dataset used for input. In contrast, we found that D-MIAT yielded more stable results as it typically only improved the performance, something other discretization algorithms did not do, especially for these two learning algorithms. We are now exploring how to make other discretization algorithms similarly more stable.

We note that the pipeline described in this paper of using D-MIAT and discretized features in addition to the original data was most effective with algorithms without discretization, namely the Naive Bayes, SVM, and Logistic Regression classification algorithms. Conversely, this approach was less effective with methods having a discretization component, particularly with C4.5, Random Forests, and AdaBoost algorithms. It has been previously noted that C4.5 has a localized discretization component (re-applied at each internal node of the decision tree) and thus may gain less from adding globally discretized features, which are split into the same intervals across all decision tree nodes Dougherty et al. [3]. Similarly, the decision stumps used as weak classifiers in AdaBoost perform essentially a global discretization and thus may have made it less likely to benefit from the globally discretized features added by D-MIAT, something that is evident from

Table 3. As can be noted from Table 6, the Random Forest algorithm benefited the least from the pipeline described in this paper— again possibly due to its discretization component. We hope to further study when and why algorithms with an inherent discretization component still benefit from additional discretization. Additionally, it seems that the K-nn algorithm, despite not having a discretization component, benefits less from the proposed pipeline than other algorithms. It seems plausible this is because this algorithm is known to be particularly sensitive to the curse of dimensionality [24, 25] and thus this specific algorithm does not benefit from the proposed approach while other classification algorithms are less sensitive to it. We plan to study this complex issue in our future research.

We believe that several additional directions should also be pursued for future work. First, this study did not consider learning from neural networks and deep learning, as these algorithms were not previously considered [11]. This is due to the relatively small size of several of the datasets within this study, which made it infeasible to obtain accurate deep learning models using this approach. We are currently considering additional datasets, particularly those with larger amounts of training data to allow us to better understand how deep learning can be augmented by discretized features. In a similar vein, we believe that interconnections likely exist between some of the generated discretized features. Multivariate feature selection and/or deep learning could potentially be used to help stress these interconnections and remove features which are redundant. Second, we propose using metacognition, or the process of learning about learning [26] to allow us to learn which discretized features should be added to a given dataset. We are also studying how one could find an optimal value, or set of values, for the *conf* and *supp* thresholds within D-MIAT. While this paper demonstrates that multiple D-MIAT thresholds can be used in combination, and each threshold does typically improve performance, we do not claim that the

thresholds used in this study represent an optimal value, or set of values, for all datasets. One potential solution to this would be to develop a metacognition mechanism for learning these thresholds for a given dataset. Similarly, it is possible that a form of metacognition could be added to machine learning algorithms as has been generally suggested within neural networks [27] to help achieve this goal.

4 Conclusions

In this paper, we present a paradigm shift in how discretization can be used. To date, discretization has typically been used as a pre-processing step which removes the original attribute values, replacing them with discretized intervals. Instead, we suggest using the features that are generated by discretization to extend the dataset by using both the discretized and non-discretized versions of the data. We have demonstrated that discretization can often be used to generate new features which should be used in addition to the dataset's non-discretized features. The D-MIAT algorithm we present in this paper is based on this assumption. This is because D-MIAT will only discretize values with particularly strong indications based on high confidence, yet relatively low support for the target class, as it assumes that classification algorithms will also be using the original data. We also show that other canonical discretization algorithms can be used in a similar fashion, and in fact a combination of the original data with D-MIAT and discretized features from other algorithms yields the best performance. We are hopeful that the ideas presented in this paper will advance the use of discretization and its application to new datasets and algorithms.

Endnotes

¹<http://www.uco.es/grupos/kdis/wiki/ur-CAIM/>

²The file "run_all.m" found in this zip was used to run D-MIAT in batch across all files in a specified directory.

Funding

The work by Avi Rosenfeld, Ron Illuz, and Dovid Gottesman was partially funded by the Charles Wolfson Charitable Trust.

Authors' contributions

AR, RI, and DG were responsible for data collection and analysis. AR and ML were responsible for the algorithm development and writing of the paper. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Author details

¹Jerusalem College of Technology, Jerusalem, Israel. ²Ben Gurion University of the Negev, Beersheba, Israel.

Received: 23 May 2017 Accepted: 4 January 2018

Published online: 18 January 2018

References

1. S García, J Luengo, JA Sáez, V Lopez, F Herrera, A survey of discretization techniques: Taxonomy and empirical analysis in supervised learning. *IEEE Trans. Knowl. Data Eng.* **25**(4), 734–750 (2013)
2. MR Chmielewski, JW Grzymala-Busse, Global discretization of continuous attributes as preprocessing for machine learning. *Int. J. Approx. Reason.* **15**(4), 319–331 (1996)
3. J Dougherty, R Kohavi, M Sahami, et al, in *Machine learning: proceedings of the twelfth international conference*, volume 12. Supervised and unsupervised discretization of continuous features (Morgan Kaufmann Publishers, San Francisco, 1995), pp. 194–202
4. H Liu, R Setiono, Feature selection via discretization. *IEEE Trans. Knowl. Data Eng.* **9**(4), 642–645 (1997)
5. LA Kurgan, KJ Cios, Caim discretization algorithm. *IEEE Trans. Knowl. Data Eng.* **16**(2), 145–153 (2004)
6. L Gonzalez-Abrial, FJ Cuberos, F Velasco, JA Ortega, Ameva: An autonomous discretization algorithm. *Expert Syst. Appl.* **36**(3), 5327–5332 (2009)
7. FEH Tay, L Shen, A modified chi2 algorithm for discretization. *IEEE Trans. Knowl. Data Eng.* **14**(3), 666–670 (2002)
8. P Yang, J-S Li, Y-X Huang, Hdd: a hypercube division-based algorithm for discretisation. *Int. J. Syst. Sci.* **42**(4), 557–566 (2011)
9. C-J Tsai, C-I Lee, W-P Yang, A discretization algorithm based on class-attribute contingency coefficient. *Inf. Sci.* **178**(3), 714–731 (2008)
10. FJ Ruiz, C Angulo, N Agell, Idd: a supervised interval distance-based method for discretization. *IEEE Trans. Knowl. Data Eng.* **20**(9), 1230–1238 (2008)
11. A Cano, DT Nguyen, S Ventura, KJ Cios, ur-caim: improved caim discretization for unbalanced and balanced data. *Soft Comput.* **20**(1), 173–188 (2016)
12. JL Lustgarten, V Gopalakrishnan, H Grover, S Visweswaran, in *AMIA. Improving classification performance with discretization on biomedical datasets* (American Medical Informatics Association (AMIA), Bethesda, 2008)
13. JL Lustgarten, S Visweswaran, V Gopalakrishnan, GF Cooper, Application of an efficient bayesian discretization method to biomedical data. *BMC Bioinformatics.* **12**(1), 309 (2011)
14. DM Maslove, T Podchyska, HJ Lowe, Discretization of continuous features in clinical datasets. *J. Am. Med. Inform. Assoc.* **20**(3), 544–553 (2013)
15. A Rosenfeld, DG Graham, R Hamoudi, R Butawan, V Eneh, S Khan, H Miah, M Niranjan, LB Lovat, in *2015 IEEE International Conference on Data Science and Advanced Analytics, DSAA. MIAT: A novel attribute selection approach to better predict upper gastrointestinal cancer* (Campus des Cordeliers, Paris, 2015), pp. 1–7
16. I Guyon, A Elisseeff, An introduction to variable and feature selection. *J. Mach. Learn. Res.* **3**(Mar), 1157–1182 (2003)
17. I Guyon, An introduction to variable and feature selection. *J. Mach. Learn. Res.* **3**, 1157–1182 (2003)
18. Y Saeys, I Inza, P Larrañaga, A review of feature selection techniques in bioinformatics. *Bioinformatics.* **23**(19), 2507–2517 (2007)
19. RA Hamoudi, A Appert, et al, Differential expression of nf-kappab target genes in malt lymphoma with and without chromosome translocation: insights into molecular mechanism. *Leukemia.* **24**(8), 1487–1497 (2010)
20. Z Zheng, R Kohavi, L Mason, in *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining. Real world performance of association rule algorithms* (ACM, New York, 2001), pp. 401–406
21. J Alcalá-Fdez, A Fernández, J Luengo, J Derrac, S García, L Sanchez, F Herrera, Keel data-mining software tool: data set repository, integration of algorithms and experimental analysis framework. *J. Mult. Valued Log. Soft. Comput.* **17**, 255–287 (2011)
22. M Lichman, *UCI Machine Learning Repository*. (University of California, School of Information and Computer Science, Irvine, 2013). <http://archive.ics.uci.edu/ml>
23. IH Witten, E Frank, *Data Mining: Practical Machine Learning Tools and Techniques, Second Edition, Morgan Kaufmann Series in Data Management Systems*. (Elsevier, Cambridge, 2005)

24. Jsh Friedman, et al., *Flexible metric nearest neighbor classification*. Technical report, Technical report. (Department of Statistics, Stanford University, 1994)
25. DW Aha, Editorial. *Artif. Intell. Rev.* **11**, 7–10 (1997)
26. C Watkins, *Learning about learning enhances performance* (Institute of Education, University of London, 2001)
27. R Savitha, S Suresh, N Sundararajan, Metacognitive learning in a fully complex-valued radial basis function neural network. *Neural Comput.* **24**(5), 1297–1328 (2012)

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)
