Advanced Review

Actionable knowledge discovery and delivery

Longbing Cao*

Actionable knowledge has been qualitatively and intensively studied in the social sciences. Its marriage with data mining is only a recent story. On the one hand, data mining has been booming for a while and has attracted an increasing variety of increasing applications. On the other, it is a reality that the so-called knowledge discovered from data by following the classic frameworks often cannot support meaningful decision-making actions. This shows the poor relationship and significant gap between data mining research and practice, and between knowledge, power, and action, and forms an increasing imbalance between research outcomes and business needs. Thorough and innovative retrospection and thinking are timely in bridging the gaps and promoting data mining toward next-generation research and development: namely, the paradigm shift from *knowledge discovery from data* to *actionable knowledge discovery and delivery*. © 2012 Wiley Periodicals, Inc.

How to cite this article:
WIREs Data Mining Knowl Discov 2012, 2: 149–163 doi: 10.1002/widm.1044

ISSUES WITH CURRENT KNOWLEDGE DISCOVERY FROM DATA

Actionable knowledge 'is not only relevant to the world of practice, it is the knowledge that people use to create that world'.[1] Actionable knowledge is not a new concept in social science and business domains. It has been discussed very intensively in areas such as business management,[2,3] organization science,[4] management science,[5,6] and alike. However, the engagement of actionable knowledge with data mining has only taken place in recent years,[7–9] especially in retrospection on deliverable effectiveness in supporting decision-making action-taking.

Data mining seeks to extract interesting patterns from data. In recent decades, data mining has boomed as an emerging discipline, and has been featured by an increasingly large number of publications. In applications, however, we see only a few commercially available data mining products (some of them including statistical aspects) on the market. These products often lead to patterns, or so-called knowledge discovered in data, which either evidence existing domain observations or commonsense, or cannot be used for taking decision-making actions at all. This reflects the emergence of an extreme and overwhelming imbalance between a massive number of research publications and rare workable products/systems.

In recent years, more and more data mining researchers with strong practices or first-hand industrial experience have recognized the critical problems and challenges associated with data mining research.[10–19] This retrospection has triggered the first round of a paradigm shift, namely, from 'data mining' to 'knowledge discovery',[18] to discover hidden and interesting knowledge from data. Typical research includes research on subjective measures,[20–22] unexpectedness,[23,24] novelty,[25] actionable rules,[26] action rules,[27,28] and interpretability.[11] However, it is argued that the 'knowledge' discovered from data is not powerful enough for 'direct' and 'decisive'[29] problem-solving. Here, 'direct' means there is no need to further manipulate the knowledge discovered and 'decisive' indicates decision-making actions for desired results when it is used for problem-solving.

In industry and business, an even more obvious trend is that more and more practitioners are urging the transformation from data to actionable knowledge, to make data mining useful for real-world applications[30] in health,[31] retail,[29] intrusion detection,[32] web logs[33] and to generally upgrade organizational competitive advantage.[5]

*Correspondence to: longbing.cao-1@uts.edu.au

Advanced Analytics Institute, University of Technology, Sydney, Australia

DOI: 10.1002/widm.1044

The reality facing the knowledge discovery in databases (KDD) community seems to be unpleasant. The imbalanced situation has never been thoroughly addressed; rather, it seems more and more serious.[14] From the system science perspective, KDD-based problem-solving is a system, which involves not only data itself, but also the environment of knowledge discovery, and delivering decisions to application needs. However, in the traditional KDD framework, data are the focus, and the environment is often simplified, if not ignored. Decisions are often too simple, straightforward or far removed from the expectations of business problem-solving. A typical and increasing scenario is that a very limited proportion of the increasingly enormous number of annual publications in the community is able to support operable/workable decision-making in the real world.

In KDD projects, we often face the following scenarios:

- Data miner: 'I have found something interesting!' 'Many patterns have been found!' 'They satisfy my technical metric thresholds very well!'
- Business people: 'So what?', 'They are just commonsense.' 'I don't care about them.' 'I don't understand them.' 'How can I use them?'

From the technical and engineering perspective, many issues seem to be overlooked in classic data mining. Let us take financial data mining as an example:

- Problem dynamics and interaction in a system: market dynamics such as coupling between two stock prices are often overlooked in modeling.
- Problem environment: a time-series model is built on closing prices which is applied to represent the market dynamics.
- Business processes, organizational factors, and constraints: a trading pattern is discovered without differentiating among types of orders, market, and investors.
- Human involvement: a trading rule captures historical trading patterns but not the investor's intention.
- Knowledge discovered: if frequent pattern mining is used to identify frequent trading patterns, we may find very many 'interesting' patterns. However, almost all of them may be unworkable, because they may just reflect the majority's trading behaviors which are not interesting to investors.
- Evaluation: a frequent trading pattern with high confidence but low *Sharpe ratio* when applied to a market.

The above problems come from the gap between academic objectives and business goals, and between academic outputs and business expectations. To utilize its unique power of enabling smart businesses and transforming business and industry by providing smart decisions, it is worthwhile asking ourselves what is 'wrong' with 'discovering knowledge from data'? What is 'inconsistent' between the underlying KDD methodologies, research intentions, and focus and the needs of real-life problem-solving? Why and where does the imbalance occur? How do these gaps arise? Most importantly, how should the existing KDD paradigm be transformed into one that can product actionable knowledge for decision-making?

GAP ANALYSIS

To gain a better understanding of knowledge actionability,[34] we explore the gaps appearing in data mining. Although it is complicated to scrutinize true reasons and to discover effective solutions for the above 'wrongness' or 'inconsistency', deep gap analysis and thorough retrospection about traditional KDD, and therefore innovative thinking and interdisciplinary interaction, are helpful to determine possible actions.

Gaps between Delivered and Desired

There may be gaps between knowledge, power, and action[35,36] in existing data mining methodologies. Let us firstly try to understand where the gaps are located. We observe this from the macrolevel by focusing on methodological issues surrounding traditional KDD research.[17]

On the one hand, from the research culture perspective, we often:

- Concentrate on innovative algorithms and patterns;
- Only check the interestingness of identified patterns from the technical significance perspective;
- Do not really perceive or care about the needs of business people;
- Do not take the business environment into account; or

- Oversimplify data, surroundings, and problem definition.

On the other hand, practitioners and business analysts value something different, for example,

- Can it solve my business problem, or will it lead to what I expect?
- Has it considered the surrounding social, environmental, and organizational factors?
- Can I interpret it in my business language, experience, and knowledge?
- Can I easily adjust it as I need by following my business rules and processes?
- Can it make my job more efficient rather than causing new issues?
- Can it be easily integrated into my business rules, operational systems, and workflow?
- What could be the impact on business if I use it? Is that manageable?

Consequently, we see the gaps between academia and business in goals, factors involved, outputs, deliverable presentation modes, evaluation, and impact:

- Gap between a converted research issue and its actual business nature;
- Gap between academic objectives and business goals;
- Gap between technical significance and business interest;
- Gap between identified patterns and deliverables expected by business; and
- Gap between the deliverables from data miners and the eventual entities deployed into problem-solvers.

The above gaps result in the imbalance emerging in the KDD community and this imbalance embodied in many aspects, for instance,

- Algorithm imbalance: many published algorithms versus few that are actually workable in the business environment;
- Pattern imbalance: many patterns mined versus a very small portion if any is eventually used;
- Evaluation imbalance: technical performance validated versus no check for business interest or business impact;
- Decision power imbalance: impressive (technical) performance claimed versus very few that can either be used directly or be converted to support decision-making actions and achieve business expectation.

The knowledge we often see is 'passive', presenting information on surface level with little context or background. Such passive knowledge really does not tell us much about how to act and on what to act upon. What decision-makers need is 'active' knowledge with power to work, which is compelling and powerful for action-taking and decision-making. To narrow them, what we need to do is to convert passive data (knowledge) into active knowledge or directly produce active knowledge.

Aspects for Narrowing Gaps

The above gap analysis shows that it is not easy to discover actionable knowledge.[37] With respect to such gaps, let us discuss what aspects can be explored further to narrow down the gaps. We observe this from both macrolevel and microlevel perspectives.[14,17,38]

On the macrolevel, aspects are related to methodological and fundamental issues, including key elements: environment, human role, process, infrastructure, dynamics, evaluation, risk policy, and deliverability.

- Environment: Refers to any factors surrounding data mining models and systems; for instance, domain factors, constraints, expert groups, organizational factors, social factors, business processes, and workflows. Some factors such as constraints have been considered in current data mining research, but many others have not. It is essential to represent, model, and involve them in KDD systems and processes.
- Human role: To handle many complex problems, human-centered and human-mining-cooperated KDD is necessary. Challenging problems related to this include how to involve domain experts and expert groups in the mining process, and how to allocate the roles between human and mining systems.
- Process: Real-world problem-solving has to cater for dynamic and iterative involvement of environmental elements and domain experts along the way.
- Infrastructure: The engagement of environmental elements and humans at run time in a dynamic and interactive way requires an

open system with closed-loop interaction and feedback.[39] KDD infrastructures need to provide facilities to support such scenarios.

- Dynamics: To deal with the dynamics in data distribution from training to testing and from one domain to another is essential, in domain and organizational factors, human cognition and knowledge, the expectation of deliverables, and in business processes and systems.
- Evaluation: Interestingness needs to be balanced between technical and business perspectives from both subjective[22] and objective[40,41] aspects; special attention needs to be paid to deliverable formats, their actionability, and generalizable capability, as well as to securing the support of domain experts.
- Risk: Risk needs to be measured in terms of its presence and magnitude, if any, in conducting a KDD project and system.
- Policy: Data mining tasks often involve policy issues such as security, privacy and trust which exist not only in the data and environment, but also in the use and management of data mining findings in an organization's environment.
- Delivery: This includes determining the right form of delivery and presentation of KDD models and findings so that end users can easily interpret, execute, utilize, and manage the resulting models and findings, and integrate them into business processes and production systems.

On the microlevel, aspects related to technical and engineering issues that support KDD need to be addressed. Listed below are a few dimensions that address these concerns: architecture, procedure, interaction, adaptation, actionability, and deliverability.

- Architecture: KDD system architectures need to be effective and flexible for incorporating and consolidating specific environmental elements, KDD processes, evaluation systems, and final deliverables.
- Procedure: Tools and facilities supporting the KDD process and workflow are necessary, from business understanding, data understanding, and human–system interaction to the assessment, delivery, and execution of deliverables.
- Interaction: To cater for interaction with business people throughout the KDD process, appropriate user interfaces, user modeling and servicing are required to support individuals and group interactions.
- Adaptation: Data, environmental elements, and business expectations change all the time. KDD systems, models, and evaluation metrics are required to be adaptive for handling differences and changes in dynamic data distributions, cross domains, changing business situations, and user needs and expectations.
- Actionability: What do we mean by 'actionability'? How can we measure it? What is the tradeoff between technical and business sides? Do subjective and objective perspectives matter? This requires essential metrics and integration mechanisms to be developed.
- Deliverable: End users certainly feel more comfortable if the models and patterns delivered can be presented in a business-friendly way and be compatible with business operational systems and rules. In this sense, it is necessary for KDD deliverables to be easily interpretable, convertible into or presented in a business-oriented way such as business rules, and to be linked to decision-making systems.

The above discussions contribute to this insight: KDD-based problem-solving is expected to be a process and system for discovering and delivering actionable knowledge.[42] Such *actionable knowledge discovery and delivery* (AKD) needs to systematically consider/involve problems, data, environment, model and decisions, as well as optimization[43] in KDD. This brings us to the methodology of *domain driven data mining*.[14]

AN AKD FRAMEWORK

To enable the discovery of actionable knowledge, AKD is proposed to narrow the gaps in KDD. As a framework for AKD, domain-driven data mining (D^3M)[14] has been proposed to analyze the underlying problems and challenges facing traditional KDD methodology and systems to develop appropriate methodology and techniques to tackle the problems and changes that will enable AKD as well as the real-life problem-solving and decision support by KDD deliverables.

AKD Problem Statement

Rather than focusing on what happens in the current KDD, we prefer to observe the nature of

KDD-based problem-solving from an interdisciplinary perspective, by integrating the methodologies from other disciplines including system sciences, cybernetics, and complex systems.[44] This perspective produces a multidimensional view of AKD-based problem-solving. AKD is a six-dimension-based optimization process o[14–17,45,46]:

$$AKD ::= optimization(problem, data, environment, model, decision). \quad (1)$$

1. *Problem*: a problem is the KDD target, composed of data, business, environment and needs, contributing to corresponding KDD objectives and business analytical goals, and to the eventual evaluation and validation of KDD findings.
2. *Data*: extracted in a business problem, reflecting a mapping between a business problem-based space and an extracted/converted object-based space for decisions; risk often starts from conducting the mapping.
3. *Environment*: a problem and its data are enclosed in a certain environment, embodied through organizational or social factors, which need to be considered for a complete and genuine understanding of the problem; however, they are easily neglected, filtered, destroyed, or simplified in data extraction and analysis.
4. *Model*: a model is an appropriate tool to connect the data to proper decisions by addressing the underlying problem within an environment; a model is biased if it is not fully reflective of the problem, data, environment, or decision.
5. *Decision*: a decision is presented in terms of identified patterns or knowledge through KDD, which is believed to be an overarching solution addressing the underlying problem.
6. *Optimization*: optimization seeks a perfect match between the model and problem, data, environment, and decision against expectations.

Let us further discuss how the above AKD framework can be executed in the real world. First, AKD is a problem-solving process that progresses business problems (Ψ, with problem status τ) to problem-solving solutions (Φ)[15,47,48]:

$$\Psi(.|\tau) \rightarrow \Phi(). \quad (2)$$

The process of finding problem-solving solutions ($\Phi()$) is a procedure to find the *actionable pattern set* $\widetilde{P}$ through employing all valid models M.

$$\mathrm{AKD}^{m_i \in M} \longrightarrow \mathrm{O}_{p \in P}\mathrm{Act}(p), \quad (3)$$

where $P = P^{m_1} U P^{m_2}, \cdots, U P^{m_n}$, $\mathrm{Act}(.)$ is the evaluation function, $O(.)$ is the optimization function to extract actionable pattern $\tilde{p}$ ($\tilde{p} \in \tilde{P} \subset P$), where $\mathrm{Act}(\tilde{p})$ beats a given benchmark.

For a pattern p, $\mathrm{Act}(p)$ can be further measured in terms of *technical actionability* ($t_a(p)$) and *business actionability* ($b_a(p)$).

$$\mathrm{Act}(p) = I(t_i(p), b_i(p)), \quad (4)$$

where $I(.)$ is an objective function for aggregating the contributions of all particular aspects of actionability from problem, data, environment, model, and decision.

Further, $\mathrm{Act}(p)$ can be described in terms of *objective* (o) and *subjective* (s) factors from both *technical* (t) and *business* (b) perspectives:

$$\mathrm{Act}(p) = t_o(\mathbf{x}, \widetilde{p}) \wedge t_s(\mathbf{x}, \widetilde{p}) \wedge b_o(\mathbf{x}, \widetilde{p}) \wedge b_s(\mathbf{x}, \widetilde{p}), \quad (5)$$

where operator '$\wedge$' indicates the 'aggregation' of the specific aspect of actionability. We say that p is truly *actionable* (i.e., $\widetilde{p}$) both to academia and business if it satisfies the following condition:
IF

$$\forall p \in \widetilde{P}, \exists \mathbf{x} : t_o(\mathbf{x}, p) \wedge t_s(\mathbf{x}, p) \wedge b_o(\mathbf{x}, p) \wedge b_s(\mathbf{x}, p) \rightarrow Act(p) \quad (6)$$

THEN:

$$p \rightarrow \widetilde{p}. \quad (7)$$

Further, let $\widetilde{P} = \{\tilde{p}_1, \tilde{p}_2, \cdots, \tilde{p}_Z\}$ be an *actionable pattern set* mined by method m_n for the given problem Ψ (its data set is DB), in which each pattern $\tilde{p}_z$ is *actionable* for the problem-solving if it satisfies the following conditions:

1.a. $t_i(\tilde{p}_z) \geq t_{i,0}$; indicating the pattern $\tilde{p}_z$ satisfying technical actionability t_i with threshold $t_{i,0}$;

1.b. $b_i(\tilde{p}_z) \geq b_{i,0}$; indicating the pattern $\tilde{p}_z$ satisfying business actionability b_i with threshold $b_{i,0}$;

1.c. $R : \tau_1 \xrightarrow{A, m_n(\tilde{p}_z)} \tau_2$; the pattern can support business problem-solving (R) by taking action A, and can correspondingly transform the problem status from initially nonoptimal state τ_1 to greatly improved state τ_2.

Actionable knowledge (patterns) can lead to effective actions for better results (decision, answer, conclusion, etc.). The process of discovering actionable knowledge, or AKD, forms a framework that engages knowledge discovery on data with problem and environment toward optimal evaluation and decisions to satisfy both technical and business expectation from objective and subjective perspectives. This differentiates AKD from the normal KDD process.

Due to the inconsistency that often exists in different aspects, we frequently find that the identified patterns only fit into one of the following subsets:

$$\mathrm{Act}(p) \rightarrow \{\{t_i^{\mathrm{act}}, b_i^{\mathrm{act}}\}, \{\neg t_i^{\mathrm{act}}, b_i^{\mathrm{act}}\}, \{t_i^{\mathrm{act}}, \neg b_i^{\mathrm{act}}\}, \{\neg t_i^{\mathrm{act}}, \neg b_i^{\mathrm{act}}\}\} \quad (8)$$

where '¬' indicates the corresponding element is not satisfactory. In real-world data mining, it is often very challenging to find the most actionable patterns that are associated with both 'optimal' t_i^{act} and 'optimal' b_i^{act}. Clearly, AKD favors patterns confirming the relationship $\{t_i^{\mathrm{act}}, b_i^{\mathrm{act}}\}$.

Actionability Computing

Actionability means the power to work, which is an optimal outcome and objective from AKD through the best integration of six core dimensions. Consequently, actionability is also embodied through each dimension and its integration:

1. *Actionability on problem* reflects the depth and width of our understanding of the underlying problem, its surroundings, constraints, and expected outcomes from AKD.
2. *Actionability on data* reflects the depth and width of our understanding of the underlying data complexity, structure, volume, dimensionality, type, speed, and dynamics.
3. *Actionability on environment* reflects the depth and width of our understanding of human, domain, organizational, and social aspects, as well as interactions and dynamics surrounding the problem and data.
4. *Actionability on model* reflects the quality of the models selected to understand the problem, data, and environment.
5. *Actionability on decision* reflects the operational power of the AKD deliverables for direct and effective problem-solving.
6. *Actionability on optimization* reflects the best mapping from the underlying problem to the expected decisions made by AKD models and the best combination of all dimensions.

In essence, actionability is the quality and power of AKD outcomes for effective decision-making and problem-solving. For different purposes, actionability may be interpreted in terms of varying terms, for instance,

- *Autonomy* of the deliverables for direct use in an unattended problem-solving process or system,
- *Deliverability* and *transferability* of the identified patterns and knowledge from data miners to business people, and from one domain to another,
- *Dependability* of the identified patterns and knowledge,
- *Explainability* and *interpretability*[49] of the identified patterns and knowledge,
- *Impact* of the deliverables leading to what is expected by business,
- *Repeatability* of the proposed algorithms and methods,
- *Semantics* and *understandability* of deliverables for seamless integration into business ontology and machine-based understanding and use, and
- *Trust* of the proposed algorithms and methods, as well as identified patterns and knowledge, without security and privacy offense and risk to the underlying problem and environment.

Actionability computing therefore emerges as an interesting research issue in AKD. While from the quantitative perspective, Formula (5) reflects the overall interpretation of actionability computing, there are many open issues to be further explored; for instance, how to represent and quantify the trust, autonomy, semantics quality of my fraud detection model for online banking fraud control. This leads to many new opportunities, as suggested by the following topics, for further exploration in creating actionable knowledge[6]:

- What are the key attributes for actionable knowledge? How can knowledge be both scientifically rigorous and practically useful?
- One aspect of actionable knowledge involves disseminating our research to practitioners so that they understand it and are willing to act on it. What forms of diffusion are most effective for this purpose? What kinds of communication and messages gain practitioners'

attention and understanding? What are the mechanisms for translating research into practice?

- What research methods are likely to contribute to actionable knowledge? How open are we to different research methods? How can research questions be formulated and examined so the subsequent findings are likely to be implemented?
- A good deal of actionable knowledge is tacit and exists only in practice. How do we capture and make sense of such knowledge? How do we study it scientifically?
- Generating actionable knowledge involves an inherent tension between two radically different cultures: science that seeks knowledge that is internally valid and generalizable, and practice that asks for useful answers to situation-specific problems. How might these competing demands be managed so that there is greater appreciation and dialogue between the two cultures? What does each culture stand to gain and lose from interacting with the other? What should be the relationship between practitioners and researchers?
- How can practitioners help researchers formulate, conduct, and disseminate their research in more actionable ways? How can they inform researchers about the tacit dimensions of their practice? What valuable lessons can practitioners teach researchers, and how can this be done so that researchers will listen?
- How can we help practitioners become better consumers of knowledge about management? Can they be inoculated against fads?
- What can our scholarly journals do to close the gap between research and practice? Should authors be held accountable for reflecting on the action possibilities of their findings? Should the implications for practice be more than an afterthought?

AKD Concept Map

A high-level concept map for developing AKD methodology and techniques consists of the following layers: domain problem, ubiquitous intelligence, theoretical foundation, supporting technique, and actionability computing:

- Domain problem: This, in general, targets complex knowledge from complex data in domain-specific applications and problems that cannot be well-handled by existing data mining and knowledge discovery techniques. Such problems may include domain problems from retail to government to social network, from either a sector or a specific business problem perspective.
- Ubiquitous intelligence: This refers to the intelligence surrounding AKD problem-solving, from data to domain, organizational, social and human aspects, and the representation, synthesis, and consolidation of respective intelligence for AKD-based problem-solving.
- Theoretical foundation: This refers to the fundamental theories to enable AKD, either borrowed from many relevant disciplines from the information sciences to social sciences, or invented for data sciences and analytics sciences, targeting the establishment of a family of scientific foundations for dealing with increasingly emergent complexities and challenges in data and analytics.
- Supporting technique: This refers to AKD techniques and tools to engage and consolidate ubiquitous intelligence, support knowledge representation and deliverables, cater for project and process management, and implement decision-making pursuant to the findings.
- Actionability computing: This refers to the quantification of the decision-making power of identified knowledge and deliverables by AKD, and the presentation, delivery, and impact of AKD findings for direct decision-making.

A serious address of the above key components in AKD demands the engagement and support needed to cater for problem, data, environment, model, decision, and optimization in KDD, and the reshaping of KDD processes, modeling and outcomes from technical, procedural, and business perspectives.

Ubiquitous Intelligence

The success of AKD relies on involving and integrating ubiquitous intelligence[14,50] in a domain-specific application. This involves data intelligence, domain intelligence, network intelligence, human intelligence, and social intelligence, as well as the synthesis of ubiquitous intelligence.

Data intelligence indicates interesting information and stories about a business problem formation or driving forces. Typical efforts are on handling data complexity such as on large-scale, multidimensional/high-dimensional, online/real time, social media, multimedia, dynamic, highly frequent, uncertain, noisy, mixed structure aspects. Apart from the usual focus on exploring complexity from data structure, quantity, speed, and characteristics from the individual data object perspective, coupling and interaction between data objects have not been seriously considered, yet they present challenges to AKD, such as what may happen if dependency between objects is considered in a similarity-based clustering process.

Domain intelligence emerges from domain factors and resources that not only wrap a problem and its target data but also assist in problem understanding and problem-solving. Domain intelligence involves qualitative and quantitative aspects. These are instantiated in terms of aspects such as domain knowledge, background information, prior knowledge, expert knowledge, constraints, organization factors, business process, and workflow, as well as environment intelligence, business expectation, and interestingness.

Human intelligence refers to explicit or direct involvement of human empirical knowledge, belief, intention, expectation, run-time supervision, evaluation, and expert groups in AKD. It also concerns the implicit or indirect involvement of human intelligence such as imaginary thinking, emotional intelligence, inspiration, brainstorm, reasoning inputs, and embodied cognition such as convergent thinking through interaction with other members in dynamic data mining and assessing identified patterns.

Network intelligence emerges from both web intelligence and broad-based network intelligence such as information and resource distribution, linkages among distributed objects, hidden communities and groups, information and resources from network, and, in particular, the web, information retrieval, searching, and structuralization from distributed and textual data. The information and facilities from the networks surrounding the target business problem either consist of the problem constituents, or contribute to useful information for actionable knowledge discovery. Therefore, they should be catered for in AKD.

Social intelligence refers to the intelligence that lies behind group interactions, behaviors[51] and corresponding regulation. Social intelligence covers both human social intelligence and animat/agent-based social intelligence. Human social intelligence is related to aspects such as social interaction, group goals and intention, social cognition, emotional intelligence, consensus construction, and group decision. Social intelligence is often associated with social network intelligence and collective interaction, as well as business rules, law, trust, and reputation for governing the emergence and use of social intelligence.

The use of ubiquitous intelligence may take one of the following two paths: *single intelligence engagement* and *multiaspect intelligence engagement*. Examples of single intelligence engagement are the involvement of domain knowledge in data mining and the consideration of user preferences in data mining. Multiaspect intelligence engagement aims to integrate ubiquitous intelligence as needed. It is very challenging but inevitable in mining complex enterprise applications. It is often very difficult to integrate every type of intelligence into one data mining system, in addition to the challenges of modeling and involving a specific type of intelligence. New data mining methodologies and techniques need to be developed to involve ubiquitous intelligence in AKD. The theory of metasynthetic engineering,[44, 52–54] agent mining,[55–58] and integration of ubiquitous intelligence[59] may provide useful clues for synthesizing ubiquitous intelligence in the AKD process.

DEPLOYMENT

Opportunities

The gaps between the aims of AKD and the existing situation of KDD research and development disclose great opportunities for AKD research and development. We list the following observations from which we can further develop KDD:

- Complex applications: While any applications could be linked to AKD, we are particularly interested in complex applications (represented in the complexity of a problem, data, or environment). Complex enterprise applications will propose major functional and nonfunctional requirements that cannot be handled by existing KDD approaches, and will drive the development of AKD toward novel and effective methodologies, algorithms, and tools.[54]
- Complex data: Data is becoming more complex in many aspects from type, volume, structure, speed, and dimensionality to dynamics. More powerful tools are needed to tackle such data complexity, as well as to consider similarity, dependency, and interaction between data points.
- Complex behaviors: Behavior can be seen everywhere, and is an essential object in analyzing applications and data. There are

limited techniques available for effective general analysis and mining of complex behaviors, especially the representation and reasoning of complex behaviors, and the learning and mining of coupled behaviors, behavior networking, group behaviors, behavior convergence and divergence, impact analysis of group behaviors, and detection of complex behavior interaction patterns in group behaviors.

- Complex environments: Next-generation knowledge discovery will have to discover knowledge in complex environments, mixing elements of human, domain, organizational, and societal factors. Environment complexities present characteristics such as dynamics that need to be catered for in the architecture, model, and process.
- Actionability measure: New performance metrics will be developed to quantify the actionability[34] of KDD deliverables and services which are dependable, user friendly and business friendly, explainable, actionable, reliable, safe, trustworthy, repeatable, and transferable.
- Deliverable semantics: Besides patterns, the delivery format and semantics will be important issues to study, so that not only are the deliverables and properties of deliverables specified, they can also be transparently and seamlessly integrated into the operational environment for decision-making. Semantic interfaces and services[60] may be developed to embed AKD deliverables into operational systems.
- Decision power: The decision power of AKD deliverables is not determined by performance such as accuracy; rather, it is the utility and usability that can be directly taken over by business people and plugged into the operational environment, leading to better decisions or expected impact.
- AKD as service: Next-generation knowledge discovery may work in a cloud environment, which will trigger the development of knowledge discovery as a service for problem-solving and decision-making in an organization. In a more standard situation, rather than fostering a knowledge discovery team in every organization, knowledge discovery services will be provided by AKD specialists in a highly advanced knowledge discovery center. This requires the development of corresponding infrastructure, networking and privacy-processing facilities, and protocols for defining, communicating, subscribing, and monitoring services.

AKD Architectures

To support the involvement of ubiquitous intelligence and the delivery of actionable knowledge, it is essential to develop effective system architectures for constructing AKD systems, and effective techniques for supporting AKD. In the following, we briefly introduce a few flexible frameworks.[14,17,47,48]

Postanalysis-based AKD (PA-AKD)[61,62] is a two-step pattern extraction and refinement exercise. First, generally interesting patterns (which we call 'general patterns') are mined from data sets by technical interestingness ($t_o()$, $t_s()$) associated with the algorithms used. The mined general patterns are then pruned, distilled and summarized into operable business rules (embedding actions) (which we call 'deliverables') in terms of domain-specific business interestingness ($b_o()$, $b_s()$) and involving domain and meta knowledge.

Unified interestingness-based AKD (UI-AKD) develops unified interestingness metrics, which are defined for capturing and describing both business and technical concerns. The mined patterns are further converted into deliverables based on domain knowledge and semantics. UI-AKD looks just the same as the normal data mining process except for three inherent characteristics. One is the interestingness system, which combines technical significance ($t_i()$) with business expectations ($b_i()$) into a unified AKD interestingness system ($i()$). This unified interestingness system is then used to extract truly interesting patterns. The second is that domain knowledge and environment must be considered in the data mining process. Finally, the outputs are actionable patterns and operable business rules.

Combined interestingness-based AKD (CM-AKD) comprises multisteps of pattern extraction and refinement on the whole data set. First, J steps of mining are conducted based on business understanding, data understanding, exploratory analysis, and goal definition. Second, generally interesting patterns are extracted based on technical significance ($t_i()$) (or unified interestingness ($i()$)) into a pattern sub-set (P_j) in step j. Third, knowledge obtained in step j is further fed into step $j+1$ or relevant remaining steps to guide the corresponding feature construction and pattern mining (P_{j+1}). Fourth, after the completion of all individual mining procedures, all identified pattern subsets are merged into a final pattern set based on environment, domain knowledge, and business expectations (b_i). Finally, the merged patterns are converted into business rules as final deliverables (patterns and business rules) that reflect business preferences and needs.

AKD Implementation

Corresponding AKD techniques are necessary to fulfill the methodology of AKD.

Constrained Knowledge Delivery Environment

Actionable knowledge is discovered in a constraint-based context that mixes environmental reality, expectations, and constraints in the knowledge discovery and delivery process. Specifically, several types of constraints play a significant role in the AKD process: *domain constraints*, *data constraints*, *interestingness constraints*, and *deliverable constraints*. Efforts are needed to develop both generic and domain-specific tools and systems to cater for these constraints.

Cooperation between Human and KDD Systems

The real-life requirements for discovering actionable knowledge in a constraint-based environment determine that real-world data mining is more likely to follow a man–machine-cooperation mode; in other words, a human–mining cooperation rather than an automated process and system. Human involvement is embodied through the cooperation between humans (including users and business analysts, mainly domain experts) and a data mining system. This is because of the complementation between human qualitative intelligence such as domain knowledge and field supervision, and the quantitative intelligence of KDD systems such as computational capabilities. Therefore, real-world complex data mining presents as a human–mining-cooperated interactive knowledge discovery and delivery process. The tasks are to develop theories and tools to support the involvement of humans and human intelligence into an AKD system.

Interactive and Parallel KDD Support

To support AKD, it is important to develop interactive mining support that involves domain experts and human–mining interaction. Interactive facilities are also useful for evaluating data mining findings by involving domain experts in a closed-loop manner. On the other hand, parallel mining support is often necessary for dealing with concurrent applications, distributed, and multiple data sources. In cases with intensive computation requests, distributed and parallel mining[63] can greatly upgrade real-world data mining performance. There are huge areas to explore in terms of developing interactive, visual, parallel, and distributed systems and tools for AKD in a complex environment such as those involving multiple sources of data, information, resources, and humans.

Closed-Loop and Iterative Refinement

Actionable knowledge discovery in a constraint-based context is more likely to be a closed-loop process rather than an open loop one. A closed-loop process indicates that the outputs of data mining are fed back to adjust/inform relevant parameter or factor tuning in particular stages. It is worthwhile to study what should be in the loop, how to engage and optimize the components in the loop, and when to terminate the iteration.

Mining In-Depth Patterns

Greater effort is essential to uncover in-depth patterns in data. 'In-depth patterns' (or 'deep patterns') are not straightforward and can only be discovered through more powerful models following thorough data and business understanding and effectively involving domain intelligence or expert guidance. An example is to mine for insider trading patterns in capital markets. Without deep understanding of the business and data, a naive approach is to analyze the price movement change by partitioning data in terms of preevent, on the event and postevent. A deeper pattern analysis on such price difference analysis may involve domain factors such as considering market or limit orders, market impact, fusion of price, index and announcement information, and checking the performance of *potential abnormal return*, *liquidity*, *volatility*, and *correlation*.

Post Mining

Post mining[62] handles the following problems: How to read and understand discovered patterns, which are often in thousands or more? What are the most interesting ones? Is the model accurate and what does the model tell us? How should we use the rules, patterns, and models? To answer these questions and present useful knowledge to users, it is necessary to conduct post mining to further analyze the learned patterns, evaluate the built models, refine and polish the built models and discovered rules, summarize them, and use visualization techniques to make them easy to read and understand.

Combined Mining

Combined mining[48] is one of the general methods of analyzing complex data for identifying complex knowledge. The deliverables of combined mining are *combined patterns*. For a given business problem (Ψ), we suppose there are the following key entities associated with it in discovering interesting knowledge for business decision-support: data set D, feature set F, method set R, interestingness set I, impact set T, and pattern set P. On the basis of the above variables, a general pattern discovery process can be described as follows: patterns $P_{n,m,l}$ are identified through data mining method R_l deployed on features F_k from a dataset D_k in terms of interestingness $I_{m,l}$[48]:

$$P_{n,m,l} : R_l(F_k) \rightarrow I_{m,l}, \tag{9}$$

where $n = 1, \ldots, N$; $m = 1, \ldots, M$; $l = 1, \ldots, L$.

From a high-level perspective, combined mining represents a generic framework for mining complex patterns in complex data.

Agent-Driven Actionable Knowledge Discovery

Agent-driven actionable knowledge discovery[53,59,64–70] can contribute to the problem-solving of many data mining issues, for example, multiagent data mining infrastructure and architecture, multiagent interactive mining, multiagent-based user interaction, automated pattern mining, multiagent-distributed data mining, multiagent dynamic mining, multiagent mobility mining, multiagent multiple data source mining, multiagent peer-to-peer data mining, and multiagent web mining. Agent technology can help with these challenges by involving autonomy, interaction, dynamic selection and gathering, scalability, multistrategy, and collaboration. Other challenges include privacy, mobility, time constraint (stream data, it is too late to extract and then mine), and computational costs and performance requests.

Knowledge Discovery as Service

This is the era of efficiently mining for actionable knowledge to support specific, ad hoc and intention-driven needs for organizational decision-making by involving distributed, highly heterogeneous, dynamic, and ubiquitous intelligence in complex environments, such as intranets, extranets, cloud, and grid. This will require breakthroughs in managing and integrating different resources from different channels, implementing knowledge discovery as services, and making services available for satisfying different and changing requirements.

Knowledge Delivery

Well-experienced data mining professionals attribute the weak executable capability of existing data mining findings to the lack of proper tools and mechanisms for implementing the ideal deployment of the resulting models and algorithms by business users rather than analysts. In fact, the barrier and gap comes from the weak, if not nonexistent, capability of existing data mining deployment systems and services, found in the presentation, deliverable, and execution aspects. They form the AKD delivery system, which is much beyond the identified patterns and models themselves:

- Deliverable: studies how to deliver data mining findings and systems to business users so that the findings can be readily re-formatted, transformed, or cut and pasted into their own business systems to be presented on demand, and ensures that the systems can be understood and taken over by end users;
- Presentation: studies how to present data mining findings that can be easily recognized, interpreted, and taken over as needed;
- Execution: studies how to integrate data mining findings and systems into production systems, and how the findings can be executed easily and seamlessly in an operational environment.

Supporting techniques need to be developed for AKD presentation, deliverability, and execution. For instance, the following lists some such techniques:

- Deliverable: business rules are widely used in business organizations, and one method for delivering patterns is to convert them into business rules; for this, we can develop a tool with underlying ontologies and semantics to support the transfer from pattern to business rules;
- Presentation: typical tools such as visualization techniques are essentially helpful; visual mining could support the whole data mining process in a visual manner;
- Execution: tools to make deliverables executable in an organization's environment need to be developed; one such effort is to generate Predictive Model Markup Language (PMML) to convert models to executables so that the models can be integrated into production systems, and run on a regular basis to provide cases for business management.

TABLE 1 | Basket Analysis Database

Sequence ID	Date	Customer Role	Address	Shopping Cart
C_1	2011-10-1	Father	A	Beer, Diaper, Banana, Harry Potter, iPhone
C_2	2011-10-2	Mum	A	Apple, Cherry, Blackberry, Plum
C_3	2011-10-3	Son	A	Pencil Case, Rubber, Lego, Scooter
C_4	2011-10-2	Mum	B	Pear, Cherry, Peach, Plum, Melon, Apple
C_5	2011-10-4	Father	B	Beer, iPhone, Fish, Meat
C_6	2011-10-1	Son	B	Scooter, Pen, Notebooks

- Communication plan: a document to craft the right information, sell the right stories, communicate the value of the deliverables for targeted clients, and set up possible goals to achieve, services, products, processes, and tools to disseminate or share.[29]

AN EXAMPLE

Basket analysis is a typical application for illustrating the power of association mining for business applications. In practice, frequent association rules such as $\{diaper, beer\}$ are arguably not profitable for a department store because putting two such items as these together would significantly reduce the potential for purchasing other items on the way from the bottle shop to the baby product section. In many countries, it is even not permitted to put beer and diapers in the same section.

Let us take a fruit retail shop as an example and explain how to make rules more actionable. We extract shopping transactions from 100 customers, and identify the following association rules by using the Apriori algorithm:

R_1: {apple, banana, 95}

R_2: {apple, apple-mango juice, 80}

R_3: {apple, Harry Potter, 20}

Which rule would you choose if you were the shop owner? According to the association rule theory, R_1 is recommended because it has the highest *support*. In practice, no shop owner would follow this rule to put apples and bananas on the same shelf, since they are normally in the same section, and owners may want to promote other fruit during a customer's search path from apple to banana. Thus, R_1 should be filtered. Let us further add the item price: apple—$2/kg, apple–mango juice—$4/bottle, Harry Potter—$39.99/copy. By adding the prices, we have R'_2 and R'_3:

R'_2: {(apple, $2), (apple–mango juice, $4), 80}

R'_3: {(apple, $2), (Harry Potter, $39.99), 20}

Suppose each customer only definitely buys 1 unit per shopping, which rule would you choose if you were the shop owner? Technically, R'_2 is recommended because $\mathrm{supp}(R_2) > \mathrm{supp}(R_3)$. From the revenue perspective, R'_3 is more profitable: $\mathrm{profit}(R_2) = \$48 > \mathrm{profit}(R_3) = \$84$. This shows the importance of involving additional data and evaluating the business impact of the findings during pattern mining.

The above findings arise from the traditional association rule framework: frequency-based similarity analysis, which treats transactions independently. In reality, we know that different family members may visit the same shopping center to purchase consumables for the family's use. This makes the transactions from the same family dependent on each other. Let us look at a retail database as shown in Table 1. If $min\_supp = 2$, based on traditional association rule mining, we have three shopping patterns for individual customers:

R_4: {Beer, iPhone}

R_5: {Cherry}

R_6: {Scooter}

However, if we consider the dependency between customers in terms of family role and address:

$$C_1\text{-Father}, C_2\text{-Mum}, C_3\text{-Son};$$
$$C_1\text{-Mum}, C_2\text{-Father}, C_3\text{-Son};$$
$$C_1.\text{Address} = C_2.\text{Address} = C_3.\text{Address};$$
$$C_4.\text{Address} = C_5.\text{Address} = C_6.\text{Address};$$

and the constraint that a family member does not purchase the same item that other members have purchased for the same week:

$$C_1.\text{Week} = C_2.\text{Week} = C_3.\text{Week};$$
$$C_4.\text{Week} = C_5.\text{Week} = C_6.\text{Week};$$

we have a single family shopping pattern:

R_7: {Father: Beer, iPhone; Mum: Cherry; Son: Scooter}

R_4, R_5, and R_6 have the same concerns as R_1. By involving object relationships, obviously this family pattern is more informative for a shopping center's marketing campaign. For instance, the following marketing strategy may be designed to promote iPhone:

Definition 1 *(Marketing Strategy 1)*

If: *Father buys Beer, Mum buys Cherry, and Son buys Scooter*
Then: *30% off for another iPhone*

or a promotion pricing package for family shopping:

Definition 2 *(Marketing Strategy 2)*

Individual: *Beer-$30/box, iPhone-$900, Cherry-$15/kg, Scooter-$100*
Package: *Beer-$30/box, iPhone-$900, Cherry-$15/kg, Scooter-$0*

The above sample marketing strategies illustrate the delivery of AKD-based association mining for improving retail business.

CONCLUSION

AKD is a challenging and fundamental task in data mining and the knowledge discovery revolution. It involves discovering and delivering knowledge that can be transparently and seamlessly incorporated into an operational environment for smart information use and decision-making. It is distinguished from traditional data mining by the extensive involvement of ubiquitous intelligence in the KDD process, and the elevation of the decision power of the findings.

AKD complement other data mining methodologies and techniques by explicitly involving human, domain, organizational and social intelligence, and their synthesis rather than simplifying or overlooking these factors in the modeling and evaluation process. This empowers the end user and decision-maker to easily understand the outcomes, and take actions on the findings to enable the transparent and seamless integration of data mining outputs into operational environment.

Direct engagement of ubiquitous intelligence and the delivery of actionable knowledge into research and development have not been investigated fully, however. As a result, we are experiencing increasing gaps and imbalance between business expectation and problem-solving needs and research focus and deliverables. We argue that this is because much of the research has been driven by innovating algorithms and exploring new data complexities, rather than focusing on the intrinsic complexities and challenges of the underlying problems and the expectations from enterprise applications.

With the already substantial development of algorithms and techniques, next-generation data mining will arguably benefit from practice-based research innovation and development, concentration on the underlying business problem itself, reduction of oversimplification and abstraction of the problem in constructing solutions, and in-depth observations of core challenges. This may lead to breakthrough methodology and techniques for data mining, and render the outputs workable and actionable for real-life problem-solving.

Many opportunities may materialize during the deep investigation of fundamental problems. These may lie in those areas widely explored, such as the involvement of domain knowledge in the dynamic data mining process, as well as emerging issues such as engaging organizational and social intelligence in the KDD modeling process. This makes furthering the paradigm shift from knowledge discovery toward AKD a very promising proposition.

ACKNOWLEDGMENTS

This work is sponsored in part by Australian Research Council Discovery Grants (DP1096218 and DP0988016) and ARC Linkage Grant (LP100200774 and LP0989721).

REFERENCES

1. Argyris C. *Knowledge for Action: A Guide to Overcoming Barriers to Organizational Change*. San Francisco, CA: Josssey-Bass; 1993.
2. Argyris C. Actionable knowledge: intent versus actuality. *J Appl Behav Sci* 1996, 32:441–444.
3. Argyris C. Actionable knowledge: design causality in the service of consequential theory. *J Appl Behav Sci* 1996, 32:390–406.
4. Cross, R, Sproull, L. More than an answer: information relationships for actionable knowledge. *Organ Sci* 2004, 15:446–462.
5. Morgan, K, Morabito, J, Merino, D. Creating actionable knowledge within the organization to achieve a competitive advantage. Available at: www.coalescentknowledge.com/Slides/Presentation-2.pdf. (Accessed January 13, 2012).
6. Cummings, T, Jones, Y. Creating actionable knowledge. Available at: http://meetings.aomonline.org/2004/theme.htm. (Accessed January 13, 2012).
7. He Z, Xu, X, Deng S. Data mining for actionable knowledge: a survey. Available at: http://arxiv.org/abs/cs/0501079. (Accessed January 13, 2012).
8. Cras Y. Turning raw data into actionable knowledge: known challenges and new complexities. In: *Proceedings of CSDM*. Springer Verlag, Heidelberg; 2010.
9. Barrett J, Caruthers E, German K, Hamby E, Lofthus R, Srinivas S, Ells E. From data to actionable knowledge: a collaborative effort with educators. In: *Proceedings of KDD*. ACM Press; 2011.
10. Adomavicius G, Tuzhilin A. Discovery of actionable patterns in databases: the action hierarchy approach. In: *Proceedings of KDD1997*. ACM Press; 1997, 111–114.
11. Aggarwal C. Towards effective and interpretable data mining by visual interaction. *ACM SIGKDD Explor Newslett* 2002, 3:11–22.
12. Ankerst M. Report on the SIGKDD-2002 panel—the perfect data mining tool: interactive or automated? *ACM SIGKDD Explor Newslett* 2002, 4:110–111.
13. Boulicaut J, Jeudy B. Constraint-based data mining. In: *The Data Mining and Knowledge Discovery Handbook*. Springer; 2005, 399–416.
14. Cao L. Domain-driven data mining: challenges and prospects. *IEEE Trans Knowl Data Eng* 2010, 22:755–769.
15. Cao L, Zhang C. The evolution of KDD: towards domain-driven data mining. *Int J Pattern Recognit Artif Intell* 2007, 21:677–692.
16. Cao L, Zhang C. Domain-driven actionable knowledge discovery in the real world. In: *Proceedings of PAKDD2006*. LNAI. Vol. 3918. Springer, Heidelberg; 2006, 821–830.
17. Cao L, Yu P, Zhang C, Zhao Y. *Domain Driven Data Mining*. New York: Springer; 2010.
18. Fayyad U, Smyth P. From data mining to knowledge discovery: an overview. In: Fayyad U, Piatetsky-Shapiro G, Smyth P, Uthurusamy R, eds. *Advances in Knowledge Discovery and Data Mining*. AAAI Press/The MIT Press, California; 1996, 1–34.
19. Fayyad U, Shapiro G, Uthurusamy R. Summary from the KDD-03 panel—data mining: the next 10 years. *ACM SIGKDD Explor Newslett* 2003, 5:191–196.
20. Silberschatz, A., Tuzihilin, A. On subjective measures of interestingness in knowledge discovery. In: *Proceedings of KDD*. ACM Press; 1995.
21. Silberschatz A, Tuzhilin A. What makes patterns interesting in knowledge discovery systems. *IEEE Trans Knowl Data Eng* 1996, 8:970–974.
22. Liu B, Hsu W, Chen S, Ma Y. Analyzing subjective interestingness of association rules. *IEEE Intell Syst* 2000, 15:47–55.
23. Padmanabhan, B., Tuzhilin, A. A belief-driven method for discovering unexpected patterns. In: *Proceedings of KDD-98*. ACM Press; 1998.
24. Suzuki, E. Autonomous discovery of reliable exception rules. In: *Proceedings of KDD*. ACM Press; 1997.
25. Al Hegami A, Bhatnagar V, Kumar N. Novelty framework for knowledge discovery in databases. In: *Proceedings of DaWaK*. Springer, Heidelberg; 2004.
26. Kaur, H. Actionable rules: issues and new directions. In: *Transactions on Engineering, Computing and Technology, World Information Society*. Vol. 5. 2005, 61–64.
27. Ras, Z.W., Wieczorkowska, A. Action rules: how to increase profit of a company. In: *Proceedings of PKDD00*. LNCS/LNAI. Vol. 1910. 2000, 587–592.
28. Tsay L S, Ras Z W. Action Rules Discovery System DEAR2, Method and experiments. *J Exp Theor Artif Intell* 2005, 17:119–128.
29. Dunoff A. Turning data into actionable knowledge, 2007. Available at: mcmorrowreport.com/WPs/WPActionableKnowledge.pdf
30. Dynamia2004, Generate actionable knowledge to improve business performance. Available at: www.statoo.com/en/PDF/Dynamia2004.pdf. (Accessed January 13, 2012).
31. Kumar, S. Transforming data for actionable knowledge, http://www.youtube.com/watch?v=DV8-Hk7HnAQ&noredirect=1. (Accessed January 13, 2012).
32. Julisch D. Intrusion detection alarms for actionable knowledge. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM Press; 2002, 366–375.
33. Yang Q., Ling, C., Gao, J. Mining web logs for actionable knowledge. In: Zhong N, J. Liu J, eds. *Intelligent Technologies for Information Analysis*. Heidelberg: Springer; 1998.

34. Cao L, Zhang C. Knowledge actionability: satisfying technical and business interestingness. *Int J Bus Intell Data Min* 2007, 2:496–514.
35. Saha B, Kakani, R. Knowledge, power and action: towards an understanding of implementation failures in a government scheme. *AI Soc* 2007, 21:72–92.
36. Montecel, R. Knowledge and action from dropping out to holding on. *IDRA Newslett*, Nov-Dec, 2006.
37. Sexton, M., Lu, S. The challenges of creating actionable knowledge: an action research perspective. *Constr Manag Econ* 2009, 27:683–694.
38. Kleinberg J, Papadimitriou C, Raghavan P. A microeconomic view of data mining. *data min knowl discovery* 1998, 2:311–324.
39. Xin D, Shen X, Mei, Q, Han J. Discovering interesting patterns through user's interactive feedback, In: *Proceedings of the 2006 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'06)*. ACM Press; 2006.
40. Freitas A. On objective measures of rule surprisingness. In: *Proceedings of PKDD98*. Springer, Heidelberg; 1998, 1–9.
41. Hilderman R, Hamilton H. Applying objective interestingness measures in data mining systems. In: *Proceedings of PKDD00*. Springer, Heidelberg; 2000, 432–439.
42. Yang Q, Yin J, Ling C, Pan R. Extracting actionable knowledge from decision trees. *IEEE Trans Knowl Data Eng* 2007, 19:43–56.
43. Freitas A. A critical review of multi–objective optimization in data mining– a position paper. *SIGKDD Explor* 2004, 6:77–86.
44. Cao L, Dai R, Zhou M. Metasynthesis: M–space, M-interaction and M-computing for open complex giant systems. *IEEE Trans Syst Man Cybern A* 2009, 39:1007–1021.
45. Cao L. Domain-driven actionable knowledge discovery. *IEEE Intell Syst* 2007, 22:78–89.
46. Cao L, Zhang C. Domain-driven data mining: a practical methodology. *Int J Data Warehousing Min* 2006, 2:49–65.
47. Cao L, Zhao Y, Zhang, H, Luo D, Zhang C. Flexible frameworks for actionable knowledge discovery. *IEEE Trans Knowl Data Eng* 2010, 22:1299–1312.
48. Cao L, Zhang, H, Zhao Y, Luo D, Zhang C. Combined mining: discovering informative knowledge in complex data. *IEEE Trans SMC Part B*. 2011, 41:699–712.
49. Yao Y, Zhao Y. Explanation-oriented data mining. In: Wang J ed. *Encyclopedia of Data Warehousing and Mining*. Hershey, PA: Idea Group Reference; 2005, 492–497.
50. Cao L, Luo D, Zhang C. Ubiquitous intelligence in agent mining. In: *Proceedings of ADMI 2009*. LNCS. Vol. 5680. Springer, Heidelberg; 2009, 23–35.
51. Cao L. In-depth behavior understanding and use: the behavior informatics approach. *Inf Sc* 2010, 180:3067–3085.
52. Qian X, Yu J, Dai R. A new discipline of science—the study of open complex giant system and its methodology. *Chin J Syst Eng Electron* 1993, 4:2–12.
53. Cao L, Dai R. Agent-oriented metasynthetic engineering for decision making. *Int J Inf Technol Decis Mak* 2003, 2:197–215.
54. Cao L, Dai R. *Open Complex Intelligent Systems*. Beijing: Posts and Telecom Press; 2008.
55. Cao L. Integrating agent, service and organizational computing. *Int J Softw Eng Knowl Eng* 2008, 18:573–596.
56. Cao L, Weiss G, Yu P. A brief introduction to agent mining. *J Auton Agents Multiagent Syst*, 2012, to appear.
57. Gorodetsky V, Karsaev O, Samoilov V. Multi-agent technology for distributed data mining and classification. In: *Proceedings of IAT 2003*. IEEE Computer Society Press; 2003, 438–441.
58. Symeonidis A, Mitkas P. *Agent Intelligence Through Data Mining*. Heidelberg: Springer; 2006.
59. Cao L. *Data Mining and Multiagent Integration*. Heidelberg: Springer; 2009.
60. Singh M, Huhns M. Service oriented computing: semantics, processes and agents. New York: John Wiley & Sons; 2005.
61. Liu B, Hsu W. Post-analysis of learned rules. In: *Proceedings of AAAI/IAAI*. Vol. 1. 1996, 828–834.
62. Zhao Y, Zhang C, Cao L. *Post-Mining of Association Rules: Techniques for Effective Knowledge Extraction*. Hershey, NY: Information Science Reference; 2009.
63. Kargupta H, Chan P, Kumar V. *Advances in Distributed and Parallel Knowledge Discovery*. California: AAAI Press/The MIT Press; 2000.
64. Cao L, Gorodetsky V, Mitkas P. Agent mining: the synergy of agents and data mining. *IEEE Intell Syst*, 2009, 24:64–72.
65. Cao L, Luo C, Zhang C. Agent-mining interaction: an emerging area. In: *Proceedings of AIS-ADM*. Springer, Heidelberg; 2007, 60–73.
66. Cao L, Gorodetsky V, Liu J, Weiss G. In: *Proceedings of ADMI2009*, LNCS. Vol. 5680. Springer, Heidelberg; 2009, 23–35.
67. Baik S, Cho J, Bala J. Performance evaluation of an agent based distributed data mining system. *Adv Artif Intell* 2005, 3501:91–99.
68. Dasilva J, Giannella C, Bhargava R, Kargupta H, Klusch M. Distributed data mining and agents. *Eng Appl Artif Intell* 2005, 18:791–807.
69. Davies W. Agent-based data-mining. Available at: http://www.agent.ai/doc/upload/200403/davi94_1.pdf. (Accessed January 13, 2012)
70. Davies W, Edwards P. Distributed learning: an agent-based approach to data-mining. In: *Proceedings of Machine Learning 95 Workshop on Agents that Learn from Other Agents*; 1995.

