On Subjective Measures of Interestingness in Knowledge Discovery

Avi Silberschatz
AT&T Bell Laboratories
600 Mountain Avenue
Murray Hill, NJ

Alexander Tuzhilin*
Information Systems Department
Stern School of Business
New York University

Abstract

One of the central problems in the field of knowledge discovery is the development of good measures of interestingness of discovered patterns. Such measures of interestingness are divided into *objective* measures – those that depend only on the structure of a pattern and the underlying data used in the discovery process, and the *subjective* measures – those that also depend on the class of users who examine the pattern. The purpose of this paper is to lay the groundwork for a comprehensive study of subjective measures of interestingness. In the paper, we classify these measures into *actionable* and *unexpected*, and examine the relationship between them. The unexpected measure of interestingness is defined in terms of the *belief* system that the user has. Interestingness of a pattern is expressed in terms of how it affects the belief system.

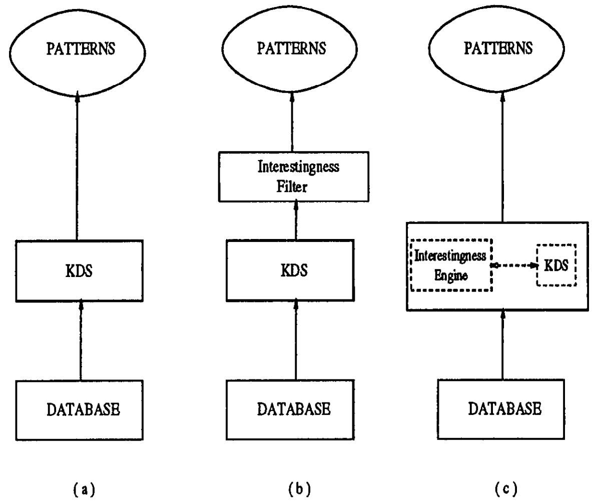

Figure 1: A Framework for Knowledge Discovery.

1 Introduction

Over the past several years, the knowledge discovery community developed a "discovery" framework that is schematically represented in Figure 1-a[1]. As Figure 1-a shows, the data is stored in some database, which can be based on any data model (e.g. relational, network, and object-based). On "top" of this database, there is a knowledge discovery system (KDS) that generates some patterns which are presented to the user. For the purpose of this paper, it does not matter how the KDS works and what the structure of the discovered patterns is. Therefore, we will treat the KDS as a "black box" and assume that a pattern is an arbitrary first-order sentence expressed in terms of the schema of the database or the user-defined vocabulary [2].

It has been recognized in the knowledge discovery literature that a discovery system can generate a glut of patterns, most of which are of no interest to the user [3]. Thus, it is of vital importance to define good measures of interestingness that would allow the system to discover only the useful patterns.

One approach to defining interestingness of a pattern is to define it in *objective* terms, where interestingness of a pattern is measured in terms of its structure and the underlying data used in the discovery process. For example, as Piatetsky-Shapiro points out in [7], the interestingness of a rule $A \rightarrow B$ is usually defined as a function of $p(A)$, $p(B)$ and $p(A \wedge B)$, where $p(\alpha)$ is the probability that condition α is true. Typical examples of such objective measure of interestingness of a rule are its "information content" based on the J-measure [10], a certainty factor [4], and a strength [2].

It has been noted in [8] that objective measures of interestingness, although useful in many respects, usually do not capture all the complexities of the pattern discovery process, and that *subjective* measures of interestingness are needed to define interestingness of a pattern. These subjective measures do not depend only on the structure of a rule and on the data used in the discovery process, but also on the *user*

*This work was supported in part by the NSF under Grant IRI-93-18773.

[1]This is a simplified version of Figure 1-1 from [3]. To avoid unnecessary complications, we present only the parts of Figure 1-1 that are relevant to our paper.

who examines the pattern. These measures recognize that a pattern that is of interest to one user, may be of no interest to another user. For example, a pattern discovering some security trading irregularities, such as insider trading, may be of great interest to the officials from the Securities and Exchange Commission (SEC). However, it is of very little use to a homeless person living in New York City.

In [8], subjective measures of interestingness were studied within the context of the discovery system KEFIR [6] that analyzes healthcare insurance claims for uncovering "key findings." The *key findings* in KEFIR are the most important changes from the norms for various *indicators* assessing different characteristics of provision of healthcare, such as cost, usage, and quality. The authors in [8] argue that "a good measure of the interestingness of a finding is the estimated benefit that could be realized by taking a specific action in response." Since the KEFIR system deals with the financial data pertaining to the insurance claims, it measures benefits in financial terms. KEFIR classifies all possible findings into a predefined set of types, each type defined in terms of some logical condition imposed on one or several indicators. KEFIR then defines a production rule for each type of finding that specifies the actions to be taken in response to the findings that typically indicate how to bring "abnormal" indicators back to their norms. Moreover, the domain expert needs to assign a probability of success to the actions in the rule. Once a new finding is discovered, the system determines all the production rules matching this finding and selects the rule with the highest probability of success. It then computes the estimated benefit of taking the action for the selected rule as potential savings realized from the action restoring the deviation back to its norm. This estimated benefit serves as a measure of interestingness in KEFIR.

The method used in KEFIR provides a good approach for defining a subjective measure of interestingness in terms of the benefits accrued from the corrective actions bringing deviated measures back to their norms. This approach, however, is very domain specific for the following reasons. First, it deals only with patterns expressed as deviations (changes of an indicator from its norm). Second, it pre-classifies all the patterns that can be discovered into a finite (hopefully small) set of classes in order to assign a corrective action to each class. This is possible in case of KEFIR because it deals with a very domain-specific problem related to healthcare. Thirdly, it makes several domain-specific assumptions about the way estimated benefits are computed.

In this paper, we also study subjective measures of interestingness. However, unlike [8], we study them in a domain-independent context. In particular, we propose a classification of measures of interestingness and identify two major reasons why a pattern is interesting from the subjective (user-oriented) point of view:

- *Unexpectedness* – a pattern is interesting if it is "surprising" to the user.
- *Actionability* – a pattern is interesting if the user can do something with it to his or her advantage. This is, essentially, the subjective measure of interestingness studied in [8].

We also examine the relationship between these two measures of interestingness with the emphasis on the first measure.

Once the concept of interestingness is defined, it can be incorporated into the KBS in one of the following two ways. First of all, an *interestingness filter* can be developed and placed at the *back-end* of the KBS module (see Figure 1-b). In this case, the KDS generates many patterns that are put through the filter that selects only few "interesting" patterns. Alternatively, an *interestingness engine* can be placed *inside* the KBS module so that it can *focus* the search only on the interesting patterns (see Figure 1-c). Clearly, the second approach is preferred since it avoids generating many non-interesting patterns, thus, speeding the discovery process.

The purpose of this paper is to lay the groundwork for a comprehensive study of domain independent subjective measures of interestingness. The main emphasis of the paper is on the attempt to understand, at the *intuitive level*, what these measures are and how they are related to each other. Having this goal in mind, we *do not* attempt to provide a complete treatment of this issue. We believe that this topic is sufficiently rich and, thus, would require a communal effort that goes well beyond the scope of a single paper.

2 Measures of Interestingness

As pointed out in the introduction, it is important to study not only objective, but also subjective measures of interestingness of patterns because a pattern that is of interest to one user may be of no interest to another one. We identify two reasons why a pattern can be interesting to a user from the subjective point of view. It is either because the pattern is *unexpected* or because it is *actionable*. We now explain, at the intuitive level, what these two concepts mean and also explore how they are related to each other. We will base our discussions of these two subjective measures of interestingness on the following example that will be used throughout the paper.

Example 1 Consider the database of student evaluations of different courses offered at some university. Each semester, students do evaluations of all the courses they take, and the summaries of these evaluations are stored in the database as reports shown in

Fall 1994

Course No :	CS101
Instructor:	John Doe
No of Students Registered:	42
No of Students Responded :	26
No of Forms Rejected :	0

Expected Grade Percentages				
A	B	C	D	F
34.8%	47.8%	13.0%	4.3%	0.0%

Average Expected Grade = 3.1

The Instructor: (1 - 10), The Course: (11 -16)	Mean	S.D.	Strong Disagreement 1	2	Percentage Responding 3	4	5	6	Strong Agreement 7
1. The instructor presented course materials in an organized fashion..	6.2	0.9	0.0	0.0	0.0	7.7	7.7	42.3	38.5
2. The instructor stimulated student interest in the course...........	5.6	1.3	0.0	3.8	0.0	15.4	11.5	38.5	23.1
3. The instructor explained difficult concepts effectively............	5.4	1.3	0.0	7.7	0.0	7.7	26.9	34.6	15.4
4. The instructor responded well to questions and comments............	6.1	0.8	0.0	0.0	0.0	3.8	11.5	50.0	30.8
5. The instructor was enthusiastic....................................	6.4	0.8	0.0	0.0	0.0	3.8	7.7	26.9	53.8
6. The instructor provided practical examples.........................	6.0	1.0	0.0	0.0	3.8	3.8	15.4	34.6	34.6
7. The instructor showed an interest in students......................	6.2	1.1	0.0	0.0	3.8	7.7	0.0	38.5	46.2
8. The instructor graded fairly.......................................	5.8	1.4	3.8	0.0	0.0	11.5	11.5	26.9	38.5
9. The instructor provided sufficient feedback........................	5.8	1.0	0.0	0.0	3.8	3.8	19.2	42.3	19.2
10. Overall, I would recommend this instructor........................	6.4	0.9	0.0	0.0	3.8	0.0	3.8	34.6	50.0
11. There was sufficient class participation..........................	5.6	1.0	0.0	0.0	0.0	15.4	26.9	30.8	23.1
12. The class sessions were worthwhile................................	5.8	0.9	0.0	0.0	0.0	3.8	38.5	26.9	26.9
13. This was a demanding course.......................................	5.9	1.0	0.0	0.0	0.0	11.5	15.4	42.3	30.8
14. The material I learned will be useful to me.......................	5.8	1.3	0.0	0.0	7.7	11.5	11.5	26.9	38.5
15. The text and other reading materials were worthwhile..............	5.3	1.6	7.7	0.0	3.8	7.7	26.9	26.9	23.1
16. Overall, I would recommend this course............................	5.7	1.5	3.8	0.0	0.0	11.5	11.5	26.9	30.8

Figure 2: A Course-Faculty Evaluation Report.

Figure 2. These reports are generated every semester for every course offered at that university. □

2.1 Actionability Measure

According to this measure, a pattern is interesting because the user can *do* something about it; that is, the user can react to it to his or her advantage. For example, the pattern that Professor X is consistently getting the overall instructor ratings (item 10 in Figure 2) below the overall course ratings (item 16 in Figure 2) can be of great interest to the chairperson of the department where X teaches because this shows to the chairperson that Professor X has "room" for improvement in his or her teaching ratings and X should be encouraged to work more on his teaching and presentation skills.

Actionability is an important subjective measure of interestingness because users are mostly interested in the knowledge that permits them to do their jobs better by taking some specific actions in response to the newly discovered knowledge. However, it is not the only important subjective measure of interestingness, as we discuss it in the next section.

2.2 Unexpectedness Measure

If a newly discovered pattern is surprising to the user, then it is certainly interesting (how many times we exclaimed "Oh, this is interesting..." when we discover something unexpected?). For example, if in most of the course evaluations, the overall instructor ratings (item 10) are *higher* than the overall course ratings (item 16), and it turns out that in most of Professor X's ratings overall instructor evaluations are *lower* than the overall course evaluations then such a pattern is unexpected and, hence, interesting. As another example, assume that in some course only 8% of the students responded with their evaluations, whereas this number is normally between 60% and 90%. This pattern is definitely interesting because it is certainly unexpected.

We maintain that unexpected patterns are interesting because they contradict our expectations which, in turn, depend on our system of *beliefs*. For example, in the "instructor evaluation" pattern above, we believe that the overall instructor ratings should be higher than the overall course ratings, whereas the pattern contradicts this belief. Similarly, the "response rate" pattern contradicts our belief that there should be a "reasonable" response rate from the students.

2.3 Relationship Between Unexpectedness and Actionability

Clearly, some patterns are unexpected *and* actionable at the same time. For example, the "instructor

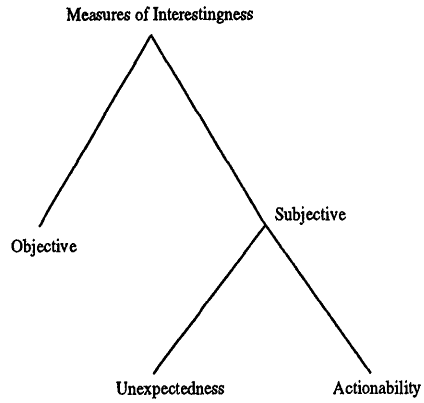

Figure 3: Classification of Interestingness Measures.

evaluation" pattern is actionable *and* unexpected at the same time.

Furthermore, some actionable patterns can be expected. For example, assume Professor Y is getting consistently high ratings, and the chairperson of the department wants to nominate him for the "Teacher of the Year" award. However before doing so, the chairperson wants to see the latest evaluations for Professor Y, and it turns out that they are also good. Such pattern (good student evaluations for Professor Y in the last semester) is expected. However, it is also actionable because the chair can nominate Professor Y for the award now without any reservations.

Also, a pattern can be unexpected and non-actionable. For example, the "response rate" pattern of only 8% is unexpected. However, it is not actionable for the chairperson of the department offering this course because student response rates are beyond his or her control and cannot be influenced by the chairperson in any way.

Thus, we can conclude that the two subjective measures of interestingness, "unexpectedness" and "actionability" of a pattern are, in general, *independent* of each other[2]. This conclusion leads to the classification of measures of interestingness as presented in Figure 3.

Both unexpectedness and actionability measures are important. However, we consider only unexpectedness for the remainder of the paper, leaving actionability as a topic of future research.

3 Defining Unexpectedness via Belief Systems

As indicated in Section 2.2, unexpectedness is related to beliefs. Intuitively, the more a pattern disagrees with a belief system, the more unexpected and hence the more interesting it is. Thus, to define interestingness of a pattern, we have to define beliefs first.

3.1 Beliefs

Beliefs and belief revision have been extensively studied in AI. In particular, there are two major approaches to defining beliefs and belief revision. In the first approach, based on the work of Alchourrón et al. [1] and others, we either believe in something or we don't. When a new belief is considered, it is either added to the set of previous beliefs in case it does not contradict them, or some of the previous beliefs have to be removed to accommodate a new one in case it contradicts some of the previous beliefs.

In contrast to the first approach, the second approach to the theory of beliefs and belief revision assumes that we can believe in certain statements only *partially*. In other words, it assigns some *degree* or *measure*, or a *confidence factor* to each belief. The two most prominent approaches to assigning such a measure to a belief are the Bayesian approach [5] and the the Dempster-Shafer theory of evidence approach [9]. In the Bayesian approach, the degree of belief is associated with the conditional probability that the belief holds given some previous "evidence" for this belief, whereas in the Dempster-Shafer theory a *belief function* [9] is assigned to beliefs which, in general, does not satisfy the axioms of probability theory.

In this paper, we follow the second approach and associate some confidence measure with each belief. We also assume that beliefs are arbitrary predicate formulae expressed in first-order logic. In addition, we assume that there are two types of beliefs, *hard* and *soft* beliefs, which are defined below.

Soft Beliefs. These are beliefs that the user is willing to change as new patterns are discovered that provide the user with new evidence. For example, the belief that the overall instructor ratings are higher than the overall course ratings is a soft belief because it can be changed as new evidence (new grades) is reported every semester. We assign a *degree* (or *measure*) to each soft belief that specifies our "confidence" in it. Furthermore, we will adopt the Bayesian approach and assume that the degree of belief is measured with the conditional probability [5]. This means that if α is a belief and ξ is the previous evidence "supporting" that belief, then the degree of belief in α is $P(\alpha \mid \xi)$. For example, α can be the belief that "the overall instructor ratings are higher than the overall course ratings," and ξ can be the "evidence" consisting of all the instructor and course ratings over several semesters of studies and over all the courses offered by that university.

Given new evidence E, we update the degree of belief in α, $P(\alpha|E,\xi)$, using the Bayes rule:

[2]Although we believe that most unexpected patterns are also actionable and most actionable patterns are unexpected.

$$P(\alpha|E,\xi) = \frac{P(E|\alpha,\xi)P(\alpha|\xi)}{P(E|\alpha,\xi)P(\alpha|\xi)+P(E|\neg\alpha,\xi)P(\neg\alpha|\xi)} \quad (1)$$

Example 2 Let α_0 be the belief that "the overall instructor ratings are higher than the overall course ratings," and assume that the degree of this belief is 0.85 ($P(\alpha_0|\xi) = 0.85$). Assume that a new course evaluation arrived containing a new evidence E_0 that the overall instructor rating (item 10 in Figure 2) for the course taught by this instructor is 5.8 and the overall course rating (item 16 in Figure 2) is 5.2. We associate two functions with belief α_0 computing the conditional probabilities $P(E|\alpha_0,\xi)$ and $P(E|\neg\alpha_0,\xi)$ for evidence E. Assume that it turns out that for the new evidence E_0 $P(E_0|\alpha_0,\xi) = 0.62$ and $P(E_0|\neg\alpha_0,\xi) = 0.42$. Then, substituting these numbers into (1), we compute $P(\alpha_0|E_0,\xi) = 0.89$. □

The Bayesian degree of belief defined above has the following properties that follow directly from the definitions:

Property 1. Positive evidence strengthens the belief: if $\alpha \models E$ then $P(\alpha \mid E,\xi) \geq P(\alpha \mid \xi)$. □

Property 2. Negative evidence weakens the belief: if $\alpha \models \neg E$ then $P(\alpha \mid E,\xi) \leq P(\alpha \mid \xi)$. □

Hard Beliefs. The hard beliefs are the *constraints* that cannot be changed with new evidence. In fact, if new evidence contradicts these beliefs, then there must be some mistakes made in acquiring this new evidence. For example, if it turns out that the number of students responding to the evaluation survey is greater than the number of students registered for that class (two items in the upper-left corner of the evaluation report in Figure 2), then it means that the accuracy of this report is highly questionable and that its data is wrong. We would like to stress that hard beliefs, as soft beliefs, are *subjective* and vary from one user to another. Also, we do not associate any degree with hard beliefs because the user never changes beliefs of this type.

3.2 Interestingness of a Pattern Relative to a Belief System

As we stated in Section 2.2, intuitively, a pattern is interesting relative to some belief system if it "affects" this system, and the more it "affects" it, the more interesting the pattern is. We distinguish between hard and soft beliefs, and thus treat these two cases separately.

Hard Beliefs. If a pattern contradicts the set of hard beliefs of the user then this pattern is *always* interesting to the user. In other words, if B_H is a set of logical sentences defining hard beliefs of a user and if p is a pattern (that is also a logical sentence), then if $B_H \models \neg p$, then pattern p is interesting. Note that a contradicting pattern does not affect hard beliefs and that hard beliefs are kept unchanged. Such a contradiction means that the data used to derive the pattern must be wrong. For example, in the "response rate" pattern described in Section 3.1 something must have been wrong with the data because the number of responding students cannot be greater than the number of students registered for the course.

Soft Beliefs. Based on our previous discussion of the intuitive meaning of interestingness, we formally define interestingness of pattern p relative to a (soft) belief system B as

$$I(p,B) = \sum_{\alpha\in B} \frac{|P(\alpha|p,\xi) - P(\alpha|\xi)|}{P(\alpha|\xi)} \quad (2)$$

where the sum is taken over all of the soft beliefs α in B. This definition of interestingness measures by how much a new pattern p changes degrees of our beliefs.

Example 3 In Example 2, it was shown that if $P(\alpha_0|\xi) = 0.85$ then $P(\alpha_0|E_0,\xi) = 0.89$. Substituting these numbers into (2), we obtain the interestingness of the pattern E_0 (that in the new faculty-course evaluation report the instructor ratings (5.8) were higher than the course ratings (5.2)) relative to belief α_0 as $I(E_0,\alpha_0) = 0.047$.

Assume that the evaluation report described in Example 2 has the overall course rating of 5.8 and the overall instructor rating 5.2 (we will denote this evidence E_1). Assume that it turns out that $P(E_1|\alpha_0,\xi) = 0.42$ and $P(E_1|\neg\alpha_0,\xi) = 0.62$. Then $P(\alpha_0|E_1,\xi) = 0.79$ and, substituting these numbers into (2), we obtain $I(E_1,\alpha_0) = 0.076$. □

Note that in Example 3 the pattern that contradicts the belief that instructor ratings are usually higher than the course ratings is more interesting than the pattern that confirms this belief (since $0.047 = I(E_0,\alpha_0) < I(E_1,\alpha_0) = 0.076$). This observation, that unexpected patterns are more interesting than expected patterns, is formalized in the following theorem.

Theorem 1 *Let α be a belief, such that $0.5 < P(\alpha|\xi) < 1$. Let p be a pattern confirming belief α ($\alpha \models p$). Also assume that the function evaluating $P(p|\alpha,\xi)$ is symmetric with respect to α, that is for all p, $P(\neg p|\alpha,\xi) = P(p|\neg\alpha,\xi)$. Then $I(p,\alpha) \leq I(\neg p,\alpha)$*[3]*. Moreover, the equality holds if and only if $P(p|\alpha,\xi) = P(\neg p|\alpha,\xi)$.*

[3]Note that pattern p is expected in this context because it follows from α and $\neg p$ is unexpected because it contradicts α.

Patterns

		Contradictory	Non-contradictory
Beliefs	**Hard**	change data	accept data
	Soft	check data	accept data

Figure 4: Actions to be Taken With the Data When a New Pattern is Discovered.

Patterns

		Contradictory	Non-contradictory
Beliefs	**Hard**	do nothing	do nothing
	Soft	depends on the data check (see text)	update degree of belief

Figure 5: Actions to be Taken With Beliefs When a New Pattern is Discovered.

Sketch of Proof: It follows from Properties 1 and 2 that, to prove the theorem, we should show that

$$\frac{P(\alpha|p,\xi) + P(\alpha|\neg p,\xi)}{2} \leq P(\alpha|\xi)$$

Substituting formula (1) into this expression and using the symmetry properties $P(\neg p|\alpha,\xi) = P(p|\neg\alpha,\xi)$ and $P(p|\alpha,\xi) = P(\neg p|\neg\alpha,\xi)$, this inequality is reduced to $(P(p|\alpha,\xi) - P(\neg p|\alpha,\xi))^2 \geq 0$. □

4 How Discovered Patterns Affect Beliefs and the Underlying Data

In Section 3, we defined interestingness of a pattern in terms of its unexpectedness. In this section, we examine what should be done with the belief system and with the underlying data on which the pattern is based, once an unexpected pattern is discovered.

In order to understand the actions that need to be taken, we divide the beliefs into hard and soft and the discovered patterns into the ones that contradict the beliefs and the ones that do not contradict them.

The actions to be taken with respect to the underlying data when an unexpected pattern is discovered are summarized in Figure 4. In particular, if a discovered pattern contradicts the hard beliefs, then this means that the underlying data on which the pattern is based is wrong and needs to be *changed.* For example, as was pointed out in Section 3.1, if the number of students responding to the evaluation survey is greater than the number of students registered for the course, this means that something is wrong with the data and that the data should be changed appropriately. If the pattern does not contradict hard or soft beliefs, then the data should be left unchanged (accepted). Moreover, if a pattern contradicts a soft constraint, then it is an indication that the data should be possibly *checked.* For example, if we have a soft belief that cars are usually driven at a speed of 50 to 80 mph on American freeways, and if we discover that somebody has driven a distance of 250 miles in two hours (which contradicts our soft belief), then we may want to check the data to make sure that there are no mistakes made in the process of recording this data.

The actions to be taken with respect to the belief system when an unexpected pattern is discovered are summarized in Figure 5. In particular, if the discovered pattern contradicts the hard beliefs, then there is nothing one can do to these beliefs because these beliefs are unchangeable constraints (laws), and the data supporting the pattern must be wrong. If the pattern does not contradict the hard beliefs then, clearly, nothing needs to be done with these beliefs. If the pattern contradicts the soft beliefs, then, as we pointed out before, a check may be necessary to test whether the pattern is based on the correct data. If it is based on the wrong data, then the beliefs remain intact. If the pattern is based on the correct data, then the degrees of soft beliefs should be updated as specified in Section 3.1. Finally, if the pattern does not contradict soft beliefs then the degrees of soft beliefs should be updated as specified in Section 3.1.

5 Future Research Directions

Since the purpose of this paper is to stimulate a discussion in the knowledge discovery community about the subjective measures of interestingness, we presented only the overall approach to subjective measures of interestingness and did it mostly informally. We plan to proceed with a formal study of these concepts only after some consensus among the KDD researchers is reached on the intuitive meanings of various measures of interestingness.

There are many important questions that have to be addressed in order to better understand subjective measures of interestingness. First, how to for-

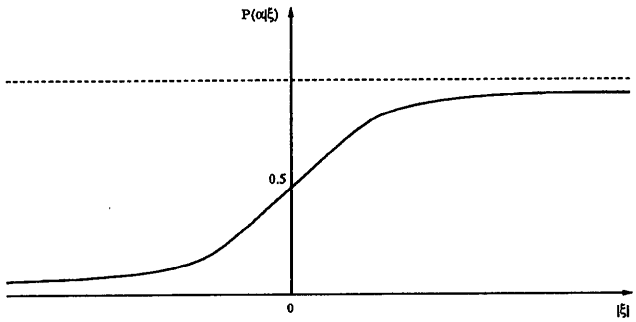

Figure 6: Relationship Between the Degree of Belief and the "Measure of Evidence."

malize actionability and understand better how it is related to unexpectedness. Second, how to compute interestingness of a pattern relative to a belief system and how to do this efficiently. Third, how to maintain a belief system in an efficient manner. Fourth, what should the structure of a belief system be, and how it affects the computations of interestingness of a pattern. Fifth, how would subjective and objective measures of interestingness be combined into one integral measure. Sixth, how can we use the belief systems discussed in this paper to discover incorrect and corrupted data (i.e. how can we do "data cleaning"). Seventh, formalize the relationship between the degree of belief and the "measure of evidence." Intuitively, if there is neither supporting nor contradicting evidence for some belief α, then the degree of belief is $P(\alpha \mid \xi) = 0.5$, where $P(\alpha \mid \xi)$ is a conditional probability and ξ is a supporting evidence. In this case the "measure of evidence," $|\xi| = 0$. Intuitively, the relationship between $P(\alpha \mid \xi)$ and $|\xi|$ is as shown in Figure 6 – the more supporting evidence ξ for α there is, the closer $P(\alpha \mid \xi)$ is getting to 1, and similarly, the more there is negative evidence ξ for belief α, the more $P(\alpha \mid \xi)$ approaches 0. One of the important issues is to find a good measure of evidence, $|\xi|$, and then, possibly, redefine the measure of interestingness (2) in terms of $|\xi|$.

These are just a few questions that need to be addressed in order to understand subjective measures of interestingness better, and we believe that we took the first step towards this goal by raising all of these issues in the paper.

6 Conclusions

We classified different measures of interestingness of patterns into objective and subjective and identified two subjective reasons why a pattern can be interesting to the user – either because it is unexpected or because it is actionable. We also argued, at the intuitive level, that these two measures of interestingness are independent of each other and that both of them are important. As a starting point, we studied unexpectedness in this paper, leaving actionability as a topic of future research. To define interestingness in terms of unexpectedness, we considered a belief system and defined interestingness of a pattern in terms of how much it affects the belief system.

Acknowledgments

We thank Gregory Piatetsky-Shapiro for useful discussions of some of the issues addressed in the paper.

References

[1] C. Alchourrón, P. Gardenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

[2] V. Dhar and A. Tuzhilin. Abstract-driven pattern discovery in databases. *IEEE Transactions on Knowledge and Data Engineering*, 5(6), 1993.

[3] W.J. Frawley, G. Piatetsky-Shapiro, and C.J. Matheus. Knowledge discovery in databases: an overview. In G. Piatetsky-Shapiro and W.J. Frawley, editors, *Knowledge Discovery in Databases.* AAAI / MIT Press, 1991.

[4] J. Hong and C. Mao. Incremental discovery of rules and structure by hierarchical and parallel clustering. In G. Piatetsky-Shapiro and W.J. Frawley, editors, *Knowledge Discovery in Databases.* AAAI / MIT Press, 1991.

[5] E. J. Horvitz, J. S. Breese, and M. Henrion. Decision theory in expert systems and artificial intelligence. *Journal of Approximate Reasoning*, 2, July 1988.

[6] C. J. Matheus, G. Piatetsky-Shapiro, and D. McNeill. An application of KEFIR to the analysis of healthcare information. In *Proceedings of the AAAI-94 Workshop on Knowledge Discovery in Databases*, 1994.

[7] G. Piatetsky-Shapiro. Discovery, analysis, and presentation of strong rules. In G. Piatetsky-Shapiro and W.J. Frawley, editors, *Knowledge Discovery in Databases.* AAAI / MIT Press, 1991.

[8] G. Piatetsky-Shapiro and C. J. Matheus. The interestingness of deviations. In *Proceedings of the AAAI-94 Workshop on Knowledge Discovery in Databases*, pages 25–36, 1994.

[9] G. Shafer. *A mathematical theory of evidence.* Princeton University Press, 1976.

[10] P. Smyth and R. M. Goodman. Rule induction using information theory. In G. Piatetsky-Shapiro and W.J. Frawley, editors, *Knowledge Discovery in Databases.* AAAI / MIT Press, 1991.