Classification with Streaming Features: An Emerging-Pattern Mining Approach

KUI YU, Simon Fraser University
WEI DING and DAN A. SIMOVICI, University of Massachusetts Boston
HAO WANG, Hefei University of Technology
JIAN PEI, Simon Fraser University
XINDONG WU, Hefei University of Technology and University of Vermont

Many datasets from real-world applications have very high-dimensional or increasing feature space. It is a new research problem to learn and maintain a classifier to deal with very high dimensionality or streaming features. In this article, we adapt the well-known emerging-pattern–based classification models and propose a semi-streaming approach. For streaming features, it is computationally expensive or even prohibitive to mine long-emerging patterns, and it is nontrivial to integrate emerging-pattern mining with feature selection. We present an online feature selection step, which is capable of selecting and maintaining a pool of effective features from a feature stream. Then, in our offline step, separated from the online step, we periodically compute and update emerging patterns from the pool of selected features from the online step. We evaluate the effectiveness and efficiency of the proposed method using a series of benchmark datasets and a real-world case study on Mars crater detection. Our proposed method yields classification performance comparable to the state-of-art static classification methods. Most important, the proposed method is significantly faster and can efficiently handle datasets with streaming features.

Categories and Subject Descriptors: I.5.2 [**Computing Methodologies**]: Design Methodology

General Terms: Classifier Design and Evaluation, Feature Evaluation and Selection

Additional Key Words and Phrases: Emerging patterns, streaming features, feature selection, classification

ACM Reference Format:
Kui Yu, Wei Ding, Dan A. Simovici, Hao Wang, Jian Pei, and Xindong Wu. 2015. Classification with streaming features: An emerging pattern mining approach. ACM Trans. Knowl. Discov. Data 9, 4, Article 30 (June 2015), 31 pages.
DOI: http://dx.doi.org/10.1145/2700409

A preliminary version of this article titled “Mining Emerging Patterns by Streaming Feature Selection” was published in the *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD’12)*, pp. 60–68.
Authors’ addresses: K. Yu and J. Pei, School of Computing Science, Simon Fraser University, 8888 University Drive, Burnaby, V5A 1S6, BC, Canada; emails: {kuiy, jpei}@cs.sfu.ca; W. Ding and D. A. Simovici, Department of Computer Science, University of Massachusetts Boston, 100 Morrissey Boulevard, Boston, 02125-3393, MA, USA; emails: {ding, dsim}@cs.umb.edu; H. Wang, Department of Computer Science, Hefei University of Technology, 193 Tunxi Road, Hefei, 230009, Anhui Province, China; email: jsjxwangh@hfut.edu.cn; X. Wu (corresponding author), Department of Computer Science, Hefei University of Technology, 193 Tunxi Road, Hefei, 230009, Anhui Province, China and Department of Computer Science, University of Vermont, 33 Colchester Avenue, Burlington, 05405, VT, USA; email: xwu@cs.uvm.edu.

DOI: http://dx.doi.org/10.1145/2700409

1. INTRODUCTION

In some real-world applications, we are facing datasets of not only very high dimensionality but also with features that keep arriving, that is, datasets with streaming features. For example, in order to monitor and analyze the environment in different areas, researchers may deploy a set of observation stations in those areas. Each station is treated as an object in the data collected. In other words, the number of data objects is fixed. However, the number of features in temporal domains keeps increasing, since new observation data is collected all the time. Even though the data collection rate may not be very high, the dimensionality of the underlying dataset can easily reach tens or hundreds of thousands after a while. As another example, in our research project in automatic detection of subkilometer craters in high-resolution planetary images [Ding et al. 2011], a fixed number of craters are continually monitored using the latest available high-resolution images, which become available and updated over time. Since impact craters are among the most studied geomorphic features in the solar system, yielding information about past and present geological processes and providing the only tool for measuring relative ages of observed geologic formations [Urbach and Stepinski 2009], it is invaluable to build and maintain robust classification models using the accumulated features in a streaming way extracted from the images available so far.

It is a novel challenge to learn and maintain a classification model on such data with streaming features, as new features keep arriving. Although classification on data streams has been well studied in the data-mining and machine-learning literature [Aggarwal 2010; Dong et al. 2003], to the best of our knowledge, the existing emerging-pattern–based classification methods only focus on scenarios in which new objects keep arriving and the features are predetermined. Those methods are orthogonal to the scenarios studied in this article.

To build and maintain effective classification models on datasets with streaming features, we have to address at least two major challenges. First, as the existing classification methods typically assume a predefined space of features, it is important to determine how to handle streaming features in an online manner. This problem is highly related to feature selection, that is, how to select and maintain a set of features from a stream of features. Second, it is essential to ascertain how to build and maintain a classification model using the features incrementally selected from a stream of features. Ideally, the model construction and updating should be seamlessly integrated with online feature selection.

In order to tackle high-dimensional data with streaming features, we adapt the well-known emerging-pattern–based classification methods [Dong and Li 1999; Novak et al. 2009]. Emerging patterns are well recognized and effective in classifying high-dimensional data, since an emerging pattern can handle a subpopulation in a subspace that deliberates a clear discriminative pattern [Dong et al. 1999; Duan et al. 2014]. However, it is still challenging to extend emerging patterns to classify datasets with streaming features. First, it is computationally expensive or even prohibitive to mine long-emerging patterns. As features are accumulated over time, many emerging patterns may become long, partially due to the redundancy among features. Second, it is challenging to integrate mining emerging patterns and feature selection. To address the problem of learning classification models on datasets with streaming features, in this article, we propose a "semi-streaming" approach. Specifically, to tackle the challenge of streaming features, we propose an online feature selection step, which is capable of selecting and maintaining a pool of effective features from a feature stream. Our feature selection step scans features one by one as they are available. Moreover, the proposed feature selection step is specially designed for emerging patterns. To tackle the challenge of classification model construction and maintenance, we propose an offline step. Periodically, we can compute and update emerging patterns from the pool of selected features that are picked by the online step. Although the emerging-pattern mining step cannot be made online in theory, this step can be conducted only periodically, and can be separated from the online step. In other words, when the classification model needs to be updated, the offline task of emerging-pattern mining can take place.

Our semi-streaming approach is fundamentally different from applying emerging-pattern mining in a straightforwardmanner on a dataset with streaming features. The online feature selection step substantially reduces the dimensionality of the feature space under which the offline emerging-pattern mining step has to handle. This becomes possible only when the feature selection step can handle streaming features in an online manner, and also can select features according to the requirements of emerging-pattern mining. Due to the effective online feature selection customized for emerging-pattern mining, the emerging patterns mined in the offline step tend to be short, since many redundant and correlated features are reduced before the patterns are mined. This practically facilitates emerging-pattern mining dramatically.

Using a series of benchmark data-sets, we evaluate the effectiveness and efficiency of the EPSF algorithm, and compare it with the baseline methods and the state-of-the-art methods. The empirical study clearly shows that our method not only achieves high accuracy, but also takes less CPU time than the existing classification methods. Most important, the EPSF algorithm can handle datasets with streaming features efficiently. In our real-world case study, we evaluate the EPSF algorithm with crater detection from planetary images. The experimental results show that our method is not only highly comparable with the existing crater detection algorithms, but also produces a concise set of emerging patterns that are interpretable for domain scientists to understand the crater data.

The remainder of this article is organized as follows. Section 2 reviews related work. Section 3 gives the preliminaries and Section 4 presents our approach. Section 5 reports our experimental results. Section 6 presents our conclusions.

2. RELATED WORK

Our work relates to emerging-pattern mining and feature selection, which we will briefly review in this section.

Dong and Li [1999] first introduced emerging patterns (EPs) to represent strong contrasts between different classes of data. In addition, a Jump Emerging Pattern (JEP) is a special type of EP whose support increases from zero in one class to nonzero in the other class [Li et al. 2001a]. Like other patterns composed of conjunctive combinations of feature-value pairs [Wang and Karypis 2005; Trépos et al. 2013; Sahoo et al. 2014], EPs can be easily understood and used directly in a wide range of applications [Song et al. 2014; Wang et al. 2013a], such as failure detection [Lo et al. 2009] and discovering knowledge in gene expression data [Fang et al. 2012].

Dong et al. [1999] proposed the first EP classifier, called Classification by Aggregating Emerging Patterns (CAEP). Based on CAEP, Li et al. [2001a] proposed a JEP classifier that is distinct from the CAEP classifier. The JEP classifier uses JEPs exclusively because JEPs discriminate between different classes more strongly than any other type of EP. Li et al. [2000] also presented a lazy EP classifier based on an instance-based EP discovery, called DeEPs, to improve the efficiency and accuracy of CAEP and JEP classifier. In addition, Fan and Ramamohanarao [2006] proposed a robust EP classifier, called SJEP-classifier, using a strong JEP. A strong JEP from class C_1 to class C_2 satisfies two conditions: (1) the support of itemset X is zero in C_1 but nonzero and satisfies a minimal support threshold in C_2, and (2) any proper subset of X does not satisfy condition (1). The SJEP classifier integrates the CP-tree miner into the EP classifier, and uses much fewer JEPs than the JEP classifier. The disadvantage of JEP and SJEP classifiers is that if a dataset contains no or very few JEP and SJEP patterns, then it will cause the classification performance of JEP and SJEP classifiers to be significantly reduced.

The well-known bottleneck of EP classifiers is that they are computationally prohibitive to dealing with a dataset with more than 60 dimensions without prior feature set reduction until the Zero-suppressed Binary Decision Diagrams (ZBDD) EP miner was proposed [Loekito and Bailey 2006]. The ZBDD EP miner can deal with a relatively high-dimensional dataset, but like previous EP mining approaches, it still suffers from an explosive number of EPs, even with a rather high support threshold. Accordingly, it is still a challenging research issue to build a robust EP classifier from high dimensionality.

Feature selection has been generally viewed as a problem of searching for a minimal subset of features in high-dimensional data that leads to the most accurate prediction model [Liu and Yu 2005]. There are two types of feature selection methods proposed in the literature: batch methods and online methods.

A batch method has to access the entire feature set on the training data and performs a global search for the best feature at each round [Brown et al. 2012]. Contrast to batch methods, online feature selection can be conducted online. Recently, Wang et al. [2013b] proposed an OSF algorithm for online feature selection. The OFS algorithm assumes that data instances keep arriving, and performs feature selection on each data instance as it is available. In contrast to OFS, the Fast-OSFS and alpha-investing algorithms were proposed to deal with the scenarios in which features keep arriving but the number of data instances is fixed [Zhou et al. 2006; Wu et al. 2010, 2013]. Zhou et al. [2006] presented alpha investing, which sequentially considers new features as the addition to a predictive model by modeling the candidate feature set as a dynamically generated stream. However, alpha investing requires prior information of the original feature set and never evaluates redundancy among the selected features as time goes on. To fix the drawbacks, Wu et al. [2010, 2013] presented the Online Streaming Feature Selection (OSFS) algorithm and its faster version, the Fast-OSFS algorithm.

In a recent study, Yu et al. [2013] integrated local causal-structure learning into EP mining to help reduce high dimensionality. The study has shown that this integration can efficiently extract a minimum number of sets of strongly predictive patterns from high-dimensional data and get highly accurate EP classifiers. Different from the work of Yu et al. [2013], which requires a complete set of features available beforehand, the proposed algorithm in this article is capable of dealing with a high-dimensional dataset with streaming features. Our new algorithm can mine emerging patterns from datasets of not only very high dimensionality but also with features that keep arriving, that is, datasets with streaming features.

3. NOTATIONS AND DEFINITIONS

3.1. Emerging Patterns

Consider a training dataset D is defined upon a feature set F and the class attribute C. F contains N features, that is, $F = \{F_1, F_2, \ldots, F_N\}$. For $\forall F_i \in F$, we assume that it is in a discrete domain $dom(F_i)$. Let I be the set of all items, that is, $I = \{F_i = f_i | F_i \in F, f_i \in Dom(F_i)\}$, the class attribute $C = \{C_1, C_2, \ldots, C_K\}$ be a finite set of K distinct class labels, and X be an itemset and $X \subseteq I$. The dataset D can be partitioned into $D_1, D_2, \ldots, D_K$, where D_j consists of instances with class label C_j, $j = 1, \ldots, K$. The mathematical notations used in this article are summarized in Table I.

The Support and Growth Rate (GR) of an itemset X, and an emerging pattern from D_l to $D_m (l, m = 1, \ldots, K, and\ l \neq m)$, are defined as follows.

Definition 3.1 (*Support*).

$$support_D(X) = count_D(X)/|D| \tag{1}$$

Table I. Summary of Mathematical Notations

Notation	Mathematical meanings
D	training dataset
D_l, D_m	subsamples of D
F	a feature set from D
F_i	a single feature (attribute), $F_i \in F$
f_i	a discrete value of F_i
$Dom(.)$	$Dom(F_i)$ denotes all discrete values of F_i
S	a feature subset $S \subseteq F$
C	the class attribute
C_i	a class label, $C_i \in C$
I	a set of all items from F
X	an itemset, $X \in I$
$count_D(X)$	the number of instances in D that supports X
$\|.\|$	$\|D\|$ returns the number of instances in D
$support_D(X)$	the support value of X on D
$GR_{D_l \rightarrow D_m}(X)$	the growth rate of X from D_l to D_m
ρ	a threshold of $GR_{D_l \rightarrow D_m}(X)$
e	an emerging pattern
E_i	a set of emerging patterns
$Rateimp(e)$	the growth rate improvement of e
$P(.\|.)$	$P(C = C_j \| S = s)$ denotes the posterior probability of C_j conditioned on s
t	a time point
T	a test instance
$CMB(C)_t$	a Markov blanket of C selected at time t

where $count_D(X)$ is the number of instances in D containing X and $|D|$ is the number of instances in D.

Definition 3.2 (*GR: Growth Rate*). [Dong and Li 1999]

$$GR_{D_l \rightarrow D_m}(X) = support_{D_m}(X)/support_{D_l}(X) \tag{2}$$

If $support_{D_m}(X) = 0$ and $support_{D_l}(X) = 0$, then $GR_{D_l \rightarrow D_m}(X) = 0$; if $support_{D_m}(X) \neq 0$ but $support_{D_l}(X) = 0$, then $GR_{D_l \rightarrow D_m}(X) = \infty$.

Definition 3.3 (*EP: Emerging Pattern*). [Dong and Li 1999] Given a threshold $\rho > 1$, an EP from D_l to D_m is an itemset X, where $GR_{D_l \rightarrow D_m}(X) \geq \rho$.

An emerging pattern e from D_l to D_m is also called an EP of D_m. If $GR(e) = \infty$, e is called a Jumping EP (JEP). The goal of EP mining is to extract EP set E_i for class C_i, which consists of EPs from $\{D - D_i\}$ to D_i, given a minimum support threshold and a minimum growth rate threshold.

A positive GR improvement threshold is introduced to ensure a concise and representative set of EPs that are not subsumed by one another and consist of items that are strong contributors to their predictive power. The GR improvement can also help to reduce the search space by eliminating EPs that are redundant. Here is the definition of GR improvement.

Definition 3.4 (*Growth Rate Improvement*). [Zhang et al. 2000a] Given an EP e, the GR improvement of e, $Rateimp(e)$, is defined as the minimum difference between its GR and the GRs of all of its subsets,

$$Rateimp(e) = \min(\forall e' \subset e, GR(e) - GR(e')). \tag{3}$$

3.2. Feature Relevance in Feature Selection

In feature selection, a feature space F in general can be divided into three disjoint groups: strongly relevant, weakly relevant, and irrelevant features [Koller and Sahami 1995]. The goal of feature selection is to select a subset from F without performance degradation on prediction models. In the following definitions, $P(C = C_j|S = s)$ is the posterior probability of class C_j given a set of values of s of a subset S.

Definition 3.5 (*Conditional Independence*). Two distinct features $F_i \in F$ and $F_k \in F$ are conditionally independent on a feature subset $S \subseteq F - \{F_i \cup F_k\}$, if and only if there exists an assignment of values f_i and f_k, s.t.

$$P(F_i = f_i|F_k = f_k, S = s) = P(F_i = f_i|S = s). \tag{4}$$

Definition 3.6 (*Strong Relevance*). A feature F_i is strongly relevant to the class attribute C, if and only if there exists an assignment of values f_i, C_j, and s for which $P(S = s, F_i = f_i) > 0$,

$$\forall S \subseteq F - \{F_i\}\ s.t.\ P(C = C_j|S = s, F_i = f_i) \neq P(C = C_j|S = s). \tag{5}$$

Definition 3.7 (*Weak Relevance*). A feature F_i is weakly relevant to the class attribute C, if and only if it is not strongly relevant, and $\forall f_i, C_j$, and s for which $P(S = s) > 0$,

$$\exists S \subset F - \{F_i\}\ s.t.\ P(C = C_j|S = s) \neq P(C = C_j|S = s, F_i = f_i). \tag{6}$$

Definition 3.8 (*Irrelevance*). A feature F_i is irrelevant to the class attribute C, if and only if it is neither strongly nor weakly relevant, and there exists an assignment of values f_i, C_j, and s for which $P(S = s, F_i = f_i) > 0$,

$$\forall S \subseteq F - \{F_i\}\ s.t.\ P(C = C_j|S = s, F_i = f_i) = P(C = C_j|S = s). \tag{7}$$

Definition 3.9 (*Markov Blanket*). [Koller and Sahami 1995] The Markov blanket of feature F_i, denoted as $M_i \subset F - \{F_i\}$ makes all other features independent of F_i given M_i, that is,

$$\forall F_k \in F - (M_i \cup \{F_i\})\ s.t.\ P(F_i|M_i, F_k) = P(F_i|M_i). \tag{8}$$

With Markov blankets, weakly relevant features can be divided into redundant features and nonredundant features [Yu and Liu 2004].

Definition 3.10 (*Redundant Feature*). A feature is redundant, hence should be removed from F, if and only if it is weakly relevant and has a Markov blanket within F.

4. MINING EMERGING PATTERN WITH STREAMING FEATURES

It is computationally expensive to evaluate the complete item combinations for a high-dimensional dataset. To mitigate this problem, we propose to effectively prune the feature space before EP mining. Recently, Wu et al. [2013] proposed streaming feature selection to deal with data with streaming features. In this model, the number of data instances is fixed, while features keep arriving and each feature is evaluated upon its arrival. Compared to traditional feature selection, the strength of streaming feature selection is that the number of features is no longer required to be fixed in advance. Feature selection is streamlined and conducted online to be able to deal with streaming features.

Although emerging-pattern mining cannot be done online in theory, in this article, we propose a "semi-streaming" approach to bridge streaming feature selection and EP mining to learn and maintain an EP classification model on data with streaming features.

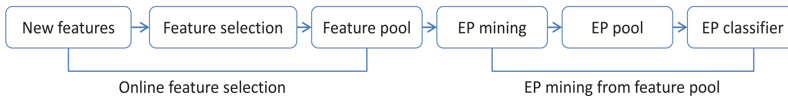

Fig. 1. The framework of a semi-streaming approach.

Definition 4.1 (*Semi-streaming approach*). A semi-streaming approach includes an online feature selection step customized for an offline step to periodically mine emerging patterns from the features that are picked by the online step.

In the semi-streaming approach, the online feature selection step is to select and maintain a pool of effective features from streaming features by scanning features one by one as they are available. An offline step is to construct and update an EP classifier by periodically mining emerging patterns from the pool of selected features that are picked by the online step. Definition 4.1 indicates that, although emerging-pattern mining cannot be fully online, the offline step needs to be conducted only periodically. When the EP classifier needs to be updated, the offline task of emerging-pattern mining will take place.

The framework of the semi-streaming approach is illustrated in Figure 1. At the first stage, we present an online feature selection step that is capable of selecting and maintaining a pool of effective features from a feature stream. Our feature selection step processes features one by one as they are available. At the second stage, we propose an offline step. Periodically, we can compute and update emerging patterns from the pool of selected features that are picked by the online step. There are two key research problems to be addressed:

(1) How to build an influential feature candidate pool to be used for EP mining as features are available over time
(2) How to build an EP pool to be used for classification by extracting EPs from this influential feature pool

4.1. Online Building an Influential Feature Pool

In Figure 1, the feature pool shall keep the features that are useful only for producing predictive EPs, and it may be updated over time as features are available one by one. To customize the online feature selection step for efficient emerging-pattern mining, we must be able to evaluate the degree of feature relevance with the discriminative power of EPs. We have theoretically proved the association of causal relevance in causal Bayesian networks and EP discriminability in EP mining [Yu et al. 2013]. Here we provide theoretically an analysis on the relationships between feature relevance (irrelevant features, strongly relevant features, and redundant features) and EP discriminability (non-EPs, strongly predictive EPs, and redundant EPs) in the following propositions.

As discussed in Definitions 3.2 and 3.3, the discriminability of an EP is determined by its support value and GR. Proposition 4.2 establishes the relations between non-EPs and irrelevant features.

PROPOSITION 4.2. *For* $\forall F_i \in F$, $\forall f_i \in dom(F_i)$, *and* $\forall C_j \in dom(C)$, $GR_{\{D-D_j\}\to D_j}(F_i = f_i) = 1$ *holds if and only if* F_i *is irrelevant to* C.

PROOF. Assume a dataset D has two classes: positive class C_p and negative class C_n, $C = \{C_p, C_n\}$. D_p represents C_p class data, D_n represents C_n class data, $sup_{D_p}(F_i = f_i)$ is the support value of the itemset $\{F_i = f_i\}$ in D_p, and $sup_{D_n}(F_i = f_i)$ is its support value in D_n. Then $GR(F_i = f_i)$ from D_n to D_p is calculated as follows.

$$\begin{aligned} GR_{D_n \to D_p}(F_i = f_i) &= \frac{sup_{D_p}(F_i = f_i)}{sup_{D_n}(F_i = f_i)} \\ &= \frac{P(F_i = f_i|C = C_p)}{P(F_i = f_i|C = C_n)} \\ &= \frac{P(F_i = f_i, C = C_p)}{P(C = C_p)} \Big/ \frac{P(F_i = f_i, C = C_n)}{P(C = C_n)} \\ &= \frac{P(C = C_p|F_i = f_i)P(F_i = f_i)}{P(C = C_p)} \Big/ \frac{P(C = C_n|F_i = f_i)P(F_i = f_i)}{P(C = C_n)} \\ &= \frac{P(C = C_p|F_i = f_i)}{P(C = C_n|F_i = f_i)} \bullet \frac{P(C = C_n)}{P(C = C_p)} \end{aligned}$$

If $GR_{D_n \to D_p}(F_i = f_i) = 1$, then the following holds.

$$\frac{P(C = C_p)}{P(C = C_n)} = \frac{P(C = C_p|F_i = f_i)}{P(C = C_n|F_i = f_i)}$$

As $P(C = C_p) + P(C = C_n) = 1$ and $P(C = C_p|F_i = f_i) + P(C = C_n|F_i = f_i) = 1$, we get

$$P(C = C_p|F_i = f_i) = P(C = C_p)$$

(with $\frac{a}{b} = \frac{c}{d}$ equivalent to $\frac{a}{b+a} = \frac{c}{d+c}$), as well as $P(C = C_n|F_i = f_i) = P(C = C_n)$.

According to Definition 3.8, for any assignments $f_i \in dom(F_i)$ and $C_j \in dom(C)$ to F and C, $P(C = C_j|F_i = f_i) = P(C = C_j)$ holds, therefore F_i is irrelevant to C. Similarly, from D_p to D_n, if $GR_{D_p \to D_n}(F_i = f_i) = 1$, we can also prove that F_i is irrelevant to C.

On the other hand, if F_i is irrelevant to C, we get

$$\begin{aligned} GR_{D_n \to D_p}(F_i = f_i) &= \frac{P(C = C_p|F_i = f_i)}{P(C = C_n|F_i = f_i)} \bullet \frac{P(C = C_n)}{P(C = C_p)} \\ &= \frac{P(C = C_p|F_i = f_i)}{P(C = C_n|F_i = f_i)} \bullet \frac{1 - P(C = C_p)}{P(C = C_p)} \\ &= 1 \end{aligned}$$

Thus, Proposition 4.2 is proven. □

According to Definition 3.4, for an EP e, if we can find an $e' \subset e$ to make $Rateimp(e) \leq 0$, then e is a redundant EP, and might be replaced by a subset within e. Thus, avoiding generation of those redundant EPs in advance will improve search efficiency. Proposition 4.3 explains the relationship between feature redundancy and EP redundancy.

PROPOSITION 4.3. *For* $\exists F_i \in F$, $\exists S \subset F - F_i, \forall f_i \in dom(F_i), \forall s \subset \bigcup_{k=1}^{|S|} dom(S_k)$, *and* $\forall C_j \in dom(C)$, $GR_{\{D-D_j\} \to D_j}(F_i = f_i, S = s) = GR_{\{D-D_j\} \to D_j}(S = s)$ *holds, if and only if* F_i *is redundant to* C *conditioning on the subset* S.

PROOF. $GR(F_i = f_i, S = s)$ from D_n to D_p is calculated as follows.

$$\begin{aligned}
GR_{D_n \to D_p}(F_i = f_i, S = s) &= \frac{sup_{D_p}(F_i = f_i, S = s)}{sup_{D_n}(F_i = f_i, S = s)} \\
&= \frac{P(F_i = f_i, S = s|C = C_p)}{P(F_i = f_i, S = s|C = C_n)} \\
&= \frac{P(F_i = f_i, S = s, C = C_p)}{P(C = C_p)} \Bigg/ \frac{P(F_i = f_i, S = s, C = C_n)}{P(C = C_n)} \\
&= \frac{P(C = C_p|F_i = f_i, S = s)P(F_i = f_i, S = s)}{P(C = C_p)} \\
&\quad \Bigg/ \frac{P(C = C_n|F_i = f_i, S = s)P(F_i = f_i, S = s)}{P(C = C_n)} \\
&= \frac{P(C = C_p|F_i = f_i, S = s)}{P(C = C_n|F_i = f_i, S = s)} \bullet \frac{P(C = C_n)}{P(C = C_p)}
\end{aligned}$$

From D_n to D_p, $GR(S = s) = P(S = s|C = C_p)/P(S = s|C = C_n)$

$$\begin{aligned}
\frac{P(C = C_p|F_i = f_i, S = s)}{P(C = C_n|F_i = f_i, S = s)} \bullet \frac{P(C = C_n)}{P(C = C_p)} &= \frac{P(S = s|C = C_p)}{(P(S = s|C = C_n)} \\
\frac{P(C = C_p|F_i = f_i, S = s)}{P(C = C_n|F_i = f_i, S = s)} &= \frac{P(S = s|C = C_p)P(C = C_p)}{(P(S = s|C = C_n)P(C = C_n)} \\
\frac{P(C = C_p|F_i = f_i, S = s)}{P(C = C_n|F_i = f_i, S = s)} &= \frac{P(C = C_p|S = s)}{P(C = C_n|S = s)}
\end{aligned}$$

Using the same reasoning in proving Proposition 4.2, we can get two equations $P(C = C_p|F_i = f_i, S = s) = P(C = C_p|S = s)$ and $P(C = C_n|F_i = f_i, S = s) = P(C = C_n|S = s)$. By Definitions 3.9 and 3.10, we can find a subset $S \subset F$ as a Markov blanket of F_i, and for any assignments $f_i \in dom(F_i), s \subseteq dom(S)$ and $C_j \in dom(C)$ to F_i, S and C, $P(C = C_j|F_i = f_i, S = s) = P(C = C_j|S = s)$ holds, thus F_i is redundant to C given S.

On the other hand, if F_i is redundant to C, from D_n to D_p, then the following holds.

$$\begin{aligned}
GR_{D_n \to D_p}(F_i = f_i, S = s) &= \frac{P(C = C_p|F_i = f_i, S = s)}{P(C = C_n|F_i = f_i, S = s)} \bullet \frac{P(C = C_n)}{P(C = C_p)} \\
&= \frac{1 - P(C = C_n|F_i = f_i, S = s)}{P(C = C_n|F_i = f_i, S = s)} \bullet \frac{P(C = C_n)}{P(C = C_p)} \\
&= \frac{1 - P(C = C_n|S = s)}{P(C = C_n|S = s)} \bullet \frac{P(C = C_n)}{P(C = C_p)} \\
&= \frac{P(C = C_p|S = s)}{P(C = C_n|S = s)} \bullet \frac{P(C = C_n)}{P(C = C_p)} \\
&= \frac{P(S = s|P(C = C_p)}{P(S = s|C = C_n)} \\
&= GR_{D_n \to D_p}(S = s)
\end{aligned}$$

Thus, Proposition 4.3 is proven. □

Proposition 4.3 shows that if F_i is redundant to C conditioned on a subset S, then an itemset $\forall f_i \in dom(F_i)$ together with an itemset $\forall s \subset \bigcup_{k=1}^{|S|} dom(S_k)$ contains the same predictive information as the itemset $\forall s \subset \bigcup_{k=1}^{|S|} dom(S_k)$.

With Propositions 4.2 and 4.3, as features are processed in a sequential scan, we can online build an influential feature pool and discarding irrelevant and redundant features to avoid generating non-EPs or redundant EPs in the EP mining process later on.

Once we understand the feature relevance with the discriminative power of EPs, the next step is to understand how to build this influential feature pool online, and then how to adjust the feature pool once a new feature is added into the pool. To build an influential feature pool, we need to assess online whether a new feature is irrelevant. If so, it is discarded. If not, we use Proposition 4.4 proposed by Wu et al. [2013]] to handle this newly arrived feature.

PROPOSITION 4.4. *A current Markov blanket of C at time t is denoted as $CMB(C)_t$. Assume a new feature F_i at time $t+1$ is weakly relevant to C, if $\exists S \subseteq CMB(C)_t$ such that $P(C|F_i, S) = P(C|S)$, then F_i can be discarded.*

After F_i is added into *CMB(C)*, we must check whether any existing features in the feature pool become redundant. We shall use Proposition 4.5 proposed by Wu et al. [2013] to update the current feature pool online, determining which of the existing features in the current feature pool can be removed as F_i is added.

PROPOSITION 4.5. *With $CMB(C)_t$ at time t, a new feature F_i arrives at time $t+1$, and there does not exist any $MB(F_i)$ within $CMB(C)_t$. If $\exists Y \in CMB(C)_t$ and $\exists S \subseteq \{CMB(C)_t \cup F_i\} - \{Y\}$ s.t. $P(C|Y, S) = P(C|S)$, then Y can be removed from $CMB(C)_t$.*

4.2. Building an EP Pool

The EP pool stores the candidate EPs that are mined from the feature pool. To make the EP pool correspond to the changes of the feature pool, we divide the construction of the EP pool into two steps. One step is to build 1-itemset EP pool online. This 1-itemset EP pool should be updated as the feature pool is updated. The other step is offline but periodically mines all EPs from the 1-itemset EP pool to construct an EP classifier.

4.2.1. Building a 1-Itemset EP Pool Online.

(1) *Building*. As a new feature F_i arrives, we first assess whether it is irrelevant; if so, it is discarded. Otherwise, we evaluate whether it is redundant to C by Proposition 4.4; if so, it is also discarded. If not, it is added to the feature pool *CMB(C)*. Then, the EPSF algorithm (discussed in detail in Section 4.4) converts feature F_i into a set of itemsets I_{F_i} and has a mapping between I_{F_i} and F_i, named map_form. This mapping can guarantee that itemsets contain items mapped from the same feature, and their supersets should be pruned. With I_{F_i} and the mapping, EPSF divides the training data by class, mines EPs for each class, and stores the EPs in a candidate EP pool named CEP.

(2) *Updating*. Due to $F_i's$ inclusion, EPSF updates the feature pool *CMB(C)* by removing redundant features according to Proposition 4.5. If feature Y is removed from *CMB(C)*, we update *CEP* and map_form online by removing EPs in CEP generated from Y and the mapping between itemsets I_y and Y in map_form, respectively.
To update *CMB(C)*, EPSF checks all subsets within *CMB(C)* to re-examine the redundancy of each feature in *CMB(C)*. To improve this updating efficiency, we only validate the redundancy of each originally existing feature in *CMB(C)* by testing the subsets created by the inclusion of the new feature F_i.

4.2.2. Periodically Mining All EPs from the 1-Itemset EP Pool. With the current 1-itemset EP pool, we propose an offline step to periodically mine all EPs, for an EP classifier construction and maintenance. At this step, we can periodically compute and update emerging patterns from the 1-itemset EP pool that is picked by the online step. Although the emerging-pattern mining step cannot be made online in theory, this step can be conducted only periodically, and can be separated from the online step. In other words, when the classification model needs to be updated, an offline task of emerging-pattern mining can take place.

4.3. A Score Function for EP Classifiers

When applying EPs to classification, we get all the EPs of each class C_i in a training set. With the EPs for K classes, we derive K scores for a test instance T, one score per class, by feeding the EPs of each class into a scoring function. In this article, we use the score function based on information theory proposed by Zhang et al. [2000b] for classifying unlabeled instances, since this function is simpler and more efficient than the score function proposed by Dong et al. [1999] by avoiding computing the base score for each class. Zhang et al. [2000b] defined the score function of a test instance T by the following equation:

$$L(T|C_i) = -\sum_{k=1}^{|E_i|} \log_2 P(X_k|C_i),\, X_k \in E_i \textit{ and } X_k \in T, \tag{9}$$

where $|E_i|$ is the number of emerging patterns in the EP set E_i and X_k is an emerging pattern in E_i. The test instance T is assigned class label C_i when $L(T|C_i)$ is the minimum. Given an itemset X, $P(X)$ is approximately computed by the following equation:

$$P(X|C_i) = (|X \cap C_i| + 2|X|/|D|)/(|C_i| + 2), \tag{10}$$

where $|X \cap C_i|$ is the number of training instances belonging to class C_i and containing X, $|X|$ is the total number of training instances containing X, $|D|$ is the total number of training instances, and $|C_i|$ is the number of training instances for class C_i.

In addition, to ensure that we can always find a partition for an instance, all single-item itemsets of each class, whether they satisfy the given thresholds or not, are taken into account when Equation (9) is used to classify a test instance.

4.4. The EPSF Algorithm

To integrate online feature selection and EP mining, we propose the algorithm mining Emerging Patterns with Streaming Features (EPSF), as shown in Algorithm 1.

EPSF builds two pools online: a feature pool and a 1-itemset EP pool, and periodically computes and updates emerging patterns from the 1-itemset EP pool for an EP classification model construction and maintenance. As a new feature arrives, if it is added into the feature pool, EPSF transforms it into a set of itemsets, and mines 1-itemset EPs online, which are then added into the 1-itemset EP pool. As the dimensions are processed one by one, in order to quickly respond to this change, EPSF only mines 1-itemset EPs online for each feature available, and updates the current 1-itemset EP pool correspondingly with the change of the feature pool. Then EPSF periodically computes and updates emerging patterns from the 1-itemset EP pool.

As for the EPSF algorithm, in its early version (we call it Pre-EPSF) [Yu et al. 2012], the Pre-EPSF algorithm needs to check all subsets within *CMB(C)* to re-examine the redundancy of each feature in *CMB(C)* due to a new feature F_i's inclusion. In Step 28 of Algorithm 1, the EPSF algorithm in the current version validates the redundancy of each originally existing feature in *CMB(C)* by checking only the subsets created by

ALGORITHM 1: The EPSF Algorithm

```
Initialize the minimum support threshold α, growth rate threshold ρ, and CMB(C)={};
repeat
    Input a new feature F_k;
    /*Discard irrelevant features*/;
    if P(C|F_k) = P(C) then
        Discard F_k and goto step 41;
    end
    /*Remove redundant features*/;
    if ∃S ⊂ CMB(C) s.t. P(C|F_k, S) = P(C|S) then
        Go to step 41;
    end
    /*Add F_k into the current feature pool CMB(C)*/;
    CMB(C) = CMB(C) ∪ {F_k};
    /*Convert feature F_k into a set of itemsets*/;
    I_{F_k} = convert(F_k), I_{F_k} ∈ Dom(F_k);
    /*Map between I_{F_k} and F_k*/;
    map_form=mapping(F_k,I_{F_k});
    for i = 1 : |C| do
        /*|C| denotes the number of classes*/;
        /*Mine 1-itemset EPs for each class with the thresholds α and ρ*/;
        EP_i=mineEP(I_{F_k}, α, ρ);
        /*Add EP_i to the current EP pool CEP*/;
        CEP = CEP ∪ EP_i;
    end
    /*Update CMB(C)*/;
    for each feature Y within CMB(C) excluding F_k do
        /*Find S ⊂ CMB(C) containing F_k*/;
        if ∃S ⊂ CMB(C) s.t. P(C|Y, S) = P(C|S) then
            CMB(C) = CMB(C) − Y;
            /*Update CEP by removing EPs generated from feature Y*/;
            for each y ∈ I_y do
                if y ∈ CEP then
                    CEP = CEP − y;
                end
            end
            /*Update map_form*/;
            map_form = map_form(I_y);
        end
    end
    Periodically mine all EPs from CEP with map_form;
until No more features are available;
Classify unlabeled instances by the mined EPs using Equation (9);
```

the inclusion of the new feature F_i at each time point. By avoiding checking all subsets within *CMB(C)* at each time point, the revised EPSF significantly reduces the number of subsets that need to be checked, thus improves the updating efficiency by avoiding performing some unnecessary calculation.

Using the Balloon dataset from UCI Repository of Machine Learning Databases [Blake and Merz 1998], we present an illustrating example to explain the EPSF algorithm. In Table II, the Balloon dataset includes 4 features (*color*, *size*, *act* and *age*) and one class attribute (*inflated*) with 20 samples. Assuming the input order of features is *color*, *size*, *act*, and *age*, and the G^2 test [Spirtes et al. 2000] is employed to

Table II. The Balloon Dataset

Color	Size	Act	Age	Inflated
yellow	small	stretch	adult	True
yellow	small	stretch	child	True
yellow	small	dip	adult	True
yellow	large	stretch	adult	True
yellow	large	stretch	child	True
yellow	large	dip	adult	True
purple	small	stretch	adult	True
purple	small	stretch	child	True
purple	small	dip	adult	True
purple	large	stretch	adult	True
purple	large	stretch	child	True
purple	large	dip	adult	True
yellow	small	dip	child	False
yellow	small	dip	child	False
yellow	large	dip	child	False
yellow	large	dip	child	False
purple	small	dip	child	False
purple	small	dip	child	False
purple	large	dip	child	False
purple	large	dip	child	False

Table III. *CEP* After Adding *act* for Class *F*

Candidate EP	Support (class T)	Support (class F)	$GR_{T\rightarrow F}(e)$
{act=dip}	0.33	1	3

Table IV. *CEP* After Adding *act* for Class *T*

Candidate EP	Support (class F)	Support (class T)	$GR_{F\rightarrow T}(e)$
{act=stretch}	0	0.67	∞

compute conditional independence defined in Definition 3.5 in Section 3.2 to determine feature relevance and feature redundancy, the EPSF algorithm is traced as follows.

(1) As feature *color* arrives, at Step 6 in Algorithm 1, *color* is discarded as an irrelevant feature. EPSF then processes the next feature *size* directly. Since feature *size* is also independent of the class attribute *inflated*, *size* is also discarded, and will never be considered again.

(2) As feature *act* is available, at Step 5 in Algorithm 1, *act* is regarded as a relevant feature. Step 9 then checks whether *act* is a redundant feature given the current feature pool *CMB(C)*. If so, *act* will be discarded and EPSF will consider a next feature available; if not, *act* will be added to *CMB(C)*. Since the current feature pool *CMB(C)* is empty, *act* is added to *CMB(C)* at Step 13 and *CMB(C)* = $\{act\}$. At Steps 14 to 17, feature *act* is converted into a set of itemsets, that is, $\{act = dip\}$ and $\{act = stretch\}$. From Steps 18 to 24, EPSF mines 1-itemset EPs of *act* from two classes and stores those 1-itemsets into the current EP pool, *CEP*, as shown in Tables III and IV, using the minimum support threshold 0.2 and the GR threshold $\rho > 1$. Since the current feature pool *CMB(C)* contains only *act*, Steps 25 to 39 are not implemented and EPSF directly processes the next feature *age*.

(3) As feature *age* comes, *age* is considered a relevant feature at Step 5. At Step 9, given the current feature pool *CMB(C)* = $\{act\}$, $P(inflated|age, act) \neq P(inflated|act)$. Accordingly, *age* is added to *CMB(C)* at Step 13, and *CMB(C)* = $\{act, age\}$. At

Table V. *CEP* After Adding *age* for Class *F*

Candidate EP	Support (class T)	Support (class F)	$GR_{T\rightarrow F}(e)$
{act=dip}	0.33	1	3
{age=child}	0.33	1	3

Table VI. *CEP* After Adding *age* for Class *T*

Candidate EP	Support (class F)	Support (class T)	$GR_{F\rightarrow T}(e)$
{act=stretch}	0	0.67	∞
{age=adult}	0	0.70	∞

Table VII. 16 Datasets Used in our Comparative Study

ID	Dataset	#	Size	ID	Dataset	#	Size
1	kr-vs-kp	36	3,196	9	dexter	20,000	300
2	spectf	44	267	10	breast cancer	17,816	286
3	promoters	57	106	11	arcene	10,000	100
4	infant	86	5,337	12	dorothea	100,000	800
5	madelon	500	2,000	13	colon	2,000	62
6	hiva	1,617	4,229	14	leukemia	7,129	72
7	ovarian cancer	2,190	216	15	lung cancer	12,533	181
8	lymphoma	7,399	227	16	prostate	6,033	102

(#: Number of Features, Size: Number of Instances).

Steps 14 to 17, *act* is converted into a set of itemsets, that is, $\{age = child\}$ and $\{age = adult\}$. From Steps 18 to 24, the current EP pool *CEP* is updated as shown in Tables V and VI.

(4) Due to *age's* addition to *CMB(C)*, Steps 25 to 39 further check whether *act* is a redundant feature. If so, *act* will be removed from *CMB(C)*, and its corresponding 1-itemsets in *CEP* also will be removed.

(5) With the current EP pool *CEP*, EPSF periodically mines all EPs by employing a level-wise, candidate generation-and-test approach to mine EPs (we use the ConsEPMiner algorithm [Zhang et al. 2000a]), then uses them to classify test instances later on.

In summary, compared to the CE-EP algorithm and other existing EP algorithms, we are the first group to mine EPs from data with streaming features. With an effective online feature selection customized for emerging-pattern mining, the EPSF algorithm is designed for datasets with streaming features as it does not need to store the whole data in the memory to mine EPs. This facilitates emerging-pattern mining dramatically.

Moreover, EPSF can online mine EPs from the features available so far and can consume new features in an online manner as they become available. Accordingly, the EPSF algorithm allows more expensive calculation, including feature redundancy checking (step 9), emerging pattern mining (steps 14-24), *CMB(C)* and *CEP* updating (steps 25-39) all to be online conducted within the current *CMB(C)*, which is usually much smaller than the whole feature space.

5. EXPERIMENT RESULTS

5.1. Experiment Setup

In order to thoroughly evaluate the proposed EPSF algorithm, 16 datasets (in Table VII) are selected including four from the UCI Machine Learning Repository (the first four) [Blake and Merz 1998], four very high-dimensional biomedical datasets (*hiva*, *ovarian cancer*, *lymphoma*, and *breast cancer*), four NIPS 2003 feature selection challenge datasets (*madelon*, *arcene*, *dorothea*, and *dexter*), and four frequently studied public microarray datasets (the last four).

Our comparative study has the following systematical design, using 10-fold cross-validation for all the experiments unless specified.

(1) Comparing EPSF with the state-of-the-art EP classifiers, CE-EP [Yu et al. 2013], the EPSF algorithm in our previous KDD conference version (we call it Pre-EPSF) [Yu et al. 2012], and the IG-EP classifier, the EP classifier with the information gain feature selection method. (We do not compare EPSF with CAEP [Dong et al. 1999], CBA [Liu et al. 1998], CMAR [Li et al. 2001b] and CPAR [Yin and Han 2003] since they fail to deal with high dimensionality in the scale of thousands or more.)
(2) Comparing the prediction accuracy of EPSF with that of the state-of-the-art nonassociative classifiers, including Naïve Bayes (NB), KNN, Decision Tree J48, SVM, Bagging, and AdaBoost using their implementation provided by the Weka tool [Hall et al. 2009].
(3) Comparing the prediction accuracy of EPSF with that of NB, KNN, J48, SVM, Bagging, and AdaBoost classifiers with the add-on information gain feature selection method in Weka.
(4) Analyzing the statistical qualities of the EPSF algorithm against the rivals mentioned earlier using the kappa statistic [Cohen 1960], the Friedman test [Friedman 1940], and the Nemenyi test [Demšar 2006].

We simulate the streaming feature setting using benchmark datasets to evaluate EPSF and Pre-EPSF by assuming that the dimensions on a benchmark training dataset are available one at a time and each dimension is processed upon its arrival. To discretize continuous features, we use the discretization method in the Causal Explorer Toolkit proposed by Aliferis et al. [2003]. In the experiments, we set the GR to 20 for EPSF and CE-EP. To test the impact of the minimum support threshold, we set seven minimum supports for EPSF, Pre-EPSF, and CE-EP, including 0.005, 0.01, 0.05, 0.1, 0.2, 0.3, and 0.4, respectively. The experiments were performed on a Window 7 Dell workstation with an Intel Xeon 2.93GHz processor and 12.0GB RAM.

5.2. Comparison of EPSF and Pre-EPSF

Table VIII gives the results of computer runtime of EPSF against Pre-EPSF (the EPSF algorithm in our previous version [Yu et al. 2012]). Comparing to Pre-EPSF, for validating the redundancy of each originally existing feature in the current feature subset, EPSF in this article avoids checking all subsets within the current feature subset at each round by testing only the subsets created by the inclusion of a new feature at each time point. Thus, in comparison with Pre-EPSF, EPSF in this article significantly improves the updating efficiency, as shown in Table VIII. The best results are highlighted in boldface.

Moreover, since EPSF needs fewer statistical tests to determine the redundancy of a feature in the current feature subset than Pre-EPSF (hence less unintended estimation errors are introduced), Table IX shows that EPSF gets higher prediction accuracy on some datasets than Pre-EPSF, especially on high-dimensional datasets with small samples, such as four NIPS 2003 feature selection challenge datasets and four frequently studied public microarray datasets. In Table IX, we select the best prediction accuracy under the seven minimum supports as the results for our comparative study.

Finally, Figure 2 reports the kappa statistics of EPSF and Pre-EPSF. In Figure 2, the x-axis denotes all of the 16 datasets corresponding to Table VII. The kappa statistic is a measure of consistency among different raters, taking into account the agreement occurring by chance [Cohen 1960]. The kappa statistic is standardized to lie on a -1 to

Table VIII. Runtime (in Seconds): EPSF and Pre-EPSF

Dataset	EPSF	Pre-EPSF
infant	**26**	41
kr-vs.-kp	**24**	43
promoters	17	**16**
spectf	**17**	**17**
madelon	**20**	23
hiva	**33**	163
ovarian cancer	**20**	68
lymphoma	**20**	44
dexter	**31**	387
arcene	**19**	30
breast cancer	**101**	958
dorothea	**146**	440
colon	**17**	18
leukemia	**19**	22
lung cancer	**30**	117
prostate	**20**	27

(Best results in boldface).

Table IX. Prediction Accuracy (%): EPSF and Pre-EPSF

Dataset	EPSF	Pre-EPSF
infant	**91.44**	91.35
kr-v.s-kp	92.39	**92.42**
promoters	**72.00**	71.00
spectf	**86.92**	**86.92**
madelon	59.80	**61.20**
hiva	90.71	**95.17**
ovarian cancer	92.38	**93.81**
lymphoma	**80.91**	76.82
dexter	**90.67**	89.67
arcene	**84.44**	80.00
breast cancer	**95.19**	92.59
dorothea	**95.06**	94.94
colon	**95.00**	91.67
leukemia	**100**	**100**
lung cancer	**99.44**	98.89
prostate	**98.00**	95.00

(Best results in boldface).

Table X. Kappa Statistic and Its Corresponding Kappa Agreement

Kappa statistic	<0	0.01–0.20	0.21–0.40	0.61–0.80	0.81–0.99
Kappa agreement	less than chance agreement	slight agreement	moderate agreement	substantial agreement	almost perfect agreement

1 scale, where 1 is perfect agreement, 0 is exactly what would be expected by chance, and negative values indicate agreement less than chance. The other values of kappa statistics and their corresponding kappa agreements are shown in Table X [Landis and Koch 1977].

From Figure 2, EPSF is better than Pre-EPSF using kappa statistics, thus we can conclude that EPSF is more reliable than Pre-EPSF. The explanation is that fewer statistical tests of EPSF than those of Pre-EPSF make EPSF have more statistical

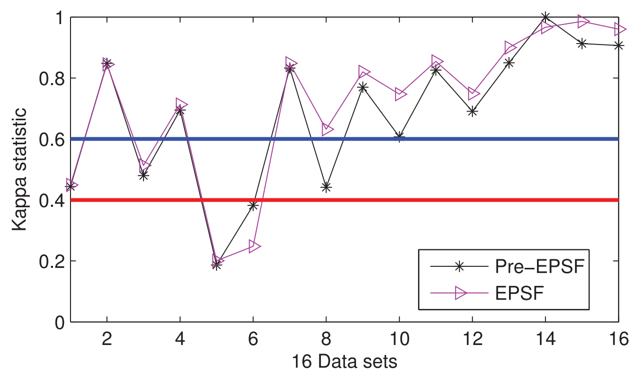

Fig. 2. Kappa statistics of Pre-EPSF and EPSF on 16 datasets.

Table XI. Prediction Accuracy (%): EPSF, CE-EP, and IG-EP

Dataset	EPSF	CE-EP	IG-EP
infant	91.44	**94.92**	91.61
kr-vs-kp	**92.39**	92.23	87.58
promoters	72.00	72.00	**75.00**
spectf	**86.92**	83.85	83.08
madelon	59.80	59.00	**60.80**
hiva	90.71	**93.70**	93.36
ovarian-cancer	92.38	**92.86**	83.33
lymphoma	**80.91**	77.73	78.18
dexter	**90.67**	88.33	79.33
arcene	84.44	**86.67**	68.89
breast-cancer	**95.19**	92.22	90.74
dorothea	**95.06**	**95.06**	93.92
colon	**95.00**	**95.00**	88.33
leukemia	**100**	**100**	**100**
lung-cancer	**99.44**	**99.44**	98.89
prostate	**98.00**	94.00	94.00
win/tie/loss	/	5/8/3	10/3/3

(Best results in boldface).

power than Pre-EPSF. We can see that both EPSF and Pre-EPSF have only two kappa statistics (the *madelon* and *hive* datasets) that are lower than 0.4 (under the red line in Figure 2), since the *madelon* dataset is a synthetic dataset including many redundant and noise features and *hive* is a very class-imbalanced dataset (the proportion of positive class is only 3.52%). Both have 11 kappa statistics that are higher than 0.6 (above the blue line in Figure 2). Pre-EPSF is a reliable emerging-pattern classifier, while the improved EPSF is more reliable and more efficient.

5.3. Comparison of EPSF with CE-EP and IG-EP

5.3.1. Comparison of Prediction Accuracy. Table XI reports detailed results in terms of prediction accuracy (the percentage of the correctly classified test instances that are previously unseen) of EPSF, CE-EP, and IG-EP on the 16 benchmark datasets. As for EPSF, CE-EP, and IG-EP, we select the best prediction accuracy under the seven minimum supports as the results for our comparative study. The best result is highlighted in boldface for each dataset.

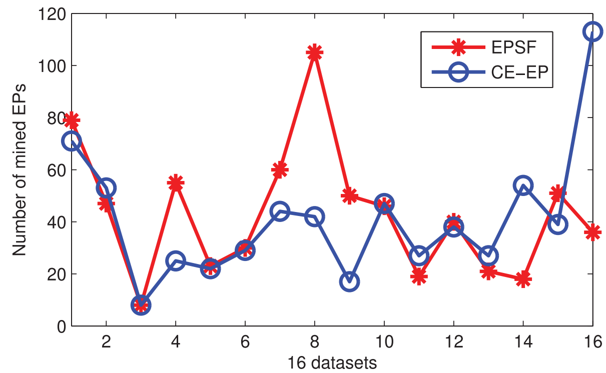

Fig. 3. Numbers of mined EPs: EPSF against CE-EP.

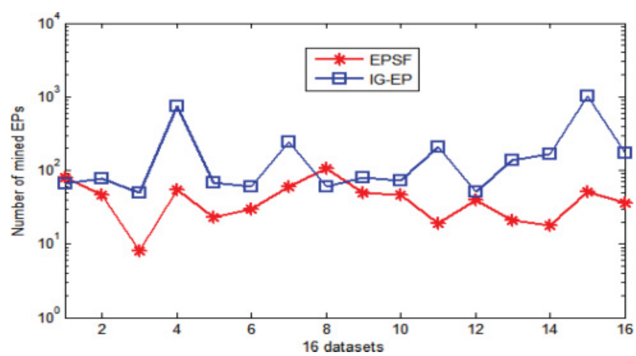

Fig. 4. Numbers of mined EPs: EPSF against IG-EP.

To further investigate the classification results, we conduct paired t-tests at a 95% significance level and summarize the win/tie/loss counts of EPSF against CE-EP and IG-EP. For example, as shown in the last row of Table XI, against CE-EP, EPSF wins five times, ties eight times, and loses three times on the 16 datasets. EPSF is always superior to or ties with IG-EP.

5.3.2. Comparison of the Numbers of Patterns and Runtime. Figures 3 and 4 compare the numbers of patterns mined by EPSF against CE-EP and IG-EP. We report the average numbers of mined patterns over all seven minimum support thresholds. In Figure 3, the x-axis denotes all of the 16 datasets corresponding to Table VII. From Figure 3, we can see that EPSF is very competitive with CE-EP on the number of mined patterns, while IG-EP selects more patterns than EPSF and CE-EP, as shown in Figure 4. These results illustrate that both EPSF and CE-EP can select a small set of strongly predictive EPs from a very high-dimensional dataset. Furthermore, we can see that even with very high dimensionality, the numbers of patterns selected by both EPSF and CE-EP do not change much in comparison with those on the first four low-dimensional datasets in Table VII.

The runtime (in seconds) of EPSF, CE-EP, and IG-EP contains all learning time, including importing datasets, and 10-fold cross-validation learning and testing. Figure 5 reports the average runtime over all seven minimum support thresholds. In Figure 5, we can see that EPSF is the fastest algorithm, while IG-EP is the slowest one.

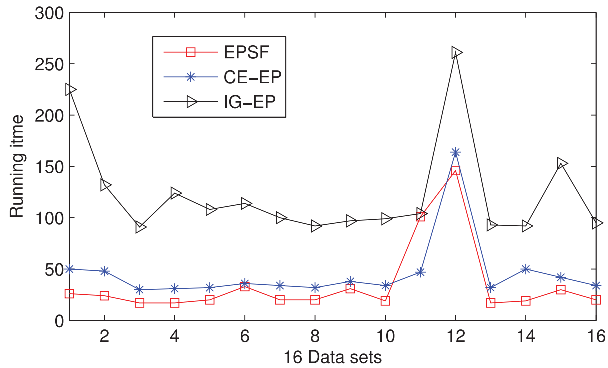

Fig. 5. Runtime (in seconds): EPSF, CE-EP, and IG-EP.

Table XII. Kappa Statistics of EPSF, CE-EP, and IG-EP

Dataset	EPSF	CE-EP	IG-EP
infant	0.4493	**0.7875**	0.4774
kr-vs.-kp	**0.8448**	**0.8448**	0.7508
promoters	0.5133	0.5133	**0.5400**
spectf	**0.7131**	0.6476	0.6306
madelon	0.2007	0.1866	**0.2191**
hiva	0.2477	**0.3497**	0.2311
ovarian cancer	0.8485	**0.8621**	0.7077
lymphoma	**0.6318**	0.5818	0.5864
dexter	**0.8200**	0.7742	0.5983
arcene	0.7467	**0.7800**	0.4867
breast cancer	**0.8548**	0.8197	0.7888
dorothea	**0.7488**	0.7362	0.6381
colon	0.9000	**0.9167**	0.7833
leukemia	0.9667	**1.0000**	**1.0000**
lung cancer	**0.9857**	**0.9857**	0.9714
prostate	**0.9600**	0.8733	0.8800

(Best results in boldface).

5.3.3. Analysis of the Statistical Qualities. To further analyze EPSF, CE-EP, and IG-EP, we compare them by the kappa statistic and Nemenyi test. To calculate the kappa statistic, we set the support threshold to 0.2 and the growth rate threshold to 20 for EPSF, CE-EP, and IG-EP. Table XII shows the kappa statistics of EPSF, CE-EP, and IG-EP. We observe that with the kappa statistics, the three classifiers have no significant difference according to Table X.

Accordingly, we further use the Friedman test [Friedman 1940] and Nemenyi test [Demšar 2006] to assess whether the performance of our algorithm EPSF is comparable to that of CE-EP and IG-EP in prediction accuracy. With the Friedman test at 95% significance level, under the null hypothesis, which states that whether the performance of EPSF and that of CE-EP and IG-EP have no significant difference in prediction accuracy, the null hypothesis is rejected. We get the average ranks for EPSF, CE-EP, and IG-EP as 2.3125, 2.1563, and 1.5313, respectively.

Then we proceed with the Nemenyi test as a post-hoc test to deal with this situation. With the Nemenyi test, the performance of the two classifiers is significantly different if the corresponding average ranks differ by at least the critical difference (for information on how to calculate the average ranks and critical difference, see Section 3.2.2

Table XIII. Comparison of Prediction Accuracy (%): EPSF, NB, IG-NB, KNN, IG-KNN, J48, and IG-J48

Dataset	EPSF	NB	IG-NB	KNN	IG-KNN	J48	IG-J48
infant	91.44	91.91	92.69	94.92	95.05	95.39	**95.43**
kr-vs.-kp	92.39	83.92	85.89	96.46	96.37	**99.31**	96.75
promoters	72.00	**74.53**	71.70	62.26	65.09	63.21	70.75
spectf	86.92	86.63	83.90	84.27	86.52	86.14	**87.27**
madelon	59.80	59.20	**62.30**	53.55	61.95	57.50	62.20
hiva	90.71	87.06	94.16	96.50	96.38	96.39	**96.62**
lymphoma	**80.91**	68.28	80.61	63.88	73.57	71.81	70.93
breast cancer	**95.19**	93.01	89.96	86.36	90.55	80.77	85.66
ovarian cancer	**92.38**	70.83	84.26	85.19	88.89	91.67	89.35
dorothea	**95.06**	90.25	93.63	90.63	93.38	89.38	93.13
arcene	**84.44**	63.00	73.00	80.00	79.00	62.00	77.00
dexter	90.67	**93.33**	78.33	63.67	86.00	82.67	87.00
colon	**95.00**	79.03	91.94	83.87	93.55	82.26	90.32
leukemia	**100**	93.06	**100**	97.22	**100**	93.06	95.83
lung cancer	**99.44**	98.34	98.90	98.34	98.34	90.61	96.69
prostate	**98.00**	69.61	94.12	93.14	94.12	88.24	93.14

(Best results in boldface).

Table XIV. Comparison of Prediction Accuracy (%): EPSF, SVM, IG-SVM, Bagging, IG-Bagging, AdaBoost, and IG-AdaBoost

Dataset	EPSF	SVM	IG-SVM	Bagging	IG-Bagging	AdaBoost	IG-AdaBoost
infant	91.44	95.45	95.48	**95.65**	95.51	95.43	95.43
kr-vs.-kp	92.39	95.06	94.02	**99.22**	96.25	93.84	93.84
promoters	72.00	**79.25**	70.75	66.98	72.64	66.04	72.64
spectf	86.92	88.02	89.89	**90.37**	87.64	84.27	85.39
madelon	59.80	56.45	**62.75**	62.20	62.45	60.50	60.7
hiva	90.71	94.70	96.26	**96.76**	96.57	96.48	96.48
lymphoma	**80.91**	77.53	79.30	68.28	78.85	62.56	68.72
breast cancer	**95.19**	92.31	90.21	84.97	88.81	84.61	87.41
ovarian cancer	92.38	**93.52**	91.67	88.89	87.96	91.67	89.35
dorothea	**95.06**	92.00	94.00	94.13	93.75	93.75	93.75
arcene	**84.44**	81.00	74.00	72.00	78.00	71.00	79.00
dexter	90.67	**91.33**	86.33	89.33	88.67	83.33	85.00
colon	**95.00**	85.48	88.71	85.48	85.48	85.48	91.94
leukemia	**100**	98.61	**100**	94.44	97.22	**100**	**100**
lung cancer	99.44	**100**	**100**	93.92	99.45	96.69	99.45
prostate	**98.00**	94.12	94.12	92.16	95.10	92.16	94.12

(Best results in boldface).

of Demšar [2006]). With the Nemenyi test, the critical difference we get is up to 0.8275.

Thus, with the critical difference and the average ranks calculated here, we conclude that the performance of EPSF and that of CE-EP have no significant difference in prediction accuracy, but are significantly better than that of IG-EP.

5.4. Comparison of EPSF against the Nonassociative Classifiers

Tables XIII through XIV give the empirical results in terms of prediction accuracy of EPSF, six nonassociative classifiers, and the same six classifiers with the add-on information gain feature selection method on the 16 benchmark datasets. We denote the six classifiers with information gain feature selection as IG-NB, IG-KNN, IG-J48, IG-SVM, IG-Bagging, and IG-AdaBoost, respectively. Table XV reports the runtime of

Table XV. Comparison of Runtime (in seconds): EPSF, NB, KNN, J48, SVM, Bagging, and AdaBoost

Dataset	EPSF	NB	KNN	J48	SVM	Bagging	AdaBoost
infant	26	5	5	10	59	20	8
kr-vs.-kp	24	5	5	5	10	5	5
promoters	17	5	5	5	5	5	5
spectf	17	5	5	5	5	5	5
madelon	20	1	21	14	770	50	21
hiva	33	5	95	185	269	430	43
lymphoma	20	6	6	6	6	15	6
breast cancer	101	10	10	21	33	51	26
ovarian cancer	20	5	5	5	5	5	5
dorothea	146	70	975	675	670	1425	625
arcene	19	1	1	1	5	14	3
dexter	31	5	5	52	27	94	17
colon	17	5	5	5	5	5	5
leukemia	19	5	6	6	10	10	10
lung cancer	30	5	5	5	10	15	10
prostate	20	5	5	5	10	10	10

Table XVI. Win/Tie/Loss Counts of EPSF Versus the Other 12 Nonassociative Classifiers

	NB	KNN	J48	SVM	Bagging	AdaBoost
EPSF	11/3/2	13/0/3	11/2/3	8/2/6	10/2/4	10/3/3
	IG-NB	IG-KNN	IG-J48	IG-SVM	IG-Bagging	IG-AdaBoost
EPSF	9/4/3	10/2/4	11/1/4	8/3/5	9/3/4	9/4/3

(Pairwise t-test at 95% Significance Level).

EPSF, NB, KNN, J48, SVM, Bagging, and AdaBoost. EPSF is very competitive with these six nonassociative classifiers. But on datasets with very high dimensionality or large sample sizes—such as the *madelon*, *hiva*, and *dorothea* datasets—the runtime of most nonassociative classifiers is more than that of EPSF.

To investigate the classification results of prediction accuracy, we conduct paired t-tests at a 95% significance level and summarize the win/tie/loss counts of EPSF against the other rivals in Table XVI. In Table XVI, we can see that EPSF is superior to NB, KNN, J48, Bagging, and AdaBoost and their variants with the information gain feature selection method, and also very competitive with SVM and IG-SVM. We have observed in the experiments that the integration of the streaming feature selection into EP mining can avoid generating non-EPs and redundant EPs. This enables EPSF not only to handle high-dimensional datasets such as the last 12 datasets in Table VII, but also to produce very promising prediction accuracy.

To further analyze prediction accuracies of these classifiers, we use the Friedman test and Nemenyi test to assess their performance. With the Friedman test at 95% significance level for EPSF, NB, KNN, J48, SVM, Bagging, and AdaBoost, the null-hypothesis is also rejected. The average ranks of EPSF, NB, KNN, J48, SVM, Bagging and, AdaBoost are 5.4063, 3.0313, 3.1875, 2.75, 5.375, 4.5, and 3.75, respectively. After the Nemenyi test, the critical difference is up to 2.252. Therefore, we can conclude that the performance of EPSF is significantly better than that of NB and J48, but is highly comparable to that of KNN, SVM, Bagging, and AdaBoost.

Finally, with the Friedman test at 95% significance level for EPSF, IG-NB, IG-KNN, IG-J48, IG-SVM, IG-Bagging, and IG-AdaBoost, the null hypothesis cannot be rejected. Thus, the performance of EPSF has no significant difference from that of the six nonassociative classifiers using the information gain feature selection method.

To further analyze the statistical qualities of EPSF, we compare EPSF with IG-NB, IG-KNN, IG-J48, IG-SVM, IG-Bagging, and IG-AdaBoost by the kappa statistics and

Table XVII. Comparison of Kappa Statistics: EPSF Against 6 Associative Classifiers with Information Gain Feature Selection Method

Dataset	EPSF	IG-NB	IG-KNN	IG-J48	IG-SVM	IG-Bagging	IG-AdaBoost
infant	0.4493	0.4548	0.4328	0.5019	0.4851	0.5012	**0.5174**
kr-vs.-kp	0.8448	0.7165	0.9273	**0.9348**	0.8800	0.9248	0.8762
promoters	**0.5133**	0.4340	0.3019	0.4151	0.4151	0.4528	0.4528
spectf	**0.7131**	0.5528	0.5933	0.5886	0.6845	0.5797	0.5314
madelon	0.2007	0.2460	0.2390	0.2440	**0.2550**	0.2490	0.2140
hiva	**0.2477**	0.2327	0.0663	0.0996	0.0077	0.0731	0.0596
lymphoma	**0.6318**	0.6124	0.4717	0.4185	0.5859	0.5571	0.3743
breast cancer	**0.8548**	0.7515	0.7668	0.6281	0.7470	0.7035	0.6853
ovarian cancer	**0.8485**	0.6862	0.7765	0.7837	0.8305	0.7552	0.7827
dorothea	**0.7488**	0.6181	0.5425	0.5502	0.6100	0.5884	0.6091
arcene	**0.7467**	0.4664	0.5728	0.5298	0.4698	0.5557	0.5770
dexter	**0.8200**	0.5567	0.7200	0.7400	0.7267	0.7733	0.7000
colon	**0.9000**	0.8256	0.8561	0.7842	0.7456	0.6729	0.8220
leukemia	0.9667	**1.0000**	**1.0000**	0.9089	**1.0000**	0.9376	**1.0000**
lung cancer	0.9857	0.9620	0.9438	0.8861	**1.0000**	0.9808	0.9808
prostate	**0.9600**	0.8825	0.8825	0.8627	0.8824	0.9020	0.8823

(Best results in boldface).

Nemenyi test. Since the prediction accuracy of IG-NB, IG-KNN, IG-J48, IG-SVM, IG-Bagging, and IG-AdaBoost is better than NB, KNN, J48, SVM, Bagging, and AdaBoost, we do not give the kappa statistics of NB, KNN, J48, SVM, Bagging, and AdaBoost. We can see that EPSF gets higher kappa statistics than the other six classifiers in Table XVII, especially on the class-imbalance datasets, such as *hiva* and *dorothea*, or the datasets with high dimensionality but small sample sizes, such as *lymphoma* and *prostate*. A possible explanation is that the emerging patterns of each class are correctly mined by EPSF from the corresponding class data and represent strong contrasts between different classes of data.

Why is the prediction accuracy of EPSF not better than that of IG-EP and 12 nonassociative classifiers on the low-dimensional datasets, such as *infant*, *kr-vs.-kp*, *promoters*, *spectf*, *madelon*, and *hiva*, while it is better than that of these algorithms on the remaining high-dimensional datasets? The explanation is that a high-dimensional dataset would have a better chance of including excessive irrelevant or redundant features than a low-dimensional dataset. Those excessive irrelevant or redundant features might significantly reduce performance of predictive models. Thus, on a high-dimensional dataset, the adverse impact of irrelevant or redundant features on predictive models is more significant than that on low-dimensional datasets. Our empirical results reveal that EPSF can deal with irrelevant or redundant features in high-dimensional datasets much better than the other rivals.

In summary, we can conclude that for datasets with streaming features, the performance of EPSF is very competitive with that of CE-EP and is better than that of IG-EP, which both need to obtain a complete set of features in advance. Furthermore, in comparison with the six nonassociative classifiers and the six classifiers with the information gain feature selection method, the prediction accuracy of EPSF is also very competitive with that of these 12 nonassociative classifiers.

5.5. Analysis of Prediction Accuracy on Support Thresholds

Figure 6 shows the prediction accuracy of EPSF and CE-EP under the seven support thresholds, 0.005, 0.01, 0.05, 0.1, 0.2, 0.3, and 0.4. We can see that in prediction accuracy, for all 16 datasets, EPSF is insensitive to the different support thresholds, even for

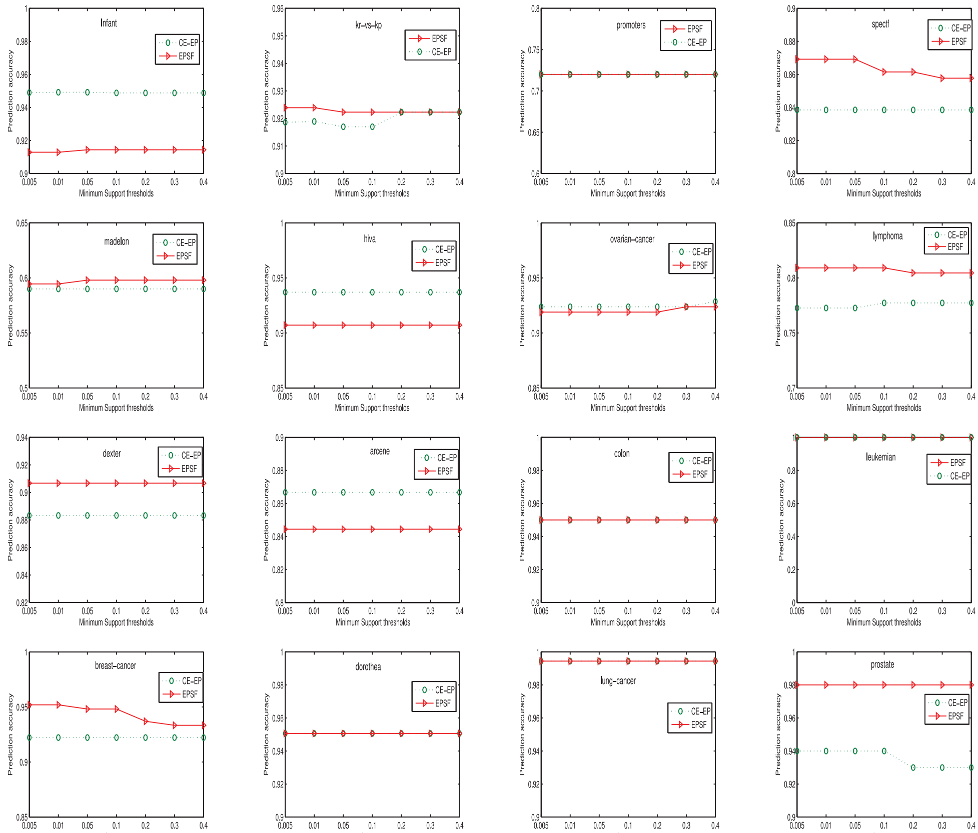

Fig. 6. Sensitivity analysis of support thresholds on prediction accuracy.

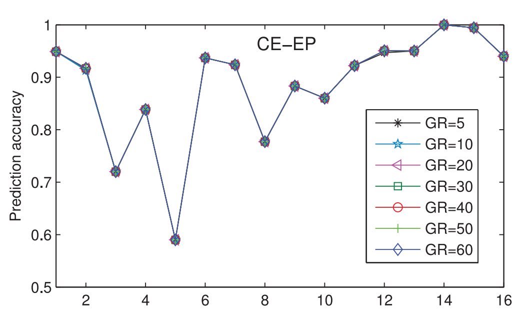

Fig. 7. The effect of GR thresholds on CE-EP.

those high-dimensional datasets. Furthermore, EPSF is not only more insensitive, but also always achieves higher accuracy under all seven support thresholds than CE-EP.

5.6. Analysis of Prediction Accuracy on Growth-Rate Thresholds

To further explore the performance of EPSF and CE-EP, we conduct an analysis on prediction accuracy of EPSF and CE-EP under seven minimum GR thresholds, as shown in Figures 7 and 8, where *GR* stands for growth rate thresholds and the minimum

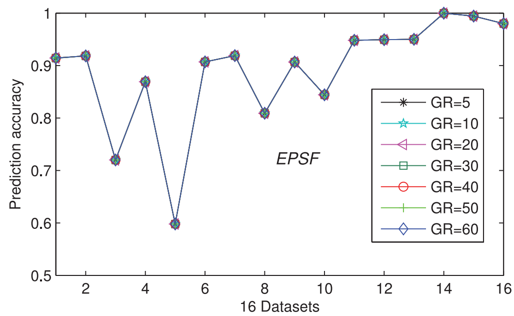

Fig. 8. The effect of GR thresholds on EPSF.

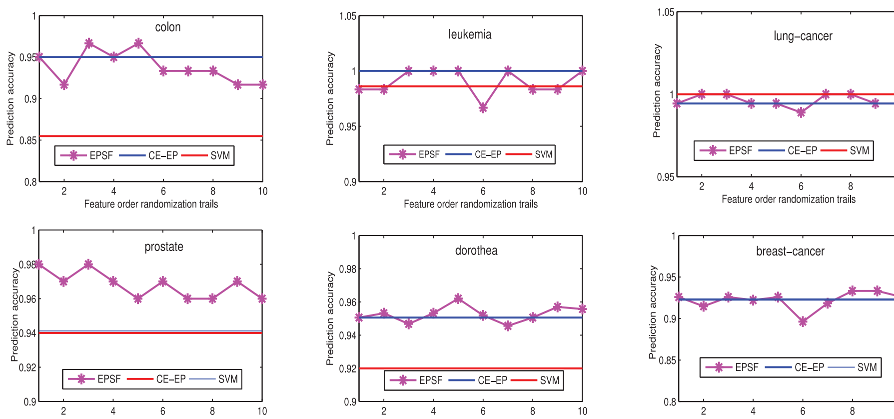

Fig. 9. Effect of the input order of features.

support threshold is fixed at 0.1. In Figures 7 and 8, the x-axis denotes all 16 datasets corresponding to Table VII. From Figures 7 and 8, we can see that CE-EP and EPSF are not sensitive to the minimum GR thresholds at all.

5.7. Effect of the Input Orders of Features

Since streaming features are processed one by one as they are available, we conduct an analysis of prediction accuracy on the input (or scan) order of features, against CE-EP and SVM as the rival algorithms. We generate a number of trials, each trial representing a random input order of features. We apply EPSF to each randomized trial and report the results in Figure 9, in which the x-axis represents each of the randomized trials and the y-axis represents the prediction accuracy from the corresponding trial. The results in Figure 9 confirm that varying the input order of features does slightly impact on the prediction accuracy; however, the results demonstrate that EPSF has a relatively stable performance.

5.8. Mining EPs with Features That Keep Arriving

When the features keep arriving, EPSF provides a solution to this problem by processing features one by one and stopping this process using the EPs seen so far with a user-specified criterion. CE-EP cannot deal with this situation since it needs to access all features in advance to identify the causes and effects of the class attribute. We evaluate this performance of EPSF in Figure 10. For four gene datasets, we select the

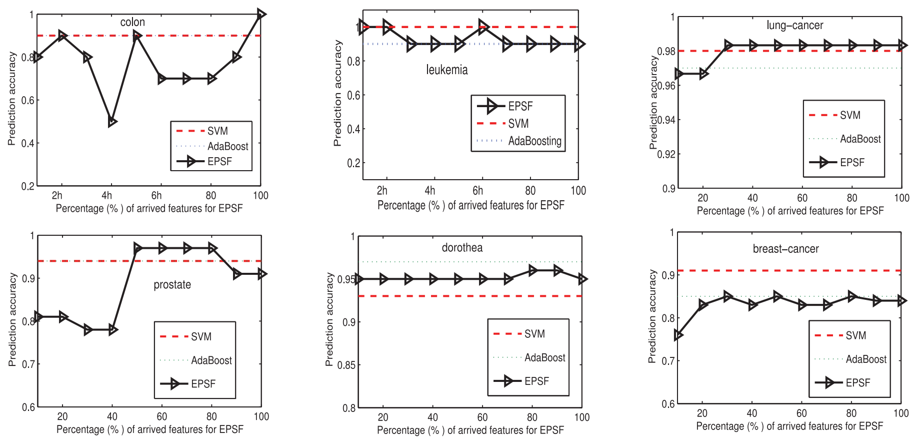

Fig. 10. Comparative performance of EPSF with features that keep arriving.

first 2/3 data instances as the training instances and the remaining for testing; for the *breast cancer* dataset, we select the first 200 data instances as the training instances and the remaining for testing. With respect to the *dorothea* dataset, we use its original training and testing datasets. SVM and AdaBoost are used as baselines on the training and testing sets with a complete set of features. With streaming features, EPSF mines EPs on the training samples as the features are available one by one and evaluates the current EPs on the testing samples.

On the *colon* dataset, when the percentage of features available is up to 20% or 50%, the prediction accuracy of EPSF is the same as SVM. When features are all available, the accuracy of EPSF is up to 100%, and is better than SVM. For the remaining datasets, EPSF can be also up to the accuracy of SVM or AdaBoost without exhaustive search over an entire feature set. This demonstrates that EPSF provides an effective and efficient solution to the EP mining problem when it is impossible to get a complete set of features in advance and must be consumed features in an online manner.

5.9. A Case Study on Automatic Impact Crater Detection

In addition to the validation of the publicly available benchmark datasets, we also apply our new approach to automatic impact crater detection in real planetary images. Impact craters, the structures formed by the collisions of meteoroids on planetary surfaces, are among the most studied geomorphic features in the solar system because they yield information about past and present geological processes and provide the only tool for measuring relative ages of planetary surfaces, that is, heavily cratered surfaces are relatively older than less cratered surfaces [Urbach and Stepinski 2009; Ding et al. 2011].

In this case study, the EPSF algorithm is plugged into the crater detection framework designed by Ding et al. [2011]. The calculation contains three steps: (1) identifying crater candidates; (2) extracting image texture features; and (3) detecting craters using supervised learning algorithms.

Crater candidates are the regions of an image that may potentially contain craters. A key insight into identifying crater candidates is that a crater can be recognized as a pair of crescent-like highlight and shadow regions in an image, as shown in Figure 11. Those highlight and shadow regions are matched so that each pair will be used to construct crater candidates, that is, the locations where craters are likely to reside. The experiments in crater detection are evaluated on Mars because it is at the center

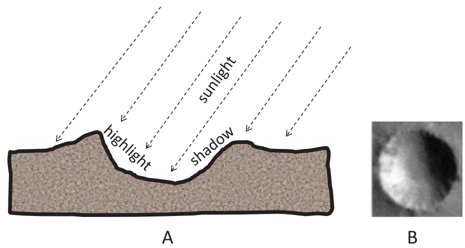

Fig. 11. (A) An illustration explaining why an image of a subkilometer crater consists of crescent-like highlight and shadow regions. (B) An image of an actual 1km crater showing the highlight and shadow regions [Ding et al. 2011].

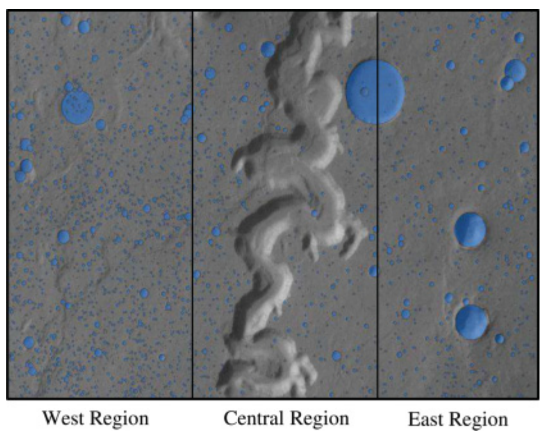

Fig. 12. Impact craters in a 37,500×56,250m^2 test image from Mars [Ding et al. 2011].

Table XVIII. Summary of Crater Datasets

	# Samples (Crater Candidates)	# Features
West region	6,708	1,089
Central region	2,935	1,089
East region	2,026	1,089

of NASA exploration efforts. A portion of the High Resolution Stereo Camera (HRSC) nadir panchromatic image h0905 is selected, taken by the Mars Express spacecraft, to serve as the test set [Ding et al. 2011]. The selected image has a resolution of 12.5 meters/pixel and a size of 3,000 by 4,500 pixels (37,500×56,250m^2). The image represents a significant challenge to automatic crater detection algorithms because it covers a terrain that has spatially variable morphology and because its contrast is rather poor (mostly noticeable when the image is inspected at a small spatial scale).

As the image arrives, it is divided into three sections denoted as the west region, the central region, and the east region (Figure 12) for the test sets summarized in Table XVIII. The central region is characterized by surface morphology that is distinct from the rest of the image. The west and east regions have similar morphology, but the west region is much more heavily cratered than the east region. A total of 1,089 image texture features are constructed. The training set consists of 204 true craters and 292 noncrater examples selected randomly from crater candidates located in the northern half of the east region.

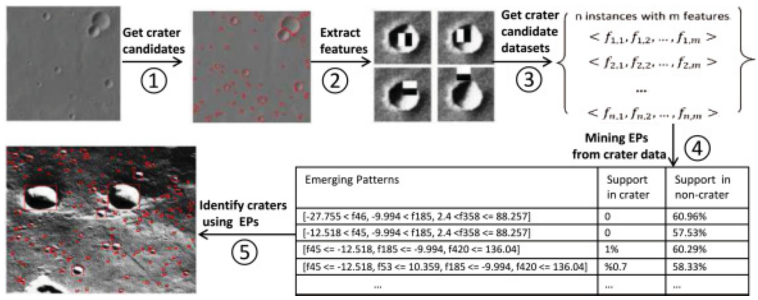

Fig. 13. Emerging patterns for crater detection.

Table XIX. Top 5 Emerging Patterns for Craters

ID	Emerging Patterns for Craters	Support in Noncraters	Support in Craters	Growth Rate
1	{f45∈[-138.46, -12.52], f46∈[-165.18, -27.76], f185∈[-102.58, -9.99], f420∈[-39.82,136.04]}	0.34%	59.31%	173.1961
2	{f45∈[-138.46, -12.52], f46∈[-165.18, -27.76], f68∈[-140.53, 2.97], f420∈[-39.82,136.04]}	0.68%	62.75%	91.6078
3	{f45∈[-138.46, -12.52], f46∈[-165.18, -27.76], f358∈[-113.51, 2.4], f420∈[-39.82,136.04]}	0.68%	62.25%	90.8922
4	{f45∈[-138.46, -12.52], f53∈[-155.16, 10.36], f185∈[-102.58, -9.99], f420∈[-39.82,136.04]}	0.68%	58.33%	85.1667
5	{f45∈[-138.46, -12.52], f185∈[-102.58, -9.99], f420∈[-39.82,136.04]}	1%	60.29%	58.6863

5.9.1. Emerging Patterns for Crater Detection. With the crate datasets summarized earlier, the framework for counting craters by emerging patterns is shown in Figure 13. In Steps 4 and 5 of Figure 13, we apply the EPSF and CE-EP algorithms to high-dimensional crater data for counting craters. Tables XIX and XX show the top 5 EPs mined from the crater training datasets for the crater class and noncrater class using the EPSF algorithm, respectively. In Tables XIX and XX, f45 denotes the 45th feature in the training crate dataset while [–138.46, –12.52] represents the value of feature f45. In Table XIX, the supports of the EPs in the crater class are much larger than in the noncrater class. An instance containing one of those EPs will favor the crater class. In Table XX, the first three EPs are the jumping EPs of the noncrater class, and denote the instances containing those EPs as the noncrater class. Accordingly, we conclude that the EPs mined by the EPSF algorithm are high-quality patterns and possess the most discriminative power. They are the best candidates to be used to construct a highly accurate classifier and also can produce an understandable classifier for crater data.

5.9.2. Comparison with Existing Crater Detection Methods. In this section, we compare EPSF with the state-of-the-art crater detection algorithms, CE-EP and Nave Boost [Ding et al. 2011]. The best results are boldfaced in Table XXI. Table XXI shows that, except for the west region, EPSF outperforms the CE-EP and Nave Boost algorithms.

Table XX. Top 5 Emerging Patterns for Noncraters

ID	Emerging Patterns for Noncraters	Support in Crater	Support in noncrater	Growth rate
1	{f46∈[-27.76,144.03], f185∈[-9.99,120.97], f358∈[2.4, 88.26]}	0	60.96%	∞
2	{f45∈[-12.52,158.41], f185∈[-9.99,120.97], f358∈[2.4, 88.26]}	0	57.53%	∞
3	{f68∈[2.97, 140.72], f185∈[-9.99,120.97], f358∈[2.4, 88.26]}	0	56.16%	∞
4	{[f53∈[10.36, 154.01], f358∈[2.4, 88.26]}	9.8%	61.30%	62.55
5	{f45∈[-12.52,158.41], f53∈[10.359,154.01], f185∈[-9.99,120.97], f420∈[136.04,276.73]}	1.47%	66.10%	44.95

Table XXI. Prediction Accuracy on Three Regions

	West Region	Central region	East region
EPSF	0.7847	**0.7959**	**0.7784**
CE-EP	**0.7852**	0.7802	0.7739
Nave Boost	0.7661	0.7888	0.7749

(Best results in boldface).

Table XXII. Prediction Accuracy on Three Regions (KNN)

	West Region	Central Region	East Region
EPSF	**0.7847**	**0.7959**	0.7784
OSFS	0.7809	0.7874	**0.7828**
HITON-PC	0.7749	0.7792	0.7813
LARS	0.7740	0.7881	0.7799
FCBF	0.7821	0.7833	**0.7828**
All features(1089)	0.7303	0.7499	0.7710

(Best results in boldface).

5.9.3. Comparison with the Other Methods. In this section, we compare the prediction accuracy of EPSF with that of the classifiers, KNN and SVM, with some state-of-the-art feature selection algorithms, OSFS [Wu et al. 2013], HITON-PC [Aliferis et al. 2010], LARS [Efron et al. 2004], and FCBF [Yu and Liu 2004]. The best results are boldfaced in Tables XXII and XXIII. From Tables XXII and XXIII, we can see that EPSF produces higher prediction accuracy than the other four algorithms in the central region and gets very competitive results with the other rivals in the remaining regions. Moreover, the classifier constructed with the mined EPs can help us understand the crated data.

6. CONCLUSIONS

In this article, to learn and maintain a classification model on data with streaming features, we have adapted the well-known emerging-pattern–based classification methods and proposed a semi-streaming approach. This new approach is fundamentally different from applying emerging-pattern mining in a straightforward manner on a dataset with streaming features. With streaming-feature selection, our approach online builds two pools: a feature pool and a 1-itemset EP pool, and periodically computes and updates emerging patterns from the 1-itemset EP pool for classification model con-

Table XXIII. Prediction Accuracy on Three Regions (SVM)

	West Region	Central Region	East Region
EPSF	0.7847	**0.7959**	0.7784
OSFS	**0.7856**	0.7874	0.7730
HITON-PC	0.7815	0.7877	0.7710
LARS	0.7840	0.7888	**0.7794**
FCBF	0.7826	0.7923	**0.7794**
All features(1089)	0.7683	0.7710	0.7754

(Best results in boldface).

struction and maintenance. The streaming feature selection step substantially reduces the dimensionality of the feature space under which the offline emerging-pattern mining step has to operate. Due to the effective streaming feature selection customized for emerging-pattern mining, the emerging patterns mined in the offline step tend to be short, and this practically facilitates emerging-pattern mining dramatically. Comprehensive experimental results on benchmark datasets and a real-world case study on automatic impact crater detection have demonstrated the effectiveness and efficiency of our approach.

ACKNOWLEDGMENTS

We thank the handling editor and the anonymous reviewers for their invaluable and constructive comments, which helped to improve the article substantially. This work is partly supported by a PIMS Post-Doctoral Fellowship Award of the Pacific Institute for the Mathematical Sciences, Canada, the Program for Changjiang Scholars and Innovative Research Team in University (PCSIRT) of the Ministry of Education, China (under grant IRT13059), the National 973 Program of China (under grant 2013CB329604), the National Natural Science Foundation of China (under grant 61229301 and 61305064), an NSERC Discovery grant and a BCIC NRAS Team Project.

All opinions, findings, conclusions, and recommendations in this article are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

Charu C. Aggarwal. 2010. Data streams: An overview and scientific applications. In *Scientific Data Mining and Knowledge Discovery*. Springer, 377–397.

Constantin F. Aliferis, Alexander Statnikov, Ioannis Tsamardinos, Subramani Mani, and Xenofon D. Koutsoukos. 2010. Local causal and Markov blanket induction for causal discovery and feature selection for classification Part I: Algorithms and empirical evaluation. *Journal of Machine Learning Research* 11, 171–234.

Constantin F. Aliferis, Ioannis Tsamardinos, Alexander R. Statnikov, and Laura E. Brown. 2003. Causal explorer: A causal probabilistic network learning toolkit for biomedical discovery. In *METMBS*, Vol. 3. 371–376.

C. L. Blake and Christopher J. Merz. 1998. UCI Repository of Machine Learning databases [http://www.ics.uci.edu/~mlearn/MLRepository.html]. Irvine, CA: University of California. *Department of Information and Computer Science* 55.

Gavin Brown, Adam Pocock, Ming-Jie Zhao, and Mikel Luján. 2012. Conditional likelihood maximisation: A unifying framework for information theoretic feature selection. *Journal of Machine Learning Research* 13, 1, 27–66.

Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* 20, 1, 37–46.

Janez Demšar. 2006. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research* 7, 1–30.

Wei Ding, Tomasz F. Stepinski, Yang Mu, Lourenco Bandeira, Ricardo Ricardo, Youxi Wu, Zhenyu Lu, Tianyu Cao, and Xindong Wu. 2011. Subkilometer crater discovery with boosting and transfer learning. *ACM Transactions on Intelligent Systems and Technology (TIST)* 2, 4, 39.

Guozhu Dong, Jiawei Han, Laks V. S. Lakshmanan, Jian Pei, Haixun Wang, and Philip S. Yu. 2003. Online mining of changes from data streams: Research problems and preliminary results. In *Proceedings of the 2003 ACM SIGMOD Workshop on Management and Processing of Data Streams*.

Guozhu Dong and Jinyan Li. 1999. Efficient mining of emerging patterns: Discovering trends and differences. In *Proceedings of the 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 43–52.

Guozhu Dong, Xiuzhen Zhang, Limsoon Wong, and Jinyan Li. 1999. CAEP: Classification by aggregating emerging patterns. In *Discovery Science*. Springer, 30–42.

Lei Duan, Guanting Tang, Jian Pei, James Bailey, Guozhu Dong, Akiko Campbell, and Changjie Tang. 2014. Mining contrast subspaces. In *Advances in Knowledge Discovery and Data Mining*. Springer, 249–260.

Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. 2004. Least angle regression. *The Annals of Statistics* 32, 2, 407–499.

Hongjian Fan and Kotagiri Ramamohanarao. 2006. Fast discovery and the generalization of strong jumping emerging patterns for building compact and accurate classifiers. *IEEE Transactions on Knowledge and Data Engineering,* 18, 6, 721–737.

Gang Fang, Gaurav Pandey, Wen Wang, Manish Gupta, Michael Steinbach, and Vipin Kumar. 2012. Mining low-support discriminative patterns from dense and high-dimensional data. *IEEE Transactions on Knowledge and Data Engineering,* 24, 2, 279–294.

Milton Friedman. 1940. A comparison of alternative tests of significance for the problem of m rankings. *Annals of Mathematical Statistics* 11, 1, 86–92.

Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The WEKA data mining software: An update. *ACM SIGKDD Explorations Newsletter* 11, 1, 10–18.

Daphne Koller and Mehran Sahami. 1995. Toward optimal feature selection. In *Proceedings of the 13th International Conference on Machine Learning*. 284–292.

J. Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics,* 159–174.

Jinyan Li, Guozhu Dong, and Kotagiri Ramamohanarao. 2000. Instance-based classification by emerging patterns. In *Principles of Data Mining and Knowledge Discovery*. Springer, 191–200.

Jinyan Li, Guozhu Dong, and Kotagiri Ramamohanarao. 2001a. Making use of the most expressive jumping emerging patterns for classification. *Knowledge and Information Systems* 3, 2, 131–145.

Wenmin Li, Jiawei Han, and Jian Pei. 2001b. CMAR: Accurate and efficient classification based on multiple class-association rules. In *Proceedings of the IEEE International Conference on Data Mining, 2001 (ICDM'01)*. IEEE, 369–376.

Bing Liu, Wynne Hsu, and Yiming Ma. 1998. Integrating classification and association rule mining. In *Proceedings of the 4th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 80–86.

Huan Liu and Lei Yu. 2005. Toward integrating feature selection algorithms for classification and clustering. *IEEE Transactions onKnowledge and Data Engineering*, 17, 4, 491–502.

David Lo, Hong Cheng, Jiawei Han, Siau-Cheng Khoo, and Chengnian Sun. 2009. Classification of software behaviors for failure detection: a discriminative pattern mining approach. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 557–566.

Elsa Loekito and James Bailey. 2006. Fast mining of high dimensional expressive contrast patterns using zero-suppressed binary decision diagrams. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 307–316.

Petra Kralj Novak, Nada Lavrač, and Geoffrey I. Webb. 2009. Supervised descriptive rule discovery: A unifying survey of contrast set, emerging pattern and subgroup mining. *The Journal of Machine Learning Research* 10, 377–403.

Jayakrushna Sahoo, Ashok Kumar Das, and A. Goswami. 2014. An effective association rule mining scheme using a new generic basis. *Knowledge and Information Systems* 1–30.

Yuanfeng Song, Wilfred Ng, Kenneth Wai-Ting Leung, and Qiong Fang. 2014. SFP-Rank: Significant frequent pattern analysis for effective ranking. *Knowledge and Information Systems* 1–25.

Peter Spirtes, Clark N. Glymour, and Richard Scheines. 2000. *Causation, Prediction, and Search*. Vol. 81. MIT Press, Cambridge, MA.

Ronan Trépos, Ansaf Salleb-Aouissi, Marie-Odile Cordier, Véronique Masson, and Chantal Gascuel-Odoux. 2013. Building actions from classification rules. *Knowledge and Information Systems* 34, 2, 267–298.

Erik R. Urbach and Tomasz F. Stepinski. 2009. Automatic detection of sub-km craters in high resolution planetary images. *Planetary and Space Science* 57, 7, 880–887.

Jianyong Wang and George Karypis. 2005. HARMONY: Efficiently mining the best rules for classification. In *SDM*, Vol. 5. SIAM, 205–216.

Jialei Wang, Peilin Zhao, S. Hoi, and Rong Jin. 2013b. Online feature selection and its applications. *IEEE Transactions on Knowledge and Data Engineering,* 26, 3, 698–710.

Xiaofeng Wang, Gang Li, Guang Jiang, and Zhongzhi Shi. 2013a. Semantic trajectory-based event detection and event pattern mining. *Knowledge and information systems* 37, 2, 305–329.

Xindong Wu, Kui Yu, Wei Ding, Hao Wang, and Xingquan Zhu. 2013. Online feature selection with streaming features. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 35, 5, 1178–1192.

Xindong Wu, Kui Yu, Hao Wang, and Wei Ding. 2010. Online streaming feature selection. In *Proceedings of the 27th International Conference on Machine Learning (ICML'10)*. 1159–1166.

Xiaoxin Yin and Jiawei Han. 2003. CPAR: Classification based on predictive association rules. In *SDM*, Vol. 3. SIAM, 369–376.

Kui Yu, Wei Ding, Dan A. Simovici, and Xindong Wu. 2012. Mining emerging patterns by streaming feature selection. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 60–68.

Kui Yu, Wei Ding, Hao Wang, and Xindong Wu. 2013. Bridging causal relevance and pattern discriminability: Mining emerging patterns from high-dimensional data. *IEEE Transactions on Knowledge and Data Engineering,* 25, 12, 2721–2739.

Lei Yu and Huan Liu. 2004. Efficient feature selection via analysis of relevance and redundancy. *Journal of Machine Learning Research* 5, 1205–1224.

Xiuzhen Zhang, Guozu Dong, and Ramamohanarao Kotagiri. 2000a. Exploring constraints to efficiently mine emerging patterns from large high-dimensional datasets. In *Proceedings of the 6th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 310–314.

Xiuzhen Zhang, Guozhu Dong, and others. 2000b. Information-based classification by aggregating emerging patterns. In *Intelligent Data Engineering and Automated Learning (IDEAL'00). Data Mining, Financial Engineering, and Intelligent Agents*. Springer, 48–53.

Jing Zhou, Dean P. Foster, Robert A. Stine, and Lyle H. Ungar. 2006. Streamwise feature selection. *Journal of Machine Learning Research* 7, 1861–1885.

Received June 2013; revised December 2014; accepted December 2014