Knowledge actionability: satisfying technical and business interestingness

Longbing Cao*, Dan Luo and Chengqi Zhang

Faculty of Information Technology,
University of Technology, Sydney, Australia
E-mail: lbcao@it.uts.edu.au
E-mail: dluo@it.uts.edu.au E-mail: chengqi@it.uts.edu.au
*Corresponding author

Abstract: Traditionally, knowledge actionability has been investigated mainly by developing and improving technical interestingness. Recently, initial work on technical subjective interestingness and business-oriented profit mining presents general potential, while it is a long-term mission to bridge the gap between technical significance and business expectation. In this paper, we propose a two-way significance framework for measuring knowledge actionability, which highlights both technical interestingness and domain-specific expectations. We further develop a fuzzy interestingness aggregation mechanism to generate a ranked final pattern set balancing technical and business interests. Real-life data mining applications show the proposed knowledge actionability framework can complement technical interestingness while satisfy real user needs.

Keywords: data mining; actionable knowledge; technical interestingness; business decision making; domain driven data mining.

Reference to this paper should be made as follows: Cao, L., Luo, D. and Zhang, C. (2007) 'Knowledge actionability: satisfying technical and business interestingness', *Int. J. Business Intelligence and Data Mining*, Vol. 2, No. 4, pp.496–514.

Biographical notes: Longbing Cao is a Senior Lecturer of the Faculty of Information Technology at the University of Technology, Sydney.

Dan Luo is a PhD candidate of the Faculty of Information Technology at the University of Technology, Sydney.

Chengqi Zhang is a Research Professor of the Faculty of Information Technology at the University of Technology, Sydney.

1 Introduction

Often mined patterns cannot support real user needs for them to take actions in business world (Ali and Wallace, 1997; Cao and Zhang, 2006, 2007; Ghani and Soares, 2006; Silberschatz and Tuzhilin, 1996; Adomavicius, G. and Tuzhilin, A. (1997). This may be due to many reasons. First, it is because business interestingness is rarely considered in existing pattern mining. For instance, in stock data mining, mined trading patterns are normally evaluated in terms of technical interestingness measures such as *correlation coefficient*. While traders, who are recommended to use these discovered patterns, basically check business-oriented metrics like *profit* and *return*. Second, initial work on developing subjective and business related interestingness mainly aims at standard and general measures.

However, due to domain-specific characteristics (Cao and Zhang, 2006, 2007; Ghani and Soares, 2006; Cao, 2007a, 2007b), it is difficult to capture and satisfy particular business expectations in a general manner. As a result, the gap between technical interestingness and business expectations has not been filled or reduced as expected by both researchers and business people. Therefore, it is essential to involve business expectations, and study a two-way significance framework highlighting not only technical but business interestingness (Cao and Zhang, 2006, 2007) in actionable knowledge discovery. The two-way significance framework is an essential component of *domain driven data mining* (Cao and Zhang, 2006, 2007; Cao, 2007a, 2007b), which aims at *the paradigm shift from data-driven hidden pattern mining to domain-driven actionable knowledge discovery*.

To care for business expectations in real-world mining applications, it is necessary to develop business interestingness metrics. For instance, in capital markets, *profit* and *return* are often used to justify whether a trade is acceptable or not from economic performance perspective. Moreover, business interestingness can also be instantiated into *objective* (Freitas, 1998; Hilderman and Hamilton, 2000) and *subjective* (Liu et al., 2000; Silberschatz and Tuzhilin, 1995) measures just like technical metrics. For example, 'beat VWAP' is used by traders to measure whether a price-oriented trading rule is confident enough to 'beat' Value-Weighted Average Price (VWAP) of the market. On the other hand, a stock may be evaluated in terms of certain psychoanalytic factors by a user to determine whether go ahead with it or not. Therefore, besides technical interestingness, actionable knowledge discovery needs to develop *objective* and *subjective business interestingness* as well.

Aiming at involving and satisfying business expectations in actionable knowledge discovery, this paper discusses a two-way significance framework, namely highlighting not only technical but business interestingness in actionable knowledge discovery. The assumption is that an extracted pattern is actionable if only it can satisfy both technical concerns but also business expectations. However, very often there is gap or incompatibility between technical and business interestingness. It is not effective to simply merge two types of metrics. To this end, we propose fuzzy aggregation techniques to combine technical and business interestingness and re-rank mined patterns through balancing two-side concerns. Finally, we demonstrate the application of two-way significance framework to develop actionable trading evidence in stock markets. The examples show that the involvement of business expectations can enhance the actionability of identified patterns when deployed into the real world.

The remainder of this paper is organised as follows. Section 2 introduces related work of knowledge actionability and business interestingness. In Section 3, a knowledge actionability framework is discussed which highlights the straightforward but significant involvement of business expectations. To balance and resolve the incompatibility of technical significance and business expectations, Section 4 proposes a fuzzy interestingness aggregation and ranking approach to create a final ranked actionable pattern list. We demonstrate the application of two-way significance in mining trading patterns in Section 5. Section 6 evaluates the performance. We conclude this paper in Section 7.

2 Related work

The actionable capability of discovered knowledge has attracted more and more attention in the retrospection of traditional trial-and-error data mining process, and the development of next-generation KDD methodology and infrastructure (Aggarwal, 2002; Ankerst, 2001; Cao and Zhang, 2007). In fact, knowledge actionability consists of an essential part of domain driven data mining. The aims and objectives of domain driven data mining research is to push data mining research toward bridging the gap between academia and business, and supporting business decision making based on the identified patterns.

Currently, initial efforts on developing effective knowledge actionability framework can be roughly categorised as follows:

- studying interestingness metrics highlighting significance of *technical subjective* concerns and performance
- developing *new theoretical framework* for valuing the actionability of extracted patterns
- *involving domain and background knowledge* into the search of actionable patterns.

Here we mainly discuss representative work of the first two aspects closely related to this paper.

The concept of *actionability* was initially investigated from the interestingness perspective (Freitas, 1998; Hilderman and Hamilton, 2000; Liu et al., 2000; Padmanabhan and Tuzhilin, 1999; Silberschatz and Tuzhilin, 1995, 1996) to filter out pattern redundancy and 'explicit' (straightforward commonsense to business people) interesting patterns through the mining process or during postprocessing (Yang et al., 2003). A pattern is *actionable* if a user can get benefits (e.g., profit (Wang et al., 2002, 2006; Wang and Su, 2002; Wong et al., 2003; Ling et al., 2006)) from taking actions on it, which may restore the deviation back to its norm. In particular, *subjective* measures such as *unexpectedness* (Silberschatz and Tuzhilin, 1995; Padmanabhan and Tuzhilin, 1998, 1999), *actionability* (Silberschatz and Tuzhilin, 1995, 1996; Yang et al., 2003; Cao and Zhang, 2006, 2007) and *novelty* (Tuzhilin, 2002) were studied to evaluate pattern actionable capability. An *unexpected* pattern contradicts and hence surprises its user's expectations, while it is not necessarily actionable. On the other hand, a *novel* one is new but not necessarily workable to its users.

In a word, the existing actionability-oriented interestingness development mainly emphasises technical and general interestingness, which does not necessarily satisfy business expectations. Even though it is straightforward, only a few researchers work on developing general business interestingness metrics (Cao and Zhang, 2006, 2007), for instance, profit mining (Wang et al., 2002, 2006; Wang and Su, 2002; Wong et al., 2003), to reduce the gap between pattern extraction and decision making. Furthermore, the development of general business interestingness is often limited in measuring colourful domain-specific applications and particular business concerns. Therefore, a more reasonable assumption is that business interestingness should be studied in terms of specific domain problems and involving domain-specific intelligence (Cao and Zhang, 2006, 2007; Cao, 2007a).

Recently, research on theoretical framework of actionable knowledge discovery emerges as a new trend. Kleinberg et al. (1998) proposed a high-level microeconomic framework regarding data mining as an optimisation of decision 'utility'. Wang et al. (2002) built a product recommender which maximises net profit. Action rules (Tzacheva and Ras, 2005) are mined through distinguishing all stable attributes from some flexible ones. Other additional work includes enhancing the actionability of pattern mining in traditional data mining techniques such as association rules (Liu et al., 2000), multi-objective optimisation in data mining (Freitas, 2004), role model-based actionable pattern mining (Wang et al., 2006), cost-sensitive learning (Domingos, 1999) and postprocessing (Yang et al., 2003), etc. All in all, the above research attempted to enhance the existing interestingness system from the technical side.

Different from the above work and traditional data-centered perspective, we propose domain-driven data mining methodology (Cao and Zhang, 2006, 2007; Cao, 2007a, 2007b), which highlights actionable knowledge discovery through involving and synthesising domain intelligence, human intelligence and cooperation, network intelligence and in-depth data intelligence. The objective is to develop theoretical and workable methodologies and techniques to convert academic outcomes directly to business use by reflecting business constraints and expectations. We believe the research on domain-driven actionable knowledge discovery can provide a concrete and practical guideline and hints for corresponding theoretical research in a general manner. In particular, knowledge actionability is one of key components in domain driven data mining.

3 Knowledge actionability framework

This section presents a framework for actionable knowledge discovery. The two-way significance framework measures knowledge actionability from not only technical but also business perspectives. It synthesises business expectation and traditional technical significance in justifying pattern interestingness.

3.1 Technical significance vs. business expectation

The development of actionability is a progressive process in data mining. In the framework of traditional data mining, the so-called actionability *act*() mainly emphasised *technical significance* by developing varying technical interestingness metrics. *Technical interestingness tech_int*() usually measures whether a pattern is of interest or not in terms of specific statistical significance corresponding to a particular data mining method. There are two steps in technical interestingness evolution. The original focus was basically on *technical objective* interestingness *tech_obj*() (Freitas, 1998; Hilderman and Hamilton, 2000), which aims to capture the complexities of pattern structure and statistical significance. Recent work appreciates *technical subjective* measures *tech_sub*() (Liu et al., 2000; Padmanabhan and Tuzhilin, 1999; Silberschatz and Tuzhilin, 1995), which also recognises to what extent a pattern is of interest to a particular user. For example, *probability-based belief* (Silberschatz and

Tuzhilin, 1996) is developed to describe user confidence of unexpected rules (Padmanabhan and Tuzhilin, 1998, 1999). We summarise these two phases as follows.

Let $X = \{x_1, x_2, \quad , x_m\}$ be a set of items, *DB* be a database that consists of a set of transactions, *x* is an itemset in *DB*. Let *e* be interesting evidence discovered in *DB* through a modelling method *M*.

$$\text{Phase 1: } \forall x \in X, \exists e : x.tech\_obj(e) \rightarrow x.act(e) \tag{1}$$

$$\text{Phase 2: } \forall x \in X, \exists e : x.tech\_obj(e) \wedge x.tech\_subj(e) \rightarrow x.act(e). \tag{2}$$

Gradually, data miners realise that the actionability of a discovered pattern must be assessed in terms of and satisfy domain user needs. To this end, we propose the concept of *business interestingness*. *Business interestingness biz_int*() measures to what degree a pattern is of interest to business needs from social, economic, personal and psychoanalytic factors. It further consists of *business objective interestingness* and *business subjective interestingness*. Recently, *business objective interestingness biz_obj*() is recognised by some researchers, say profit mining (Wang et al., 2002). In this stage, we get the Phase 3:

$$\text{Phase 3: } \forall x \in X, \exists e : x.tech\_obj(e) \wedge x.tech\_subj(e) \wedge biz\_obj() \rightarrow x.act(e). \tag{3}$$

For instance, in existing trading pattern mining, even though there are none of business interestingness metrics available, domain knowledge, traders' experience and suggestions, and abovementioned business metrics often used by traders provide the foundation for us to create business interestingness for assessing trading patterns. For example, we use *sharpe ratio SR* as fitness function to evaluate the performance of a mined trading rule in terms of both return and risk. If *SR* is high, the rule likely leads to high return with low risk. Where R_p is expected portfolio return, R_f is risk free rate, σ_p is portfolio standard deviation.

$$SR = (R_p - R_f) / \sigma_p. \tag{4}$$

Moreover, *business subjective* interestingness *biz_sub*() also plays essential roles in assessing *biz_int*(). For instance, empirical measures are usually used by experienced traders to roughly evaluate business performance. As we will discuss in Section 5, the following *index return IR* is used to measure whether *trade return TR* triggered by a mined trading rule can *beat* market *index return IR* or not.

$$TR = \frac{\sum_{i=1}^{u} AskPrice_i \times AskVolume_i - \sum_{j=1}^{v} BidPrice_j \times BidVolume_j}{TotalInvestment} \tag{5}$$

$$IR = \left(\sum_{i=1}^{n} (Index_{i+1} - Index_i) / Index_i \right) \Big/ n. \tag{6}$$

In a specified trading period, there are *n* corresponding index values $Index_i$ ($i = 1, \quad , n$) of a market. $AskPrice_i/AskVolume_i$ and $BidPrice_j/BidVolume_j$ are selling and buying price/volume at trading $time_i/time_j$, which are used when executing a mined trading rule in market data. The number of sells *u* (at ask price and volume) is supposed to be equal to buys *v* (at bid price and volume). *TotalInvestment* is the total investment value. If $IR > TR$, the mined pattern is not actionable enough. These business measures are used in our trading evidence discovery as discussed in Sections 5 and 6.

3.2 *Two-way significance of actionable knowledge*

Based on the above discussion, we believe knowledge actionability should highlight both academic and business concerns (Cao and Zhang, 2007), and satisfy a *two-way significance*. The *two-way significance* indicates that *actionability* recognises, from academic perspective, the technical significance of an extracted pattern through satisfying nominated criteria; while, from business aspect, it should also permit users to specifically react to it to better service their business objectives. In this case, the satisfaction of technical interestingness is the antecedent of checking business expectation. In summary, we view *actionable knowledge* as that satisfies not only technical interestingness *tech_int*() but also user-specified business interestingness *biz_int*(). We have the following definition.

Definition 1 (Knowledge Actionability)**:** Given an mined pattern *e*, its actionable capability *act*(*e*) is described as the extend of its satisfaction with both technical and business interestingness.

$$\forall x \in X, \exists e : x.tech\_int(e) \wedge x.biz\_int(e) \rightarrow x.act(e). \tag{7}$$

Further, it is instantiated in terms of *objective* and *subjective* factors from both *technical* and *business* sides. As a result, we are migrating into the Phase 4 of actionable knowledge discovery, as advocated in domain driven data mining (Cao and Zhang, 2006, 2007).

$$\begin{aligned}&\text{Phase 4: } \forall x \in X,\\ &\exists e : x.tech\_obj(e) \wedge x.tech\_subj(e) \wedge biz\_obj() \wedge biz\_sub() \rightarrow x.act(e).\end{aligned} \tag{8}$$

In this case, there are two sets of interestingness measures needed to be calculated when a pattern is extracted. For instance, we say a mined association trading rule is technically interesting if it satisfies requests on *support* and *confidence*. Moreover, if it can also beat the expectation of user-specified *market index return IR* then it is a generally *actionable* rule.

Unfortunately, it is often not easy to identify patterns satisfying both technical and business interestingness. In the real world, in some cases, the business expectation *biz_int*() of mined patterns may differ from or conflict with their technical significance *tech_int*(). Quite often a pattern with convincing *tech_int*() does not match with *biz_int*() standard. Contrarily, it is not a rare case that a pattern with unsatisfactory *tech_int*() generates interesting *biz_int*(). The relationship between them for a particular pattern *e* may present as one of four scenarios as listed in Table 1.

To illustrate the scenarios outlined in Tables 1 and 2 further presents examples we identified in mining activity patterns in social security data (Cao et al., 2007). For instance, the *support* and the *confidence* of the activity pattern *I*, *J* → \$ (where *I*, *J* and \$ refer to activity codes in social security area) are 0.0003 and 0.0057 respectively, which is very insignificant in statistics view, while its averaged duration amount is very high (46012.43 dollars) with a duration of around nine business days. It shows that the business impact of this pattern is high enough to be worthy of taking actions, even though it might be a very rare business case. This example also shows how business objective interestingness of a pattern can be measured, in which $\overline{d\_amt}()$ and $\overline{d\_dur}()$ are defined in Section 3.3.

Table 1 Relationship between technical significance and business expectation

Scenario	*Relationship type*	*Explanation*
*S*1	$tech\_int() \ll biz\_int()$	A pattern *e* does not satisfy technical significance but satisfies business expectation
*S*2	$tech\_int() \gg biz\_int()$	A pattern *e* does not satisfy business expectation but satisfies technical significance
*S*3	$tech\_int() \cong biz\_int()$	A pattern *e* satisfies business expectation as well as technical significance
*S*4	$tech\_int() \triangleq biz\_int()$	A pattern *e* satisfies neither business expectation nor technical significance

Table 2 Examples showing relation between technical significance and business expectation

		tech_int()		*biz_int*()			
Scenario	*Pattern*	*Support*	*Confidence*	$\overline{d\_amt}()$ (cent)	$\overline{d\_dur}()$ (min)	$risk_{amt}$	$risk_{dur}$
*S*1	I, J → \$	0.0003	0.0057	4601243	12781	0.505	0.203
*S*2	A, R, U → \$	0.386	0.757	2093	8397	0.047	0.03
*S*3	I, F, J → \$	0.257	0.684	639923	9478	0.185	0.094
*S*4	P, S → \$	0.0005	0.0013	1835	6582	0.084	0.028

Clearly, the scenario *S*4 is of no interest to us. Even though, in business world, one only cares about the satisfaction of business expectation, we advocate the goal of actionable knowledge mining is to extract patterns confirming the scenario *S*3: $tech\_int() \cong biz\_int()$, which highlights *two-way significance*. The two-way rather than one-way significance scheme indicates that a pattern has both a solid technical foundation and a robust deployment capability in business world.

Under the scenarios of *S*1 and *S*2, it is a kind of artwork to tune thresholds and their difference between *tech_int*() and *biz_int*(). In real-world data mining, besides developing proper technical and business interestingness measures, there are many other things to do to reach and enhance knowledge actionability. These issues can be studied in the domain-driven data mining methodology (Cao and Zhang, 2006, 2007). Domain-driven data mining highlights the roles and involvement of human intelligence, domain intelligence, data intelligence and network intelligence relevant to a complex problem solving (Cao et al., 2006a). In particular, it is domain users and their knowledge that play essential roles in tuning the thresholds and manage the difference between *tech_int*() and *biz_int*() (Cao and Zhang, 2007).

3.3 Developing business interestingness

In this section, we take the example in Table 2 to illustrate how to build up business interestingness. The following social security activity sequence illustrates a set of activities leading to debt, where letters A–Z represent different activities, \$ indicates the occurrence of a debt. For instance, frequent activity pattern {ACB → \$} can be mined. We here define how to measure the business interestingness of this pattern.

Suppose the total number of itemsets in this data set is $|D|$, where the number of the pattern ACB is $|ACB|$, then we define debt statistics in terms of the following aspects.

< (DABACEKB$), (AFQCPLSWBTC$), (PTSLD$), (QWRTE$), (ARCZBHY) … >.

Definition 2: The total debt amount *d_amt*() is the sum of all individual debt amounts d_amt_i ($I = 1, \ldots, f$) in f itemsets holding the pattern ACB. Then we get *pattern average debt amount* $\overline{d\_amt}()$ for the pattern ACB as:

$$\overline{d\_amt}() = \sum_{1}^{f} d\_amt() \Big/ f. \tag{9}$$

Definition 3: Debt duration *d_dur*() for the pattern ACB is the average duration of all individual debt durations in f itemsets holding the pattern ACB. Debt duration *d_dur*() of an activity is the number of days a debt keeps valid, *d_dur*() = *d.end_date* – *d.start_date* + 1, where *d.end_date* is the day a debt is completed, while *d.start_date* is the day a debt is activated. *Pattern average debt duration* $\overline{d\_dur}()$ is defined as:

$$\overline{d\_dur}() = \sum_{1}^{f} d\_dur \Big/ f. \tag{10}$$

Definition 4: A *pattern's debt amount risk* risk_{amt} is the ratio of the total debt amount of activity itemsets containing ACB to the total debt amount of all itemsets in the data set, donated by $\text{risk}(\text{ACB} \rightarrow \$)_{\text{amt}}$. $\text{risk}(\text{ACB} \rightarrow \$)_{\text{amt}} \in [0, 1]$, the larger it is the higher risk of leading to debt.

$$\text{risk}(ACB \rightarrow \$) = \frac{\sum_{1}^{|ACB|} d\_amt()_i}{\sum_{1}^{|D|} d\_amt()_i}.$$

Definition 5: A *pattern's debt duration risk* risk_{dur} is the ratio of the total debt duration of activity itemsets containing ACB to the total debt duration of all itemsets in the data set, donated by $\text{risk}(\text{ACB} \rightarrow \$)_{\text{dur}}$. Similar to debt amount support, $\text{risk}(\text{ACB} \rightarrow \$)_{\text{amt}} \in [0, 1]$, the larger it is the higher risk of leading to debt.

$$\text{risk}(ACB \rightarrow \$) = \frac{\sum_{1}^{|ACB|} d\_dur()_i}{\sum_{1}^{|D|} d\_dur()_i}. \tag{12}$$

The above metrics can serve on business performance evaluation to measure the impact of a pattern on business outcome. Similar metrics can be developed for specific domains, for instance, in Section 5, we illustrate some business metrics for actionable trading pattern mining in stock market.

4 Aggregating technical and business interestingness

The interestingness gap between academia and business as shown in scenarios *S*1 and *S*2 in Table 1 indicates different objectives of two stakeholders. As also shown in Table 2, there may be inconsistence between multiple metrics belonging to the same categories.

For instance, high $\overline{d\_amt}$ does not mean a definitely corresponding high $risk_{dur}$. This indicates that it is necessary to develop techniques to balance technical and business interestingness, so that a tradeoff and a final ranking can be set up to measure the global performance of pattern set from both sides.

To fill the gap or resolve the conflict, there are some actions to take. First, it is very helpful to involve domain knowledge and user cooperation as much as possible during the whole mining process (Cao and Zhang, 2006). Special assistance is needed in defining and refining interestingness measures and their thresholds, filtering and pruning initially extracted pattern set. Second, developing a hybrid interestingness measure integrating both business and technical interestingness may reduce the burden of requiring domain users to understand those jargons and merge both-side expectations into a one-stop actionability measure. However, the potential incompatibility between technical significance and business expectation makes it difficult to aggregate the two sides of metrics. A simple weight-based integration does not work due to internal inconsistence among measures. Therefore, rather than taking conventional weighted-formula approach (Freitas, 2004), we develop fuzzy interestingness aggregation method combining *tech_int*() and *biz_int*() to re-rank the mined pattern set.

The idea of fuzzy aggregation of technical and business interestingness is as follows. Even though simple fuzzy aggregation of interestingness measures can be viewed as a fuzzily weighted approach, we deal with this from pattern rather than measure perspective. Through defining fuzzy sets supervised by business users, we first fuzzify the extracted patterns into two sets of fuzzily ranked pattern sets in terms of fitness functions *tech_int*() and *biz_int*(), respectively. We then aggregate these two fuzzy pattern sets to generate a final fuzzy ranking. This final fuzzily ranked pattern set is recommended to users for their consideration. Although this strategy is a little bit fuzzy, it combines two-side interestingness while balances individual contributions and diversities.

In fuzzifying pattern set in terms of specific interestingness measures, the universe of discourse of a fuzzified measure must be in [0, 1]. In the following, we explain the fuzzy aggregation and ranking principle through a simple example. As illustrated in Figures 1 and 2, suppose we use two sets of five ascending linguistic values to fuzzify and segment the whole pattern set, we then get a fuzzily ranked technical pattern class $T_s = \{a, b, c, d, e\}$ and a fuzzy business pattern class $B_s = \{A, B, C, D, E\}$ in terms of the fuzzification of *tech_int*() and *biz_int*(), respectively.

Figure 1 Fuzzily ranked technical pattern class

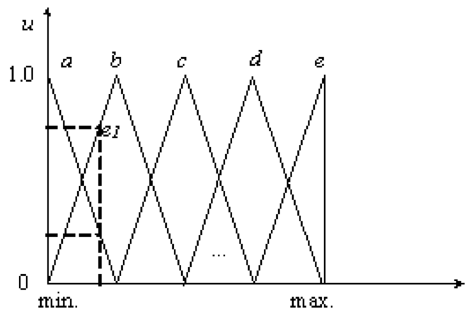

Even though class T_s and class B_s may present the same number of linguistic term, their semantics, say *b* and *B*, likely are varying. This means that we cannot simply aggregate corresponding items from two classes into one integrative output through weighted-formula approach. Instead, we develop fuzzy aggregation rules to aggregate these two fuzzy groups to generate a final recommendation list.

Figure 2 Fuzzily ranked business pattern class

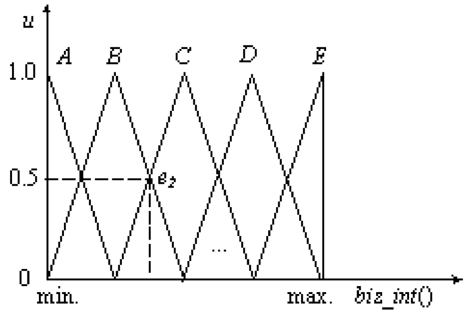

Definition 6 (Fuzzily aggregated pattern ranking)**:** Given a pattern ranked as *m*-th in pattern class T_s, while it is ranked as *n*-th in class B_s, its aggregated ranking is $(m + n - 1)$-th.

For example, if a pattern *e* technically is ranked as *c* (3rd), while as *D* (4th) from business perspective, then its final fuzzy ranking is 6th $(3 + 4 - 1)$ in the final aggregated pattern set. Further, the following definition specifies the length of the final aggregated pattern set.

Definition 7 (Length of final aggregated pattern set)**:** If mined target patterns are ranked in terms of t_s levels based on technical interestingness, while they are ranked into b_s levels from business perspective, then after aggregating these two ranking classes, the length of the final aggregated ranking set is $(t_s + b_s - 1)$.

The above aggregation and ranking is based on the fuzzification of fitness and membership functions. As a result, pattern ranking presents uncertainty. For example, as shown in Figures 1 and 2, the trading pattern e_1 can be ranked into group *b* with membership grade $= 0.75$ or *a* with grade $= 0.25$. Similar thing happens to business class, the pattern e_2 can be segmented into *B* or *C* with same grade $= 0.5$. In this example, the outcomes of the fuzzily aggregated ranking present three options, namely 2nd, 3rd or 4th in a total of nine candidates. This will bring users big inconvenience.

To manage the above uncertainty emerging in fuzzy aggregation and ranking, a ranking coefficient ρ based on moment defuzzification is introduced to defuzzify a fuzzy set to convert into a floating point that represents the final position.

$$\rho = \frac{\sum_{l=1}^{m} \eta_l \; {}^{T}_{l} \; {}^{B}_{l}}{\sum_{l=1}^{m} {}^{T}_{l} \; {}^{B}_{l}} \tag{13}$$

where, m refers to the number of triggered linguistic values, $l = 1, 2, \ldots, m$ corresponds to each triggered linguistic value. ${}^{T}_{l}$ is the membership grade of No. l linguistic term relevant to the technical fitness of a pattern. ${}^{B}_{l}$ is the membership grade of No. l linguistic term corresponding to the business interestingness of a pattern. η_l is the centroid of the No. l triggered linguistic value, it is calculated in terms of the moment and the area of each subdivision.

A real number can be calculated to measure a fuzzily aggregated pattern ranking in a relatively crisp manner. For instance, we can calculate and get $\rho = 0.125$ in the above example, which clearly indicates that the pattern is ranked as 3*rd of nine* in the finally aggregated ranking set since its membership grade is 0.75 much larger than grade 0.25 as ranked 4th.

5 A case study

In this section, we introduce a real-world data mining application that utilises the two-way significance framework in actionable knowledge discovery. This is to mine actionable trading patterns in stock markets. Here 'actionable' indicates that the discovered trading patterns should not only satisfy certain level of technical significance based on used data mining methods, but also have potential in making good money when they are deployed into the market.

In stock markets, there exist many trading strategies designed either in house or by financial researchers. A trading strategy actually indicates certain trading pattern existing in a market, which can be used to guide the trading to a trader's advantage, for instance automatic trading. However, trading patterns available from financial literature are too general to be used for specific market dynamics. Ordinary traders are normally not confident in deploying such trading rules to achieve their expectations. On the other hand, data miners are short of market microstructure knowledge and generally not sensitive to dynamic market factors. As a result, it is very difficult to discover trading patterns purely based on machine learning models, or even though patterns may be discovered, but they are unnecessarily workable in real market.

Our experience of stock data mining (Cao et al., 2006b; Lin and Cao, 2006), collaborated with Capital Markets Cooperative Research Center[1] and Securities Industry Research Centre of Asia-Pacific,[2] has disclosed a more practical way for discovering actionable trading patterns using data mining methods. This is to mine in-depth trading patterns based on general pattern set generated by financial producers, rather than discovering patterns purely from the scratch. In the following, we illustrate this approach to mine trading rule-stock pairs.

The brief process of mining trading rule-stock pairs is as follows.

> First, we collect general trading strategies from our financial partners and resources. For instance, Sullivan et al. (1999) presented some classic trading rules. In general, due to constraints of multiple organisational factors in real market (O'Hara, 1998), these trading rules cannot generate good performance if they are used to directly trade stocks.

Second, we extend a basic trading rule model, say Double Moving Average Trading Strategy, designed by financial experts, through modelling market constraints. The extension adds more parameters or organisational factors into the basic model based on historic market trend analysis and profit/cost analysis. For instance, to remove the impact of noise in market movement on trading rule modelling, we add features such as 'volatility' and 'holding days' into the model. The outcome is an enhanced trading rule model which reflects more practical situations in the market.

Third, in mining trading rule-stock associations, we measure the correlation between each instantiated trading rule of the enhanced model and tradable stocks in a market by calculating coefficients from both technical and business perspectives. In general, an enhanced trading rule has multiple numeric attributes. The value ranges of these attributes are set under the supervision of domain knowledge. Even in a limited value range, the search space of all combinations of a trading rule-stock pair is huge in a market. In such a searchable space, some rule-stock pairs are with better performance than others.

Fourth, genetic algorithm and the above fuzzy aggregation technique (Lin and Cao, 2006; Cao et al., 2006b) are used to discover trading rule-stock pairs that have potential in achieving both sides of significance. Those rule-stock pairs with high fuzzy ranking are likely with decent technical and business significance. They are searched in the combinational space and are recommended to traders.

To support the above work, we develop an intelligent agent and service based stock data mining system: F-Trade (Cao et al., 2004).[3] The *F*-Trade is connected to data warehouse in AC3, where global market stock data is accessible for our research under the collaborative agreement with CMCRC and SIRCA. With initial setting by domain users, for instance, the value range of a parameter, the *F*-Trade automatically searches and generates trading rule-stock pairs with decent technical and business significance. It also pops out outputs in both summarised and detailed manners for users to scrutinise. Further, users can tune the parameters through the pop-out parameter window at run time. The *F*-Trade then re-calculates and generates the output for the new settings. Figures 3 and 4 illustrate some screenshots for the parameter tuning and output windows.

Figure 3 Tuning parameters at run time in *F*-trade

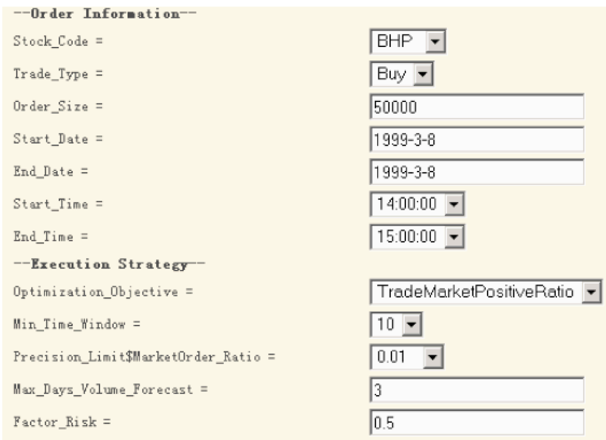

Figure 4 Merged multi-level outputs of actionable trading pattern in *F*-trade

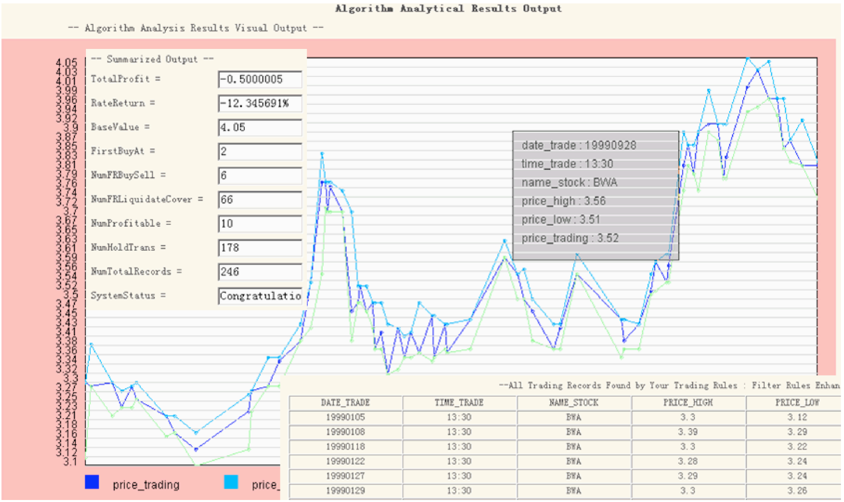

To discover trading rule-stock pairs (Cao et al., 2006b), we test the proposed approaches in Australian Stock eXchange (ASX) intraday orderbook data in 2001. Sliding window strategy was used for constructing training and testing sets, namely using data in month a as in-sample, the data of the next month $a + 1$ as out-of-sample. After one round of training-testing, month $a + 1$ becomes the in-sample, and the out-of-sample is month $a + 2$. This process is repeated till the ending month of the data set.

Figure 5 illustrates an excerpt of fuzzily ranked rule-stock pairs discovered in ASX intraday orderbook data in June 2001. Three classes of rules (coded as 1–3) and 27 ASX stocks (coded from 1 to 27) are selected for this experiment. Total 81 rule-stock pairs are classified into five fuzzy groups in terms of *sharpe ratio SR*: VP-very positive, P-positive, W-watchable, N-negative and VN-very negative. We can see that Rule 2-Stock 24, Rule 2-Stock 26, and Rule 1-Stock 24 are three *very positive* pairs in June 2001. These high-ranking pairs may be useful to support trading.

Figure 5 Fuzzily ranked trading rule-stock pairs in ASX market

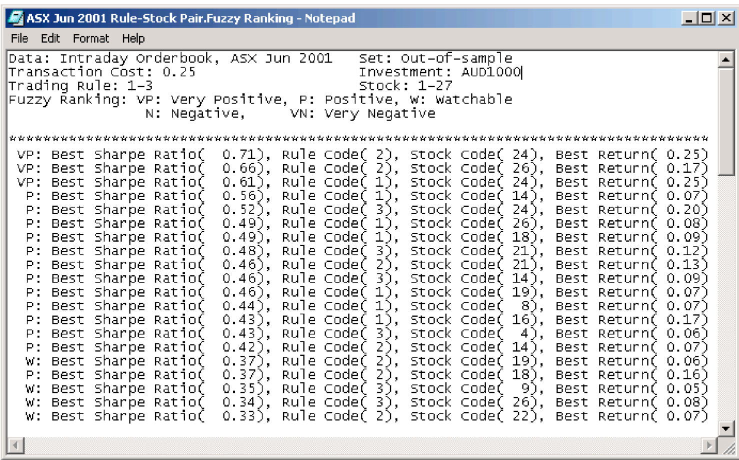

```
ASX Jun 2001 Rule-Stock Pair.Fuzzy Ranking - Notepad
File  Edit  Format  Help
Data: Intraday Orderbook, ASX Jun 2001     Set: Out-of-Sample
Transaction Cost: 0.25                     Investment: AUD1000
Trading Rule: 1-3                          Stock: 1-27
Fuzzy Ranking: VP: Very Positive, P: Positive, W: Watchable
                N: Negative,      VN: Very Negative

**********************************************************************************
 VP: Best Sharpe Ratio(  0.71), Rule Code( 2), Stock Code( 24), Best Return( 0.25)
 VP: Best Sharpe Ratio(  0.66), Rule Code( 2), Stock Code( 26), Best Return( 0.17)
 VP: Best Sharpe Ratio(  0.61), Rule Code( 1), Stock Code( 24), Best Return( 0.25)
  P: Best Sharpe Ratio(  0.56), Rule Code( 1), Stock Code( 14), Best Return( 0.07)
  P: Best Sharpe Ratio(  0.52), Rule Code( 3), Stock Code( 24), Best Return( 0.20)
  P: Best Sharpe Ratio(  0.49), Rule Code( 1), Stock Code( 26), Best Return( 0.08)
  P: Best Sharpe Ratio(  0.49), Rule Code( 1), Stock Code( 18), Best Return( 0.09)
  P: Best Sharpe Ratio(  0.48), Rule Code( 3), Stock Code( 21), Best Return( 0.12)
  P: Best Sharpe Ratio(  0.46), Rule Code( 2), Stock Code( 21), Best Return( 0.13)
  P: Best Sharpe Ratio(  0.46), Rule Code( 3), Stock Code( 14), Best Return( 0.09)
  P: Best Sharpe Ratio(  0.46), Rule Code( 1), Stock Code( 19), Best Return( 0.07)
  P: Best Sharpe Ratio(  0.44), Rule Code( 1), Stock Code(  8), Best Return( 0.07)
  P: Best Sharpe Ratio(  0.43), Rule Code( 1), Stock Code( 16), Best Return( 0.17)
  P: Best Sharpe Ratio(  0.43), Rule Code( 3), Stock Code(  4), Best Return( 0.06)
  P: Best Sharpe Ratio(  0.42), Rule Code( 2), Stock Code( 14), Best Return( 0.07)
  W: Best Sharpe Ratio(  0.37), Rule Code( 2), Stock Code( 19), Best Return( 0.06)
  P: Best Sharpe Ratio(  0.37), Rule Code( 2), Stock Code( 18), Best Return( 0.16)
  W: Best Sharpe Ratio(  0.35), Rule Code( 3), Stock Code(  9), Best Return( 0.05)
  W: Best Sharpe Ratio(  0.34), Rule Code( 3), Stock Code( 26), Best Return( 0.08)
  W: Best Sharpe Ratio(  0.33), Rule Code( 2), Stock Code( 22), Best Return( 0.07)
```

6 Evaluation

With regard to the above trading pattern mining following two-way significance framework, the performance of identified trading patterns is evaluated in terms of the following strategies:

- the comparison of predictive performance of those patterns discovered based on fuzzily aggregated two-way significance framework vs. those identified simply considering business performance
- the business performance comparison of identified patterns based on fuzzily aggregated two-way significance, equally weighted two-way significance vs. those discovered purely based on technical criteria
- the business performance of patterns discovered based on fuzzily aggregated two-way significance framework by checking traders' justification criteria.

6.1 Predictive performance comparison

The first experiment is to compare predictive performance in terms of trading rule-stock pairs identified using fuzzily aggregated two-way significance framework vs. those identified simply considering good *market return*.

The predictability is calculated as follows. Let D be the number of total rule-stock pairs found in in-sample and out-of-sample sets satisfying certain interestingness criteria. $A(B)$ be the number of total pairs in in-sample (out-of-sample) data set. AB be the number of pairs existing in both in-sample and out-of-sample sets. Then, the following statistics measures the predictability of identified trading rule-stock pairs in terms of in-sample and out-of-sample comparison.

Definition 8 (Probability of rule-stock pairs)**:** The following probability functions are defined for rule-stock pairs in or across in-sample and out-of-sample data sets: $P(A) = A/D$, $P(B) = B/D$, $P(AB) = AB/D$, $P(A|B) = P(AB)/P(B)$, $P(B|A) = P(AB)/P(A)$.

Definition 9 (Predictability metrics of rule-stock pairs)**:** The following metrics: Confidence, All_Confidence, Cosine and Coherence are defined for measuring the predictability of the patterns identified in in-sample set when they are tested in out-of-sample data.

$$\text{Confidence} = \max(P(A/B), P(B/A)) \quad (14)$$

$$\text{Coherence} = P(AB)/(P(A) + P(B) - P(AB)). \quad (15)$$

In general, the value range of the above metrics is in [0, 1]. Larger value denotes bigger actionability of pairs when they are deployed into the real trading.

We calculate the predictability of mining trading rule-stock pairs in 2001 ASX intraday orderbook data using sliding window strategy for training and testing. Table 3 lists the coherence and confidence of top 10% pairs extracted by considering high monthly return only. Table 4 provides the summarised statistics for those patterns identified based on fuzzily aggregated two-way significance framework.

Table 3 Predictability of top 10% rule-stock pairs maximising business expectation

	April (%)	*May* (%)	*June* (%)	*July* (%)	*August* (%)	*September* (%)	*October* (%)
Coherence	14.3	23	14.3	6.7	23	6.7	14.3
Confidence	25	37.5	25	12.5	37.5	12.5	25

Table 4 Predictability of top 10% rule-stock pairs based on fuzzy two-way significance

	April (%)	*May* (%)	*June* (%)	*July* (%)	*August* (%)	*September* (%)	*October* (%)
Coherence	21.3	37.5	12.8	9.4	18.6	10.5	23.2
Confidence	50	75	50	25	50	25	50

The above performance analysis shows that the confidence of rule-stock pairs based on fuzzily aggregated two-way significance seems better than that of pairs identified satisfying business performance only. It shows that the two-way significance can present patterns with more solid statistical significance and higher predicative capability.

6.2 Business performance comparison: three pattern interestingness frameworks

The second experiment is as follows. Three groups of trading rule-stock pairs are mined. The first group of pairs is identified only based on technical interestingness measures. The second group is identified based on the fuzzy two-way significance framework as discussed in Section 4. The third group is extracted using equally weighted sum of technical and business interestingness. Top 10% pairs are selected from each of the above groups. It is worthy of noting that the comparison of each particular pair's performance in the top 10% pair groups may not be very reasonable, because there may only be a small proportion of, if not none, pair overlapping in three groups. To eliminate the error triggered by pair difference, we trade the top 10% pairs from each of groups as a whole based on same investment in the same market. We then calculate the monthly return of the whole bundle of top 10% pairs in each group. Figure 6 shows the monthly return of three groups from March to October in ASX 2001 orderbook data (total investment = AUD100k).

Figure 6 Monthly return of top 10% rule-stock pairs (technical interest vs. fuzzy two-way interest vs. equally weighted two-way significance)

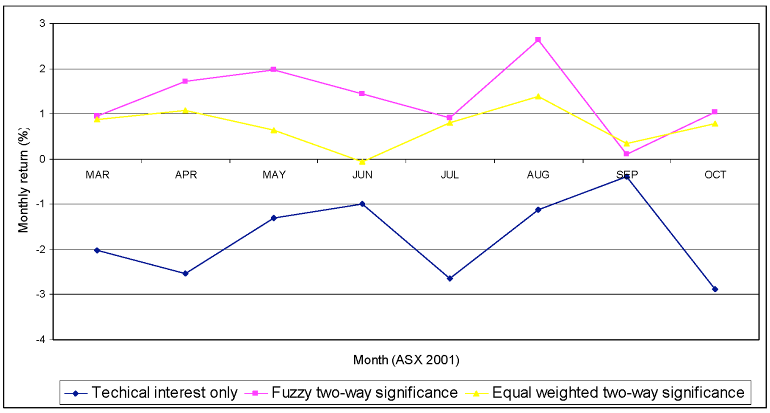

As shown in Figure 6, the business performance of those pairs only based on technical interestingness is much worse than that of pairs extracted by following two-way significance framework. The averaged monthly return of technical interest only is –1.737%, which is much lower than that of fuzzy two-way (1.343%) and equally weighted two-way (0.731%) frameworks. On the other hand, pairs based on fuzzy two-way significance framework get slightly better business performance than equally weighted two-way strategy.

6.3 Business performance comparison: beat market

The third experiment is to measure the business performance of identified trading rule-stock pairs in terms of traders' expectations and market dynamics. From objective perspective, we calculate and evaluate economic performance such as *trade return TR*, *index return IR*, and *sharpe ratio SR* as indicated in Formulas (5), (6) and (4). On the subjective side, empirical and psychoanalytic indicators are studied. For instance, 'beat market', such as beating transaction costs, beating market return, is nominated as a real-world benchmark of judging whether traders can take actions on mined trading patterns or not. *Beat transaction costs* means that specific economic measures, say *TR*, must positively surpass user-specified thresholds after deducting the impact of transaction costs on each trading transaction. *Beat market return* means that the actual monthly mean trade return *TR* generated by a trading rule must be better than the market index return *IR*.

Based on fuzzy interestingness aggregation and ranking, Figure 7 shows the monthly *TR* gained by trading different levels of trading rule-stock pairs after deducting 0.25% transaction cost in February 2001. The monthly return is basically positive, which demonstrates that trading on these mined trading rule-stock pairs can beat transaction costs. Traders welcome this encouraging result because their experience tells them "transaction costs can 'kill' most of trading rules" in real trading.

Figure 7 The monthly return *TR* to top x% rule-stock pairs

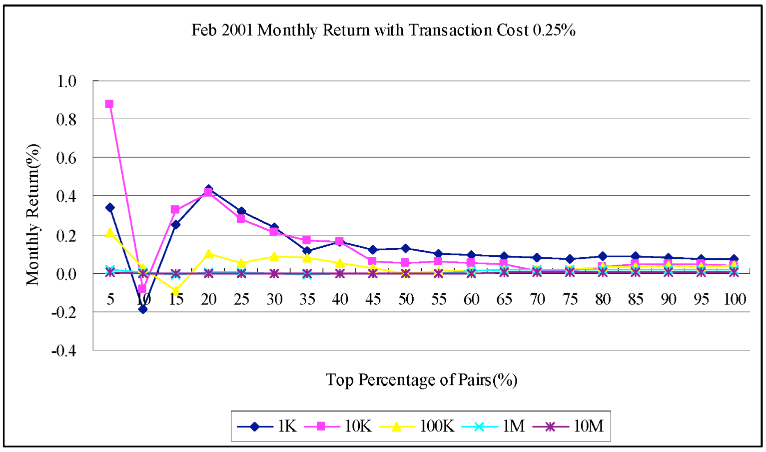

Figure 8 further compares the market index return *IR* with the monthly trade return *TR* of top 5% pairs in ASX 2001 orderbook data after deducting 0.25% transaction cost. This result shows that under the bear market situation in 2001 ASX market, trading these trading rule-stock pairs still could beat market index return. From another perspective, it demonstrates that the discovered rule-stock pairs are very promising for supporting traders' decision making.

Figure 8 The monthly trade return *TR* vs. monthly market index return *IR*

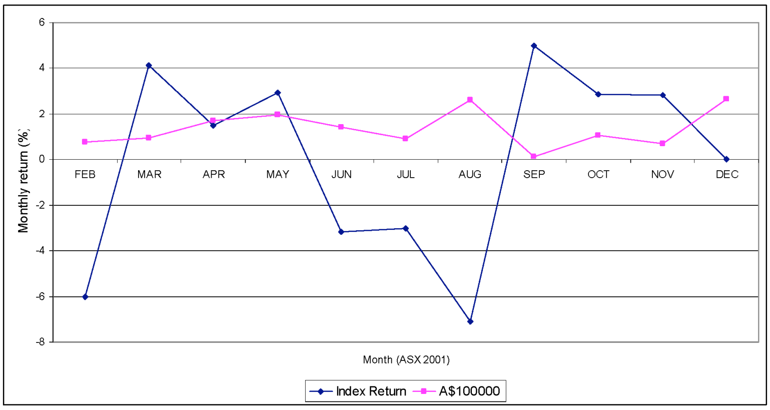

7 Conclusions

Current data mining interestingness framework mainly relies on technical aspects. Both subjective and objective interestingness have been studied. However, the evaluation of knowledge discovered whether it can be used to take actions in business world or not cannot just depend on technical measures. In practice, there often emerges a gap between technical significance and business expectations. To bridge the gap, both technical and business interestingness needs to be considered in measuring knowledge actionability. In fact, enhancing knowledge actionability has become one of most significant and challenging work in KDD. It is nominated as one of grand challenges of KDD in the future. The research on this issue may change the existing situation where a great number of rules are mined while few of them are interesting to business, and promote the wide deployment of data mining into business.

In this paper, we highlight the significance of explicitly involving business expectations and effectively bridge the gap and incompatibility between academic and business in mining actionable knowledge. We have proposed a two-way significance framework, which highlights the significant role of business expectations in evaluating knowledge actionability. Based on real-life lessons, we have analysed possible scenarios when both technical and business interestingness are concerned in data mining. It has been disclosed that very often there emerges incompatibility and uncertainty issues in balancing business and academic interestingness. To this end, and to bridge the gap, we have developed fuzzy aggregation mechanism to generate a final ranking pattern list satisfying two-side concerns.

We have also demonstrated how to design business interestingness in social security mining. In addition, experiments of mining actionable trading patterns in stock markets have been conducted, in which we have involved both business objective measures such as *sharpe ratio* and *trade return* and business subjective such as 'beat market'.

These examples show that the involvement of business interestingness can greatly enhance knowledge actionability.

Acknowledgements

This work is sponsored by Australian Research Council Discovery Grant (DP0773412, LP0775041, DP0667060), and UTS internal grants.

References

Adomavicius, G. and Tuzhilin, A. (1997) 'Discovery of actionable patterns in databases: the action hierarchy approach', *SIGKDD*, pp.111–114.

Aggarwal, C. (2002) 'Towards effective and interpretable data mining by visual interaction', *ACM SIGKDD Explorations Newsletter*, Vol. 3, No. 2, pp.11–22.

Ali, O.F.G. and Wallace, W.A. (1997) 'Bridging the gap between business objectives and parameters of data mining algorithms', *Decision Support Systems*, Vol. 21, pp.3–15.

Ankerst, M. (2001) 'Human involvement and interactivity of the next generation's data mining tools', *Workshop on Research Issues in Data Mining and Knowledge Discovery joint with DMKD2001*, Santa Barbara, CA, from http://www.cs.cornell.edu/johannes/papers/dmkd 2001-papers/p12_ankerst.pdf.

Cao, L. (2007a) 'Actionable knowledge discovery, Mehdi Khosrow-Pour', *D.B.A., Encyclopedia of Information Science and Technology*, 2nd ed., IGI Global Publications, Hershey, USA.

Cao, L. (2007b) 'Domain-driven, actionable knowledge discovery', *IEEE Intelligent Systems*, Vol. 22, No. 4, pp.78–89.

Cao, L. and Zhang, C. (2006) 'Domain-driven data mining – a practical methodology', *Int. J. Data Warehousing and Mining*, Vol. 2, No. 4, pp.49–65.

Cao, L.B. and Zhang, C.Q. (2007) 'The evolution of KDD: towards domain-driven data mining', *Int. J. Pattern Recognition and Artificial Intelligence*, Vol. 21, No. 4, pp.677–692.

Cao, L.B., Wang, J., Lin, L. and Zhang, C. (2004) 'Agent services-based infrastructure for online assessment of trading strategies', *Proceedings of the IAT'04*, IEEE Computer Society Press, Beijing, China, pp.345–349.

Cao, L.B., Zhang, C., Luo, D., Dai, R. and Chew, E. (2006a) 'Intelligence metasynthesis in building business intelligence systems', *Int. Workshop on Web Intelligence meets Brain Informatics – WImBI'06*, to appear in Lecture Notes in Computer Sciences, Springer, Berlin.

Cao, L., Luo, D. and Zhang, C. (2006b) 'Fuzzy genetic algorithms for pairs mining', *Proceedings of PRICAI2006*, LNAI 4099, Guilin, China, pp.711–720.

Cao, L., Zhao, Y. and Zhang, C. (2007) 'Mining impact-targeted activity patterns in unbalanced data', to appear in *IEEE Transactions on Knowledge and Data Engineering*.

Domingos, P. (1999) 'MetaCost: a general method for making classifiers cost-sensitive', *KDD'99*, ACM Press, pp.155–164.

Freitas, A.A. (1998) 'On objective measures of rule surprisingness', in Zytkow, J. and Quafafou, M. (Eds.): *PKDD'98*, Berlin, Springer, pp.1–9.

Freitas, A.A. (2004) 'Critical review of multi-objective optimization in data mining – a position paper', *SIGKDD Explorations*, Vol. 6, No. 2, pp.77–86.

Ghani, R. and Soares, C. (2006) *Proc. Workshop on Data Mining for Business Applications, Joint in KDD2006*, Philadephia, USA.

Hilderman, R.J. and Hamilton, H.J. (2000) 'Applying objective interestingness measures in data mining systems', *PKDD'00*, pp.432–439.

Kleinberg, J. Papadimitriou, C. and Raghavan, P. (1998) 'A microeconomic view of data mining', *Journal of Data Mining and Knowledge Discovery*, Vol. 2, No. 4, pp.311–324.

Lin, L. and Cao, L. (2007) 'Mining in-depth patterns in stock market', to appear in *Int. J. Intelligent System Technologies and Applications*.

Ling, C., Sheng, W., Bruckhaus, T. and Madhavji, N. (2006) 'Maximum profit mining and its application in software development', *Proc. SIGKDD 2006*, ACM Press, pp.929–934.

Liu, B., Hsu, W., Chen, S. and Ma, Y. (2000) 'Analyzing subjective interestingness of association rules', *IEEE Intelligent Systems*, Vol. 15, No. 5, pp.47–55.

O'Hara, M. (1998) *Market Microstructure Theory*, Blackwell Publishing Professional.

Padmanabhan, B. and Tuzhilin, A. (1998) 'A belief-driven method for discovering unexpected patterns', *KDD-98*, pp.94–100.

Padmanabhan, B. and Tuzhilin, A. (1999) 'Unexpectedness as a measure of interestingness in knowledge discovery', *Decision and Support Systems*, Vol. 27, pp.303–318.

Silberschatz, A. and Tuzhilin, A. (1995) 'On Subjective Measures of Interestingness in Knowledge discovery', *Knowledge Discovery and Data Mining*, pp.275–281.

Silberschatz, A. and Tuzhilin, A. (1996) 'What makes patterns interesting in knowledge discovery systems', *IEEE Transactions on Knowledge and Data Engineering*, Vol. 8, No. 6, pp.970–974.

Sullivan, R., Timmermann, A. and White, H. (1999) 'Data-snooping, technical trading rule performance, and the bootstrap', *The Journal of Financial*, Vol. 54, No. 5, pp.1647–1692.

Tuzhilin, A. (2002) 'Knowledge evaluation: other evaluations: usefulness, novelty, and integration of interesting news measures', *Handbook of Data Mining and Knowledge Discovery*, pp.496–508.

Tzacheva, A. and Ras, W. (2005) 'Action rules mining', *Int. J. Intelligent Systems*, Vol. 20, pp.719–736.

Wang, K. and Su, M. (2002) 'Item selection by 'hub-authority' profit ranking', *SIGKDD 2002*, ACM Press, pp.652–657.

Wang, K., Zhou, S. and Han, J. (2002) 'Profit mining: from patterns to actions', *EBDT 2002*, Prague, Czech, pp.70–87.

Wang, K., Jiang, Y. and Tuzhilin, A. (2006) 'mining actionable patterns by role models', *ICDE 2006*, IEEE Computer Science Press, p.16.

Wong, R., Fu, A. and Wang, K. (2003) 'MPIS: maximal-profit item selection with cross-selling considerations', *Proc. ICDM 2003*, IEEE Computer Science Press, pp.371–378.

Yang, Q., Yin, J., Lin, C. and Chen, T. (2003) 'Postprocessing decision trees to extract actionable knowledge', *Proc. ICDM2003*, IEEE Computer Science Press, pp.685–688.

Notes

[1]www.cmcrc.com.au.

[2]http://www.sirca.org.au/.

[3]www.f-trade.info.