Chapter 14
A Multiagent Approach to Adaptive Continuous Analysis of Streaming Data in Complex Uncertain Environments

Igor Kiselev and Reda Alhajj

Abstract The data mining task of *online unsupervised learning* of streaming data continually arriving at the system in complex dynamic environments under conditions of uncertainty is an $\mathcal{NP}$-hard optimization problem for general metric spaces and is computationally intractable for real-world problems of practical interest. The primary contribution of this work is a multi-agent method for continuous agglomerative hierarchical clustering of streaming data, and a knowledge-based self-organizing competitive *multi-agent system* for implementing it. The reported experimental results demonstrate the applicability and efficiency of the implemented adaptive multi-agent learning system for continuous online clustering of both synthetic datasets and datasets from the following real-world domains: the RoboCup Soccer competition, and gene expression datasets from a bioinformatics test bed.

14.1 Introduction

14.1.1 Problem Definition

Continuous decision-making and *anytime data analysis* in dynamic uncertain environments represent one of the most challenging problems for developing robust intelligent systems. It is indispensable for intelligent applications working in such

Igor Kisev
The David R. Cheriton School of Computer Science, University of Waterloo,
200 University Avenue West, Waterloo, Canada, ON N2L 3G1, e-mail: `ipkiselev@cs.uwaterloo.ca`

Reda Alhajj
Department of Computer Science, University of Calgary,
2500 University Drive NW, Calgary, Canada, AB T2N 1N4, e-mail: `alhajj@cpsc.ucalgary.ca`

[*] Supplementary materials, software demonstration, and video recordings of the developed multi-agent system are available at: `\sf\footnotesize\ttfamilyhttp://www.multiagent.org/mining`

L. Cao (ed.), Data Mining and Multi-agent Integration, DOI: 10.1007/978-1-4419-0522-2_14,

complex environments as autonomous robotic systems, dynamic manufacturing and production processes, and distributed sensor networks to be capable for successfully responding to environmental dynamics and making time critical decisions online under conditions of uncertainty. To develop a probabilistic theory of the operating environment (a representation of the world), an intelligent system solves the problem of *unsupervised learning*, that is the process of discovering significant patterns or features in the input data when no certain output or response categories (classes) are specified. The task of unsupervised clustering in statistical learning requires the maximizing (or minimizing) of a certain similarity-based objective function defining an optimal segmentation of the input data set into clusters [7].

There are two types of unsupervised clustering algorithms: *partitional* and *hierarchical*. The main goal of a partitional optimization algorithm can be defined by finding such assignments $\mathcal{M}^*$ of observations $\mathcal{X}$ to output subsets $\mathcal{S}$ that minimize a mathematical energy function, which characterizes the degree to which the clustering goal is not met and is the sum of the cost of the clusters: $\mathcal{W}(\mathcal{M}^*) = \sum_{i=1}^{k} c(\mathcal{S}_i)$. The problem of partitional clustering is known to be computationally challenging ($\mathcal{NP}$-hard) for general metric spaces and is computationally intractable for real-world problems of practical interest. In comparison to partitional clustering algorithms, the goal of a hierarchical optimization algorithm is to extract an optimal multi-level partitioning of data by producing a hierarchical tree $\mathcal{T}(\mathcal{X})$ in which the nodes represent subsets $\mathcal{S}_i$ of $\mathcal{X}$. The time and space complexities of the hierarchical clustering are higher than partitional one: in standard cases a typical implementation of the hierarchical clustering algorithm requires $\mathcal{O}(\mathcal{N}^2 \log \mathcal{N})$ computations.

The task of online learning in complex dynamic environments assumes near real-time mining of streaming data continually arriving at the system, which imposes additional requirements for continuous data mining algorithms of being sensitive to environmental variations to provide a *fast dynamic response to changes* with an event-driven incremental improvement of mining results (cf. Table 14.1).

Table 14.1: Required properties of online unsupervised learning methods

Feature	Classical methods	Required functionality
Model: • Data set • Decision criteria • Learning parameters	**Static:** • ***Static*** input data sets • ***Fixed*** and ***Single-objective*** criteria for learning • ***Invariable*** (cannot be changed at run-time)	**Dynamic:** • ***Dynamic*** and ***Streaming*** • ***Dynamic*** and ***Multi-objective*** quality metrics with trade-off balancing • ***Adjustable*** at run-time during algorithm execution
Method: • Learning mode • Availability of results • Reaction to changes	**Batch-oriented:** • ***Batch-oriented*** processing of a static data set • ***Only after full completion*** (needs time to get a result) • ***Must be restarted again*** from scratch with full retraining of models or extra repair methods	**Continuous:** • ***Near-real time learning*** of streaming data continually arriving at the system • ***At anytime*** during algorithm execution (always see a result and its improvement) • ***Without restarting***, and with event-driven ***incremental improvement*** of results, trading off operating time and result quality
Environment: • Uncertainty of learning • Data location	**Centralized and Deterministic:** • Assume predictable outcome (***Rigid***) • ***Single or Decentralized***, but with an additional centralized algorithm of aggregating partial results	**Distributed and Stochastic:** • Consider random effects (***Resilient***) • ***Both centralized and decentralized*** without any additional aggregation procedures

Additionally, a dynamic data mining algorithm, operating in complex uncertain environments with incomplete knowledge about parameters of the learning problem, should be suitable for *online exploratory data analysis* using different measures of similarity in order to be able to continually operate on the basis of various dynamic learning criteria.

The application of the proposed multi-agent solution to continuous online learning is appropriate for various scenarios of near real-time data processing: online intrusion detection, emergency response in hazardous situations (e.g. forest fires, chemical contaminants in drinking water), control of military operations with time critical targets, online learning of distributed robotic systems (e.g. in the RoboCup Soccer and Rescue domains), and run-time detection of previously unknown dispatching rules and effective scheduling policies in transportation logistics.

14.1.2 Related Work

Continuous and anytime data analysis imposes requirements for *adaptability* of learning methods that are simply not addressed by traditional data mining techniques. Conventional methods of unsupervised learning address the issue of statistical fluctuations of the incoming data by means of continual retraining of models that is computationally intractable or inappropriate in time-critical scenarios. Clustering results of batch-oriented methods are available only after their full completion, and must be started again from scratch in order to react on environmental variations.

To address this issue of effective online learning, various approaches have been considered in the literature (refer to online supplementary materials for a complete overview of related work [10]). Decentralized clustering algorithms were proposed to speeds up centralized learning by dividing it onto a set of processors and allowing them to learn concurrently with an additional centralized algorithm of aggregating partial mining results to the global solution [21], [12], [23]. To handle the complexity of the problem, approximate clustering algorithms were developed to search for a feasible solution in incomplete decision space, by applying approximation heuristics that reduce the problem dimension, but lead to worse result quality [7]. Clustering methods for unsupervised learning of streaming data were developed, which support incremental update of the mining result by applying additional repairing methods [1], [3]. Uncertainty of the operating environment is approached by feedback-directed clustering algorithms that apply reinforcement learning techniques to guide the search towards better cluster quality [4], [13]. The distributed-constraint reasoning formalism was proposed to approach optimization and learning problems in a decentralized manner, which is better suited to deal with changes in a localized fashion [5], but can be too expensive in large-scale dynamic environments [16] and restrictive to provide a fast response to environmental variations [22].

As opposed to previous work, we present a different anytime *multi-agent approach* to online unsupervised learning, which is different from conventional methods by being *dynamic*, *incremental* and *distributed*, rather than parallel . We demonstrate that the task of continuous unsupervised learning, when formulated as a *dynamic distributed resource allocation problem*, can be effectively approached by a decentralized market-based method of multi-agent negotiation [11].

14.2 Continuous Online Unsupervised Learning in Complex Uncertain Environments

14.2.1 Market-based Algorithm of Continuous Agglomerative Hierarchical Clustering

As opposed to previous work, we propose a different multi-agent approach to continuous online learning of streaming data by modeling the task of unsupervised clustering as a *dynamic distributed resource allocation problem* [15]. The online learning algorithm implements the concept of *clustering by asynchronous message-passing* [6], whereby the any-time solution to the continual constrained optimization problem of clustering is obtained (inferred) by satisfying a dynamic distributed constraint network of agent interests. Thus, the continual distributed learning process is carried out by means of asynchronous quasi-parallel processes of negotiation between the competitive agents of records and clusters, defined for data elements. Mining agents negotiate (act) with each other in the virtual learning marketplace in order to satisfy their individual goals and maximize their criteria values. Searching for the most profitable allocation variants (semantic links with the highest utility) to enhance their satisfaction levels with minimal costs, the agents of clusters and records dynamically establish and reconsider ontological relationships with other agents, thereby dynamically establishing ontological multilevel virtual market communities.

A distributed computational environment for self-interested agents (a virtual learning marketplace) is formally defined according to the game-theoretic notation of a *marketplace system* [8], [14]. To develop a multi-agent system capable of operating in dynamic distributed environments, we solve the task of online computational mechanism design in distributed computational settings, which assumes that *self-interested agents* can arrive at and depart from the multi-agent system dynamically over time, and there is no a trusted central mechanism to control their behavior. *Computational mechanism design* provides a mathematical framework that defines each agent's decision-making model and specifies the protocols that govern the agent interactions (the market mechanism through which agents interact) [17], [19]. A market-based algorithm of continuous agglomerative hierarchical clustering designs a multi-agent system in which rational self-interested agents with privately known preferences interact in a way that leads to equilibriums with desired system-wide properties (socially desirable outcome). Fig. 14.1 depicts an overview flowchart for the market-based algorithm of continuous agglomerative hierarchical clustering algorithm.

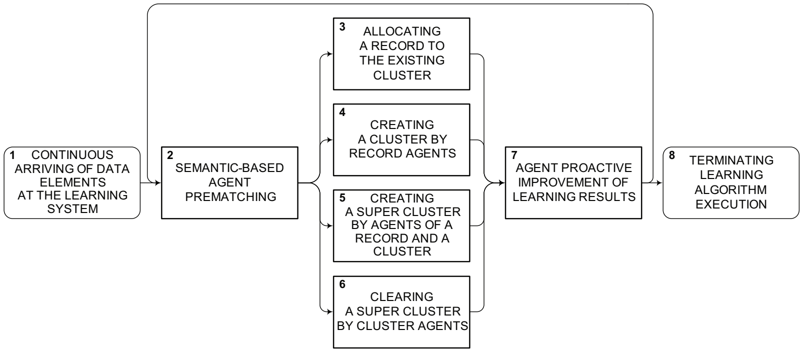

Fig. 14.1: A flowchart for the market-based algorithm of continuous hierarchical clustering

A global decision (macroscopic solution) of the dynamic distributed data clustering is implicitly achieved (performed) by the competitive agents that maintain a dynamic balance among the interest of all participants in the interaction according to the following algorithmic process (cf. Fig. 14.1): continuous arriving of data elements at the system (algorithmic step #1), locating candidate agents for allocation negotiations (algorithmic step #2), satisfying a dynamic *distributed constraint network* of agent interests for each type of agent negotiations (algorithmic steps #3, #4, #5, #6), agent proactive improvement of learning results (algorithmic step #7), and terminating the execution of the learning algorithm (algorithmic step #8). The complete description of each algorithmic step of the method is omitted in the chapter due to space limitations (refer to online supplementary materials).

The implemented auction-based negotiation method of agent negotiations is based on a modified Contract-Net Protocol, where agents dynamically submit bids based on the cost of possible allocation variant [20]. Each participant of the negotiation evaluates new allocation options and sends an approval only when the criteria value of a new agent state with a new established link and broken previous relationships with other agents (if any) is better than the value of the current agent state. If new ontological instance is created as a result of a negotiation process of the synthesis type a new mining agent of a corresponding type is produced in the virtual clustering marketplace and assigned to it.

The proposed computationally efficient multi-agent algorithm for online agglomerative hierarchical clustering is different from conventional unsupervised learning methods by being *distributed*, *dynamic*, and *continuous*. Distributed clustering process provides the ability to perform efficient run-time learning from both centralized and decentralized data sources without an additional centralized algorithm of aggregating partial mining results. Both the input dataset of decentralized sources and decision criteria for learning (e.g. similarity matrices and expert knowledge) are not fixed and can be changed at run-time during execution of the dynamic algorithm. Clustering results of the *adaptive learning algorithm* are available at any time and continuously improved to achieve a global quasi-optimal solution to the optimization problem, trading-off operating time and result quality.

14.2.2 Agent Decision-making Model

Goal-driven behavior of autonomous agents is supported by the developed microeconomic *multi-objective decision-making model*, which defines for each agent in the virtual marketplace ontology its individual goals, criteria, preference functions, and decision-making strategies [9]. Semantic agents of records have the goal to establish the most profitable allocation with the agents of clusters according to their individual agent criteria (”to be allocated”). To accomplish the allocation goal, a record agent can either send a membership application to the existing cluster to join it (algorithmic step #3) or be allocated to a new cluster, which can be created as a result of a negotiation process of the synthesis type with either another record agents or existing cluster agents (algorithmic steps #4 and #5). To support hierarchical clustering, there are two different goals defined for a cluster agent within its decision-making model (bidirectional $\lambda - \pi$ inference). The first goal of a cluster agent (”allocate”) is to establish links with the agents of records to create the most profitable ontological cluster of the best quality. The second goal of a cluster agent (”to be allocated”) has the same notion as the goal of a record agent, to establish the most profitable allocation with the agents of clusters, and is defined through the task of establishing the most effective relationship of the ”part-of” type with another cluster agent (algorithmic step #6). Fig. 14.2 illustrates the agent learning goals and basic types of agent negotiations in the virtual clustering marketplace.

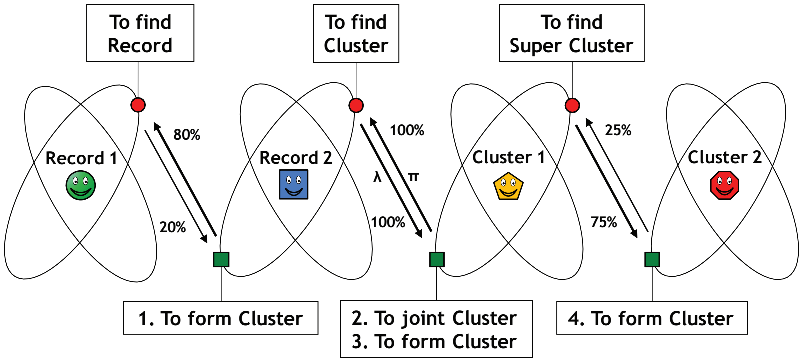

Fig. 14.2: Hierarchical architecture of a virtual clustering marketplace: agent goals and learning tasks

The developed multi-objective decision-making model makes it possible for the learning algorithm to continually operate on the basis of *dynamic learning criteria*, and to be suitable for online exploratory data analysis in complex uncertain environments with incomplete knowledge about parameters of the learning problem (cf. Fig. 14.3). Competitive agents of records and clusters act in the virtual clustering marketplace to satisfy their individual goals and maximize their criteria values according to the chosen decision-making strategy. Currently supported agent decision-making strategies are based on the following agent criteria: the Euclidian distance-based measure of similarity, the Chebychev similarity metric, and the angle metrics defining polarization (”shape”) of agent communities (multilevel and multicultural) in decision-space. Agent decision-making strategies can be applied dynamically at run-time to the whole agent society (global level), to a single agent (individual level), or to agent groups in different areas of the virtual clustering marketplace (several polarization vectors).

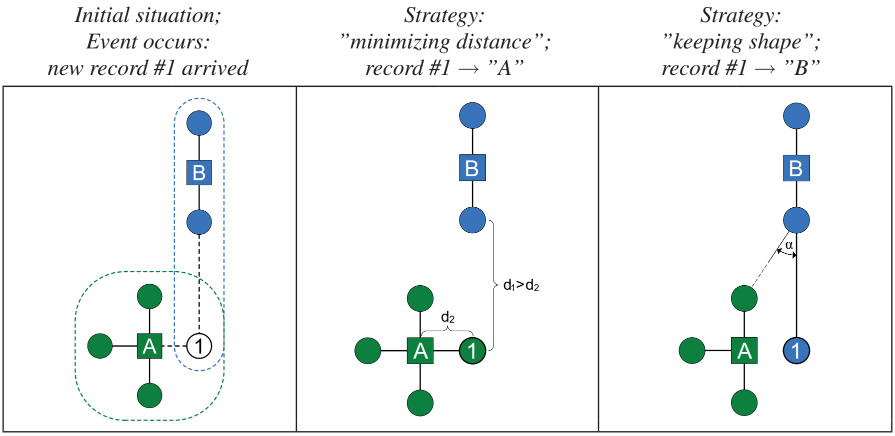

Fig. 14.3: Dynamic support for various agent decision-making strategies and learning criteria

Utilizing the agent criterion, based on the Euclidian distance-based measure of similarity, the learning algorithm relies on the ability to calculate the centroid of each cluster. For the input spaces where it is not possible, we propose to use the Chebychev similarity metric as a basis for a similarity measure between objects. The Chebychev similarity measure is based on the definition of the Chebychev distance metric, where the distance between two vectors is the greatest of their differences along any coordinate dimension. Using the Chebychev similarity metric as a similarity measure has the advantage of being less computational expensive in comparison to the distance-based metric of similarity since there is no need for estimating parameters of all records allocated to the cluster during negotiations. The complete distributed learning model, representing a full set of agent learning criteria for three types of similarity metrics (the Euclidian distance-based similarity, Chebychev similarity metric, and the angle metrics), consists of 24 equations in total and is omitted due to space limitations (refer to online supplementary materials [10]).

In order to be effective in solving time-sensitive data mining problems in complex uncertain environments, the developed multi-agent learning system additionally addresses the following challenges of online learning of streaming data: the processing of large number of input records and online tractability, the continual directed adaptation of the learning system parameters to environmental variations and a fast dynamic response to them in a real-time fashion, and the communication complexity of dynamic large-scale networks of autonomous agents [8].

14.3 Experimental Analysis

14.3.1 Datasets

The proposed multi-agent method of online learning was experimentally evaluated for continuous agglomerative hierarchical clustering of both synthetic datasets and datasets from the following real-world domains: the RoboCup Soccer competition and a gene expression datasets. The experimental datasets for the RoboCup Soccer domain were obtained by analyzing files of the final 2006 game (simulation league) between the teams ”Brainstormers” and ”WrightEagle”. The datasets for data mining were obtained using a data preparation framework, which parses log files of previous RoboCup Soccer games and generates knowledge representation structures of agent action scenes suitable for data mining purposes (459 instances with 10 attributes) [8]. To evaluate the performance of the developed solution to cluster datasets of high-dimensional data, we used a reduced cancer dataset [2]. The acute myeloid leukemia (AML)/acute lymphoblastic leukemia (ALL) dataset contains 192 gene and 73 patient samples.

14.3.2 Experimental Results

The reported experimental results demonstrate the applicability and efficiency of the implemented adaptive multi-agent learning system for continuous online clustering of both synthetic datasets and datasets from the following real-world domains: the RoboCup Soccer competition, and gene expression datasets from a bioinformatics test bed. The major experimental results of the conducted experimental analysis are summarized in Table 14.2 and graphically presented as fourteen charts (cf. Fig. 14.4 – Fig. 14.5 and Fig. 14.8 – Fig. 14.17 listed in Appendix), which demonstrate dynamics of the distributed learning process within and across various dimensions of the performance radar (cf. Fig. 14.5). Table 14.2 consists of four column groups, and reports solution quality for different algorithm parameters and agent decision-making strategies.

Both centralized and distributed (local) performance metrics was used to evaluate solution quality. We use the *Cophenetic (Pearson) coefficient* as a centralized performance metrics to measure quality of hierarchical clustering (computed in the shared memory primarily for comparison purposes) [18]. We also consider the agent decision-making criteria to be distributed performance metrics to evaluate the solution quality. Such ”personified” performance indicators allow for identifying quality ”bottle-necks” across all clustering hierarchy. The table reveals the dominance of different performance metrics when applying certain agent decision-making strategies, and also demonstrates which parameters of the multi-agent algorithm should be used to increase system performance along specific dimensions of the performance radar (the absolute best and worst parameter values for each performance metrics are emphasized in bold and italic types respectively).

Table 14.2: Performance comparison of different agent strategies

Learning Strategy	Euclidian Similarity	Top Candidates	Reference Point	Level Penalties	Stochastic Prematching	Proactive Roles	Continual Learning	Chebychev Similarity	Angle Metrics
Operating Time									
T, ms	3375	**2781**	3165	3450	*3553*	3484	3297	11937	2859
Cophenetic (Pearson) coefficients									
PMR	**0.9029**	0.8779	0.8866	*0.7873*	0.8939	0.9487	0.9029	0.6395	0.7227
PBC	0.9029	**0.9353**	0.9313	0.9277	*0.8939*	0.9624	0.9029	0.8835	0.7227
PBA	*0.9414*	0.9799	0.9799	**0.9965**	0.9650	0.9813	0.9414	0.9957	0.7790
PM	0.4213	0.4249	0.4261	**0.4270**	*0.4166*	0.5807	0.4213	0.3501	0.2813
Cluster Agent Values "Contains"									
CBA_CS	0.5974	0.5974	0.5974	0.5974	0.5974	0.5974	0.5974	0.9605	0.9930
CBC_CS	0.2273	0.2273	0.2273	0.2273	0.2273	0.3142	0.2273	0.8259	0.7690
CM_CS	**0.1715**	0.1621	0.1665	*0.1596*	0.1630	0.2086	0.1715	0.7086	0.7102
CMR_CS	0.0334	0.0334	0.0334	0.0334	0.0334	0.0334	0.0334	0.2489	0.3880
CWA_CS	0.0184	0.0184	0.0184	0.0184	0.0184	0.0334	0.0184	0.2489	0.0290
Cluster Agent Values "Contained"									
CBA_CD	0.1663	*0.1239*	**0.3046**	**0.3046**	**0.3046**	0.3098	0.1663	0.7383	0.8286
CBC_CD	0.0974	*0.0663*	**0.2273**	0.1149	0.1250	0.1373	0.0974	0.6851	0.6649
CM_CD	0.0618	0.0467	**0.0835**	*-0.4568*	0.0772	0.0654	0.0618	0.5479	0.5321
Record Agent Values "Contained"									
RBA	0.5974	0.5974	0.5974	0.5974	0.5974	0.5974	0.5974	0.5974	0.9930
RM	0.2273	0.2273	0.2273	0.2273	0.2273	0.2644	0.2273	0.2273	0.7690
RWA	0.0184	0.0184	0.0184	0.0184	0.0184	0.0325	0.0184	0.0184	0.0290

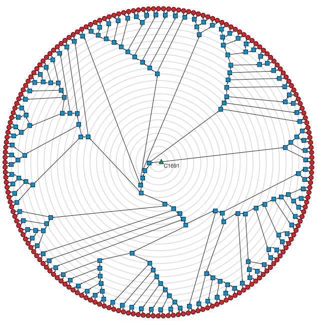

Fig. 14.4: A radial dendrogram of learning results for clustering the gene expression dataset with 192 genes and 73 patients (the equilibrium state)

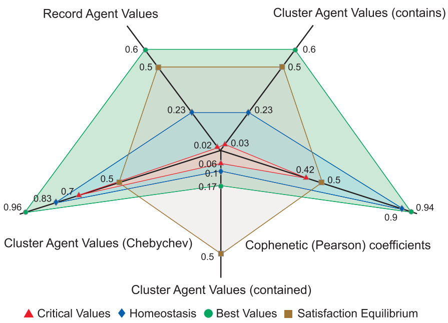

Fig. 14.5: A performance radar of major agent decision-making criteria (the multi-criteria model of quality metrics for the agent society)

The first set of experimental charts presents learning results for clustering the gene expression dataset and RoboCup Soccer dataset, and solution quality of the multi-agent algorithm across different performance metrics, agent decision-making criteria, and different parameters of the multi-agent algorithm. Learning results for clustering the gene expression dataset from the bioinformatics testbed are presented in Fig. 14.4 as a interactive radial hierarchical dendrogram, which at any point in the computation graphically represents a dynamic reconfigurable network of the semantic instances of mining agents in the virtual learning marketplace. Red circles and blue squares of the radial dendrogram represent clustered record and cluster agent respectively, and a green triangle in the center of the dendrogram constitutes the root cluster agent at the top level of the learning hierarchy. Additionally, it can be seen from the dynamic radial dendrogram that the semantic instances of mining agents are not numbered successively in a stable equilibrium state which is explained by the incremental nature of the distributed learning process. The multi-agent system responds to an external event by locally reorganizing only those areas of the global decision space that are affected by the event (*incremental optimization* in the local context). All decisions and established links are not fixed in the system and locally reconsidered when needed during reaction on environmental perturbations.

Fig. 14.5 represents the overall performance radar of five major quality metrics (agent decision-making criteria) for the agent society in the virtual learning marketplace, which define several areas of solution quality: critical and best values, homeostasis and satisfaction equilibrium. Critical values of the performance dimensions define the area of solution quality, where mining agents are not satisfied with established relationships with other agents and actively look for other allocation options to enhance their satisfaction levels with minimal costs. On the contrary, best values define the area of solution quality, where autonomous agents have absolute best satisfaction levels for each performance dimension. Satisfaction equilibrium defines the area of solution quality, where satisfaction levels of the agents are sufficiently high such that they turn to the inactive states and do not exhibit proactivity to improve their values, thereby releasing computational resources of the multi-agent system to allow other agents to be active and enhance their satisfaction values (computationally efficient algorithm implementation). Homeostasis area defines the stationary point of the multi-agent system with Pareto optimality of its quality criteria, where enhancing values of one performance metric could not be achieve without decreasing the overall solution quality. The multi-criteria quality model enables the continuous learning algorithm to operate on the basis of several performance criteria, and to be suitable for exploratory data analysis by dynamically balancing the evaluation criteria at run-time during execution of the online multi-agent system (supporting the selection and regulation of appropriate learning parameters).

Fig. 14.6 depicts the distribution of quality levels of the learning hierarchy (cuts of the agent society) across different performance metrics and over operating time of the multi-agent algorithm. The anytime multi-agent learning algorithm not only extracts an optimal multi-level partitioning of data using various measures of similarity, but also defines optimal segmentations of the input data set into clusters by dynamically selecting the quality levels (cuts) of the learning hierarchy for each performance metric defined in the multi-criteria quality model. Thus, performance comparison of Euclidian and Chebychev similarity metrics as agent learning criteria is presented in Fig. 14.13, which demonstrates that their major performance metrics have the same final configuration of quality levels, and suggests that due to described advantages of the Chebychev similarity metric it should be used where appropriate. Fig. 14.7 demonstrates the result of a conducted experiment on the analysis of the distribution of performance metrics by hierarchical levels of the agent society, and represents a stable equilibrium state of the distributed learning process (static situation). By comparing the learning results of a conventional hierarchical single-linkage clustering algorithm with performance metrics of the developed continual multi-agent algorithm, we can conclude that solution quality obtained by the multi-agent learning system in its final stable configuration is as good as one of the batch algorithm.

Fig. 14.14 illustrates the dynamic characteristics of agent instances during the continual distributed learning process and emphasizes the agent *ripple effect*, which is a decision reconsideration chain that improves the overall clustering results due to *agent proactivity*. The grey shaded square area of the chart, which is formed by the intersection of the vertical (time of the ripple effect during which agents improve their satisfaction) and horizontal (a value gained as a result of the agent proactivity) grey areas, emphasizes the property of the multi-agent algorithm to proactively trade-off solution quality and operating time. A back dotted line of the chart represents the dynamics of the Cophenetic coefficient over operating time of the multi-agent algorithm, and is displayed on its own scale to demonstrate incremental improvement of solution quality as a result of the ripple effect. A red line of the chart represents the number of non-allocated record agents, which is reduced over time and approaches a zero value in the final stable configuration of the multi-agent system. Behavior of the red line inside the shaded square area demonstrates the moment of switching agent proactivity, during which the system transit from one dynamic state of balance into a new economically more effective one by the breaking of previously established ontological relationships between record and cluster agents and establishing new semantic links. Thus, it can be seen from the diagram that the number of non-allocated record agents becomes zero in the beginning of the grey shaded area (all agents are allocated). However, during the proactivity stage the agents, seeking to increase their satisfaction values, break previous ontological relationships (temporary becoming unplanned) and lead the multi-agent system to a new configuration of the dynamic equilibrium with better solution quality (the number of non-allocated record agents becomes zero again at the end of the grey shaded area, but during the period of agent proactivity the value of the Cophenetic coefficient is increased as a result of the agent ripple effect).

Fig. 14.9 demonstrates the results of the conducted experiments on introducing the proactive agent ”Record-Cluster” role into the agent decision-making model. The chart reveals a significant increase of the solution quality across all performance metrics (though taking slightly more time to settle down to a new quasi-optimal state) when the record agents not only exhibit reactive behavior by simply responding to system events, but also proactively search for profitable allocation variants. Fig. 14.11 demonstrate that cluster centroids, produced by the developed multi-agent algorithm, can be considered as *significant representative features* (”reference points”) for the algorithm since the latter provides representative clustering results with the satisfactory overall quality of partitioning. Fig. 14.10 demonstrates the ability of the multi-agent learning algorithm to avoid local optima and enhance the overall learning results by means of the incremental stochastic agent pre-matching algorithm (”random sampling” of the search space), which restrictively regulates the agent pre-matching radius to incrementally increase a depth of agent vision and stochastically select candidate agents within the agent pre-matching radius. Fig. 14.15 represents performance metrics of cluster agents evolving in the virtual learning marketplace (with quality levels shown as numbers above the lines, and final values of the performance metrics in the stable equilibrium state of the multi-agent system shown on the right side side of the chart).

The second set of experimental charts demonstrates the performance of the developed multi-agent algorithm to conduct continual online learning of stream data. Fig. 14.16 demonstrates the functional dependence of online learning performance on the number of records being continuously clustered. It can be seen from the diagram that in situations in which a new record arrived when previous records have been allocated, the time required to incrementally incorporate a new record into the hierarchical learning structure is not exponential (and approximately linear) due to the implemented agent memory and directed search mechanisms. Each allocation moment, incremental planning a new record takes different amount of time, since mining agents do not consider all allocation options at once, but only those available in their local context, and subsequently increase their field of vision to incrementally improve the initial (previous) solutions. Thus, curve dips (troughs) of the graph demonstrate the situations where new allocation happens with minimum re-learning of previously established relationships (the ripple effect of minimal length), and picks of the graph indicate the moments when planning a new agent leads to considerable reallocation of previously formed agent links. A red line of the chart represents the accumulative number of reallocation of previously planned agents for each stable equilibrium state, which is increased over time.

Fig. 14.17 presents the results of the conducted comparison analysis of the developed continuous multi-agent method and a conventional hierarchical single-linkage clustering algorithm (the Alias ”LingPipe” software library). The chart demonstrates that for the incremental approach the time required to react on changes and maintain the clustering hierarchy valid reduces as the number of agents affected by environmental variations decreases (the length of the ripple effect is context dependent), while re-learning time remains approximately constant for the batch algorithm. Thus, the conducted comparison analysis demonstrates the strong advantage of the developed incremental multi-agent approach over the classical batch clustering algorithm in dynamic settings as a result of the developed matching memory mechanism ensures a quick response time to changes by directly adapting of only those areas of the global decision space that are affected by them. Nevertheless, it should be noted that due to its distributed nature the continuous multi-agent algorithm takes more time for batch data mining in static settings than the centralized algorithm (the blue line is higher than the read one in the beginning of the chart). Fig. 14.8 provides a performance comparison chart for clustering in a dataset in batch and continuous learning modes. This chart and Fig. 14.12 demonstrate that the algorithm ensures deterministic learning results for various input sequences of arriving records at the system (the results converge to the same values in the final stable equilibrium state of the distributed learning process).

14.4 Summary, Conclusion and Future Work

14.4.1 Summary

This work considers the problem of continuous data mining of streaming data in complex dynamic environments under conditions of uncertainty. The primary goal of the developed multi-agent approach is continuous learning in distributed environments from decentralized data sources across a heterogeneous data environment with a view to effectively responding to environmental dynamics and performing online data analysis using various dynamic learning criteria and measures of similarity. The main contributions of this research can be summarized as follows.

1. With a view to responding to rapid changes in the environment, we developed a multi-agent method of continuous online learning, by modeling the task of unsupervised clustering as a *dynamic distributed resource allocation problem*. A game-theoretic decentralized market-based method of competitive and implicit multi-agent negotiation, and an asynchronous message-passing algorithm are developed to obtain an implicit global quasi-optimal solution to the distributed constrained learning problem, which requires market-based negotiation between different self-interested agents, defined for data elements, to satisfy a dynamic distributed constraint network of agent interests by maintaining a dynamic balance among the interests of all participants in the interaction.
2. A knowledge-based competitive *multi-agent learning system* is developed to enable the data-driven self-organizing distributed process of dynamic continuous data mining. The implemented multi-agent platform of the system provides a distributed computational environment (virtual learning marketplace) and a runtime support for asynchronous quasi-parallel negotiations between the agents in a virtual marketplace. To consider personal preferences and expert knowledge, the developed virtual marketplace ontology (a semantic knowledge base located in the shared memory) has been set up to contain conceptual knowledge of the problem domain and to support the dynamic regulation of various control parameters of the learning system at run-time during its execution, such as different agent decision-making criteria, active agent negotiation roles, nature of the operating environment and various properties of the learning algorithm.
3. A multi-objective *agent decision-making model* is developed to support goal-driven behavior of autonomous agents in the virtual learning marketplace, which defines for each agent in the virtual marketplace ontology its individual goals, criteria, preference functions, and decision-making strategies. The multi-objective decision-making model enables the learning algorithm to continually operate on the basis of non-standard optimization criteria, and to be suitable for online ex-

ploratory data analysis in complex uncertain environments using various measures of similarity for situations with incomplete knowledge about parameters of the learning problem.

4. The proposed multi-agent method of online learning was experimentally evaluated for continuous agglomerative hierarchical clustering of both synthetic datasets and datasets from the following real-world domains: the RoboCup Soccer competition and a gene expression datasets from a bioinformatics test bed. The conducted comparison analysis demonstrates the superior advantage of the incremental multi-agent learning approach over conventional batch clustering algorithms.

14.4.2 Conclusions and Future Directions

We support our conclusions by conducting the experimental analysis, which demonstrate the applicability and efficiency of the developed continuous multi-agent learning system to respond to environmental dynamics and to perform online data analysis using various dynamic learning criteria and measures of similarity. The reported experimental results demonstrate the strong performance of the developed multi-agent learning system for continuous agglomerative hierarchical clustering of both synthetic datasets and datasets from the real-world domains.

According to the conducted experimental analysis, we conclude that the proposed online multi-agent approach has the advantage over conventional unsupervised learning methods of being *dynamic*, *incremental* and *continuous*. The input dataset, decision-making criteria and various control parameters of the learning system are not fixed and can be changed at run-time during execution of the dynamic algorithm. Clustering results of the adaptive learning algorithm are available at any time, and continuously and incrementally improved to achieve a global quasi-optimal solution to the optimization problem, trading-off operating time and result quality. All decentralized decisions of the method are not fixed and locally reconsidered when needed in response to environmental events, taking advantage of domain semantics.

The implemented multi-objective decision-making model enables the learning algorithm to continually operate on the basis of non-standard optimization criteria and suitable for online exploratory data analysis using various measures of similarity. Additionally, the implemented adaptive agent pre-matching mechanism of regulating a depth of agent vision makes it possible for an agent to restrictively consider only those allocation options that are inside its limited field of vision, thereby preventing enumeration of all allocation possibilities and making learning of a large amount of data computationally tractable.

Nonetheless, the current implementation of the multi-agent learning system has the following limitations. The conducted comparison analysis of the developed multi-agent method and a conventional hierarchical single-linkage clustering algorithm (the Alias "LingPipe" software library) demonstrates the superior performance of the incremental approach in comparison to the batch algorithm. The developed matching memory mechanism on the agent protocol level ensures a quick response time to changes by directly adapting of only those areas of the global decision space that are affected by them. However, keeping such records of agent statuses in the shared memory of the system in order to decrease communication costs introduces the increase in the memory complexity of the algorithm. To reduce memory consumption, a special service role of the World agent periodically validates agent matching memories to delete non-valid history records of instances with dead agents.

The conducted experiments demonstrate that introducing the proactive agent ”Record-Cluster” role significantly increases the solution quality across all performance metrics. The effective regulation of agent activities in the distributed environment is crucial to obtain a global quasi-optimal solution to the problem of better quality and to enhance the performance of the distributed learning algorithm. The developed agent activity control framework provides the decentralized mechanism of regulating proactivity for certain types of agent negotiation roles based on the observation of the hormonal level in the multi-agent environment. Nevertheless, balancing proactivity of agent negotiation roles in the virtual learning marketplace is a challenging issue. Thus, some performance charts of the conducted experimental analysis reveal a slight decrease in solution quality during the moments of switching agent proactivity for certain types of agent negotiation roles. Although the strong advantage of the developed incremental multi-agent approach over classical batch clustering algorithms in dynamic settings was demonstrated, the conducted experimental analysis revealed the current limitation of the developed learning algorithm to efficiently perform massive data processing of high-dimensional data (e.g. genome-wide gene expression data) in centralized and batch settings due to communication costs of additional message-passing and decision synchronization algorithmic steps. Nevertheless, to increase the performance of the developed approach for batch data processing in static settings, we suggest conducting combinatorial auctions in the local agent context and propose a hybrid online learning approach that combines distributed constraint optimization techniques and decision theoretic approaches [9].

We consider the developed solution to be an important step in our research towards the development of effective online learning algorithms in dynamic uncertain environments, and plan to explore how to learn non-stationary dynamics. Future work will be directed towards addressing the following challenging problems that have arisen from this research: (1) extending the adaptive learning approach to support online automatic *semi-supervised classification* by continuously deducing semantic-based classification rules from clustering results and performing automatic rule-based classification and subsequent pattern verification at run-time, and (2) developing a *hybrid online learning approach*, which reduces the problem dimension while maintaining the essential characteristics of the original system, and provides a dynamic bidirectional cyclic feedback on using the market-based distributed resource allocation (bottom-up) and Bayesian reinforcement learning of decentralized partially observable Markov decision processes (Dec-POMDPs) (top-down).

Appendix

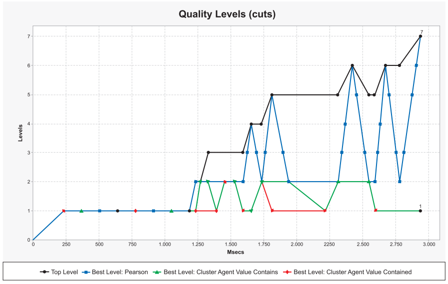

Fig. 14.6: Quality levels (cuts) of the learning hierarchy across different performance metrics

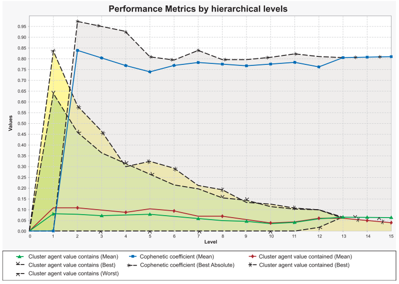

Fig. 14.7: Distribution of performance metrics by hierarchical levels (stable equilibrium state)

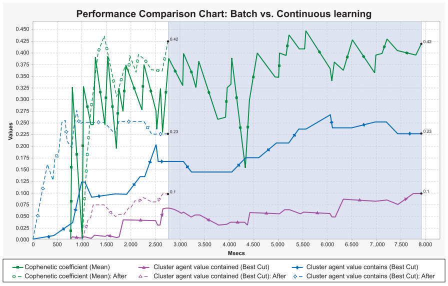

Fig. 14.8: Time required for Batch vs. Continuous learning (performance comparison)

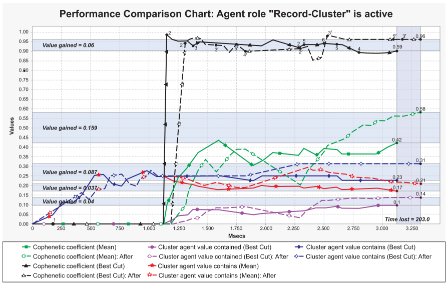

Fig. 14.9: Trading-off operating time & learning quality(agent role ”Record-Cluster” is proactive)

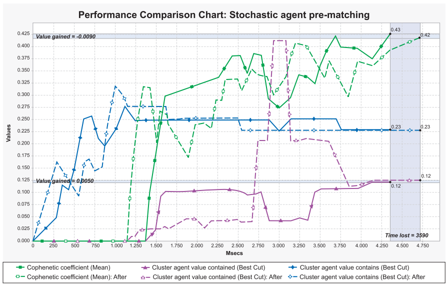

Fig. 14.10: Time required for learning with incremental and stochastic agent prematching

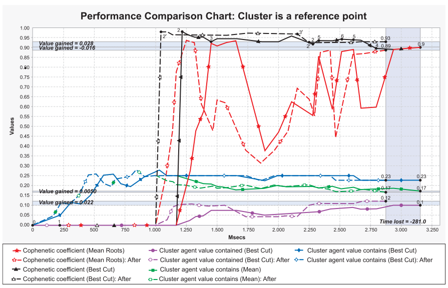

Fig. 14.11: Performance comparison chart: cluster centroids as representative ”reference points”

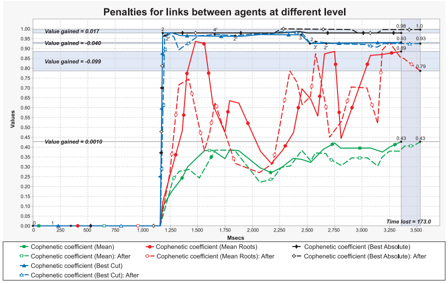

Fig. 14.12: Time required for learning with penalties, established for links between agents at different levels of the learning hierarchy

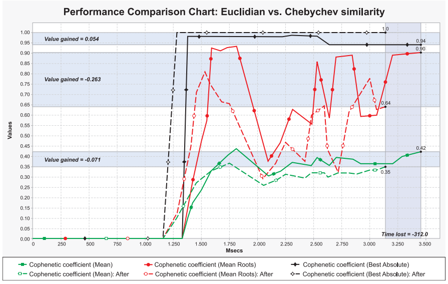

Fig. 14.13: Performance comparison chart: Euclidian vs. Chebychev similarity metrics as agent learning criteria (with quality levels shown)

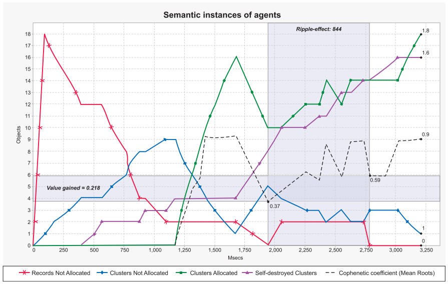

Fig. 14.14: The agent dynamics of the distributed learning process (with incremental improvement of solution quality as a result of the ripple effect)

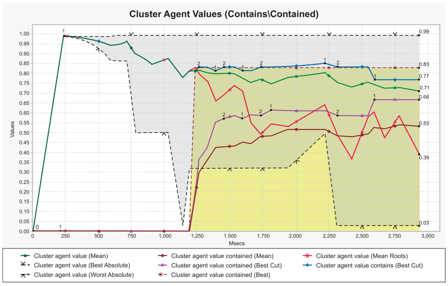

Fig. 14.15: Performance metrics of cluster agents evolving in the virtual learning marketplace (with quality levels and equilibrium metric values)

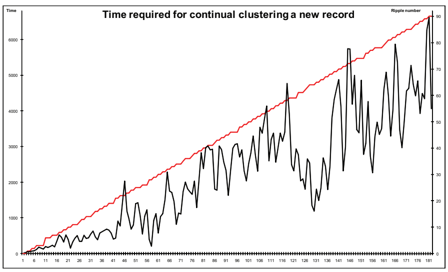

Fig. 14.16: Time required for continual learning a new record when previous records allocated (with the ripple effect for each equilibrium state)

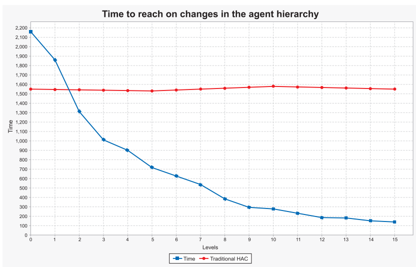

Fig. 14.17: Time required to respond to changes in clustering hierarchy (comparison with the single-linkage clustering Alias ”LingPipe” algorithm)

References

1. Aggarwal, C.C. (ed.): Data Streams: Models and Algorithms. Advances in Database Systems. Springer, New York, NY, USA (2007)
2. Al-Shalalfa, M., Alhajj, R.: Application of double clustering to gene expression data for class prediction. In: AINA Workshops (1), pp. 733–738. IEEE Computer Society (2007)
3. Auroop R. Ganguly Joao Gama, O.A.O.M.M.G.R.R.V. (ed.): Knowledge Discovery from Sensor Data. CRC Press Inc - Taylor and Francis Ltd, New York, NY, USA (2008)
4. Bagherjeiran, A., Eick, C.F., Chen, C.S., Vilalta, R.: Adaptive clustering: Obtaining better clusters using feedback and past experience. IEEE International Conference on Data Mining **0**, 565–568 (2005)
5. Davidson, I., Ravi, S.S.: Agglomerative hierarchical clustering with constraints: Theoretical and empirical results. In: PKDD-05, *LNCS*, vol. 3721. Springer (2005)
6. Frey, B.J., Dueck, D.: Clustering by passing messages between data points. Science **315**, 972–976 (2007)
7. Jain, A.K., Murty, M.N., Flynn, P.J.: Data clustering: A review. ACM Comput. Surv. **31**(3), 264–323 (1999)
8. Kiselev, I., Alhajj, R.: An adaptive multi-agent system for continuous learning of streaming data. IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT'08) **2**, 148–153 (2008)
9. Kiselev, I., Alhajj, R.: Online dynamic optimization under conditions of uncertainty. In: A.P.S.D.R. Nicholas R. Jennings Alex Rogers (ed.) 1th International Workshop "Optimisation in Multi-Agent Systems" (OptMAS), pp. 52–59. AAMAS-08 (2008)
10. Kiselev, I., Alhajj, R.: Supplementary materials, software demonstration, and video recordings of the developed multi-agent learning system, described in this work. Website (2008). `\footnotesize{\ttfamily{http://www.multiagent.org/mining}}`
11. Kiselev, I., Glaschenko, A., Chevelev, A., Skobelev, P.: Towards an adaptive approach for distributed resource allocation in a multi-agent system for solving dynamic vehicle routing problems. In: AAAI-07, pp. 1874–1875 (2007)
12. Klusch, M., Lodi, S., Moro, G.: Issues of agent-based distributed data mining. In: AAMAS-03, pp. 1034–1035 (2003)
13. Likas, A.: A reinforcement learning approach to online clustering. Neural Computation **11**(8), 1915–1932 (1999)
14. MacKie-Mason, J.K., Wellman, M.P.: Handbook of Computational Economics, vol. Volume 2, chap. Chapter 28 Automated Markets and Trading Agents, pp. 1381–1431. Elsevier (2006)
15. Modi, P.J., Jung, H., Tambe, M., Shen, W.M., Kulkarni, S.: Dynamic distributed resource allocation: A distributed constraint satisfaction approach. In: ATAL: Revised Papers from the 8th Intern. Workshop on Intelligent Agents VIII, pp. 264–276. Springer-Verlag, UK (2002)
16. Modi, P.J., Shen, W.M., Tambe, M., Yokoo, M.: Adopt: asynchronous distributed constraint optimization with quality guarantees. Artif. Intell. **161**(1-2), 149–180 (2005)
17. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.): Algorithmic Game Theory. Cambridge University Press, New York, NY, USA (2007)
18. Rodrigues, P.P., Gama, J., Pedroso, J.P.: Hierarchical clustering of time-series data streams. IEEE Trans. Knowl. Data Eng. **20**(5), 615–627 (2008)
19. Russell, S.J., Norvig, P.: Artificial Intelligence: A Modern Approach, 2nd international ed. edn. Prentice-Hall, Upper Saddle River, NJ, USA (2003)
20. Smith, R.G.: The contract net protocol: High-level communication and control in a distributed problem solver. IEEE Trans. Computers **29**(12), 1104–1113 (1980)
21. Symeonidis, A.L., Mitkas, P.A.: Agent Intelligence Through Data Mining. Multiagent Systems, Artificial Societies, and Simulated Organizations. Springer-Verlag New York (2005)
22. Theocharopoulou, C., Partsakoulakis, I., Vouros, G.A., Stergiou, K.: Overlay networks for task allocation and coordination in dynamic large-scale networks of cooperative agents. In: E.H. Durfee, M. Yokoo, M.N. Huhns, O. Shehory (eds.) AAMAS, p. 55. IFAAMAS
23. Zhang, S., Zhang, C., Wu, X.: Knowledge Discovery in Multiple Databases. Springer-Verlag (2004)