On Complex Event Processing for Real-time Situational Awareness

Nenad Stojanovic[1], Alexander Artikis[2]

[1] FZI Karlsruhe, 76131 Karlsruhe, Germany
Nenad.Stojanovic@fzi.de
[2] NCSR Demokritos, Athens 15310, Greece
a.artikis@iit.demokritos.gr

Abstract. In this paper we give an overview of the existing research results and open research challenges in applying complex event processing for real-time situational awareness. We consider two different viewpoints: better detection of emerging complex situations and prediction of future situations. In order to illustrate these viewpoints we consider two application areas: activity recognition from the video content and social media observation, respectively.

Keywords: Situation awareness, Social media observation, Activity recognition

1 Introduction

Real-time data processing has become very important for many applications, such as item-tracking in RFID-supported logistics, social-media channels observation, activity recognition from video content, computer network monitoring, patient monitoring and trader behavior evaluation in financial markets. In all of these applications the amount of data being generated requires on–the-fly processing and immediate reaction in order to be managed in an efficient way. Indeed, such real-time orientation enables the detection of problems (e.g. a damaged item in a delivery, or bad image of a company in recently posted tweets) as soon as they happen, so that a corresponding reaction can be successfully performed.

In the nutshell of this mechanism is the ability to recognize in real-time[1] (or even ahead of time) some interesting situations, what is called “real-time situational awareness”. Note that this goes beyond the traditional (static) situational awareness that is focused on the understanding a situation (if possible in real-time). Real-time situational awareness introduces the notion of real-time emergency: the main goal is to recognize a situation of interest as soon as possible in order to be able to react to it properly.

Such a process introduces several challenges for the processing of data:

1. it should be very efficient in order to deal with a huge amount of events,

[1] We consider „business real-time“ as the criteria for declaring something to be processed in real-time

2. it should be intelligent in order to enable an early recognition of interesting situations,
3. it should be very flexible in order to deal with various and dynamically changing patterns (situations) of interests (that should be recognized in real-time), and
4. it should be tolerant to various types of noise.

Due to its real-time processing orientation, complex event processing (CEP) is a technology that aims to resolve these challenges. Indeed, CEP is in the nutshell about efficient management of the pattern detection process in the huge and dynamic data streams and as such it is very suitable for detecting/recognizing complex real-time situations. However, CEP can be also viewed as a technology for detecting trends in the flow of data that leads to predicting some future situations. In this paper we give an overview of the application of complex event processing for the real-time situational awareness from these two different viewpoints:

- better detection of emerging complex situations (complex real-time-situation awareness) and
- prediction of future situations (future-situation awareness).

In order to illustrate these viewpoints we consider two application areas: a) activity recognition from the video content, and b) social media observation, respectively.

Indeed, in the area of activity recognition one has to deal with several of the problems faced by traditional event processing systems such as the real-time recognition of long-term activities given large amounts of detected short-term activities the continuous, automated refinement of long-term activity representation for increased recognition accuracy, as well as the accurate recognition in the presence of noise.

Social media observation is dealing with detecting interesting information in the streams of data coming from various social media channels. These channels have become very important sources for monitoring public opinion in order to, e.g. (re)define marketing strategies. Recently the term real-time marketing intelligence has been coined for describing the need for reacting immediately on the threats/opportunities for having more and more satisfied customers.

Moreover, this paper discusses about the application of two machine learning approaches in CEP:

- using unsupervised learning for predicting novel situations and
- using supervised learning for automatically refining complex event representations and handling noise.

2 Activity Recognition

2.1 Motivation

One of the objectives of computer vision research is to recognize in an automated way what happens in scenes depicted by video sequences. An automatic video interpretation system takes video sequences as input and produces as output the interpretation of these sequences. Such a system is typically composed of two main components. The first component produces a symbolic representation of the mobile objects detected on raw video sequences, as well as of the 'short-term activities' of the mobile objects – simple events, in the terminology of event processing – such as when a person is walking, running, stays inactive, moves abruptly, and so on. The second component recognizes various types of 'long-term activity' – complex events – such as when a person leaves an object unattended, two people are fighting, etc. A long-term activity is defined as a set of temporal, spatial and/or logical constraints on a set of short-term activities.

Automatic video interpretation systems have to deal with several of the problems faced by traditional event processing systems:

- They have to recognise in real-time long-term activities given thousands of short-term activities per second.
- They are required to recognize long-term activities that are possibly about to take place, such as when a person is about to leave an object unattended.
- They need machine learning techniques to automatically and continuously refine what constitutes a long-term activity. Such a refinement increases recognition accuracy.
- They have to deal with uncertainty, as the detection of short-term activities from video content is often noisy.

In following we discuss how some of the above issues are being addressed in the field of activity recognition, as we believe that several of the methods of this field are directly applicable to complex event processing.

2.2 A Representative Approach

Several approaches have been proposed in the literature for long-term activity recognition – [3, 5, 7, 7, 8] are but a few recent examples. To address the issue of uncertainty mentioned above, various probabilistic reasoning frameworks have been adopted. One such framework, that is increasingly gaining attention, is Markov Logic Networks (MLN) [4]. MLN combine first-order logic and probabilistic graphical models. The use of first-order logic allows for the representation of long-term activities including complex temporal constraints – this is opposed to sequential graphical models that allow for restricted temporal representation. The main concept behind MLN is that the probability of a world increases as the number of formulas it violates decreases. Therefore, a world violating formulas becomes less probable, but not impossible as in first-order logic. Syntactically, each formula F_i in Markov logic is represented in first-order logic and it is associated with a weight w_i. The higher the value of the weight, the stronger the constraint represented by F_i. Semantically, a set of Markov logic formulas (F_i, w_i) represents a probability distribution over possible worlds.

MLN have been recently used for long-term activity recognition – consider, for example, [3, 8]. (A detailed account of the use of MLN for activity recognition and complex event processing in general may be found in [1].) The MLN inference algorithms take into consideration the weights attached to the short-term activities detected by the underlying computer vision component. As mentioned above, it is often the case that short-term activity detection is noisy and, therefore, this feature of MLN is very helpful for long-term activity recognition. Furthermore, using MLN one may attach weights to the first-order logic rules expressing a long-term activity. Strong weights are given to rules that are almost always true, while weak weights are given to rules that describe exceptions.

An important feature of MLN is that they are supported by a variety of machine learning algorithms. More precisely, it is possible to estimate the weights of the first-order rules expressing a long-term activity and/or the rules themselves, given a set of training data, that is, short-term activities annotated with long-term activities. Weight learning in MLN is performed by optimising a likelihood function, which is a statistical measure of how well the probabilistic model (MLN) fits the training data. Weights can be learned by either generative or discriminative estimation. Generative learning attempts to optimise the joint distribution of all variables in the model. In contrast, discriminative learning attempts to optimise the conditional distribution of a set of outputs, given a set of inputs.

In addition to weight learning, the structure of a MLN, that is, the rules expressing long-term activities, can be learned from training data. A variety of structure learning methods have been proposed for MLN. These methods build upon the techniques of Inductive Logic Programming (ILP) [6]. In brief, the MLN structure learning methods can be classified into top-down and bottom-up methods [4]. Top-down structure learning starts from an empty or existing MLN and iteratively constructs clauses by adding or revising a single predicate at a time, using typical ILP operations and a search procedure. However, as the structure of a MLN may involve complex long-term activities, the space of potential top-down refinements may become intractable. For this reason, bottom-up structure learning can be used instead, starting from training data and searching for more general hypotheses. This approach usually leads to a more specialised model, following a search through a manageable set of generalisations.

Long-term activity recognition in MLN involves querying a ground Markov network about long-term activities. A ground Markov network is produced by grounding all first-order logic rules expressing long-term activities using a finite set of constants concerning the application under consideration. Complete grounding of MLN, even for simple long-term activity knowledge bases, results in complex and large networks, compromising inference efficiency. For this reason, we may employ 'lazy inference methods', or 'lifted inference methods' which can answer queries without grounding the entire network [4].

2.3 Research Challenges

Markov Logic Networks (MLN) combine the strengths of logical and probabilistic inference. Consequently, they address, to a certain extent, the following issues that often arise in activity recognition and complex event processing in general: incomplete simple event (short-term activity) streams, erroneous simple event detection, and inconsistent annotation of simple events and complex events (long-term activities). MLN also offer a very expressive framework, as the full power of first-order logic is available.

Being based on logic, MLN benefit from a formal and declarative semantics, a variety of inference mechanisms, and methods for learning a knowledge base of complex events from data. Compared to procedural methods, logic-based ones facilitate efficient development and management of complex event representations, which are clearly separated from the generic inference mechanism. Moreover, compared to methods exhibiting informal semantics, logic-based approaches support validation and traceability of results.

Although MLN are increasingly being used for activity recognition and, in general, complex event processing, there are several issues that need to be resolved still, such as the incorporation and use of numerical temporal constraints in MLN inference, and the simultaneous learning of numerical parameters — for example, weights and numerical temporal constraints — and the logical structure of the knowledge base expressing complex events.

3 Future-situation awareness in social media

3.1 Motivation

The power of data processing has migrated in recent years from explaining the past using very efficient (web) search mechanisms (answering on the question "What happened?") into understanding the present through different social media-based filtering mechanisms (answering on the question "What is happing now?"). Indeed, in last two years there is a huge proliferation of systems for the real-time analyses of the streams from news portals, discussion forums, Twitter, etc. leading to different types of services for filtering interesting information, e.g. automatic notification in the case that interesting information has appeared. This feature enables developing of very efficient reactive systems that will react as soon as a problem/opportunity appears. It boosts the competitiveness through the real-time awareness that increases the speed of reaction. In this case the data is used for an efficient detection of some already known and already-happened situations. However, an extreme dynamicity and complexity of the modern business implies the need for creating awareness of that what can happen in order to be able not only to react on problems, but to avoid that they will happen at all (the task of the so-called proactive systems).

On the other hand, there are new, community-driven, sources of information and influences that can be considered and exploited for estimating what can happen, like the wisdom of crowds, the power of an individual blogger to impact a company's image, the fact that consumers now buy as communities, and the potential radical restructuring of classic functions like product design because consumers are designing their own products. In that case the data should be used for an efficient prediction that some situations might happen, including the detection of prior unknown situations (answering on the question "*What will happen?*"). This data-driven proactivity opens possibilities to provide new solutions for new situations and from the business point of view it will boost *competitiveness trough innovativeness*. Many added value services based on this proactivity can be envisioned, starting from calculating specific, context-based predictions that can point to new business threats or opportunities, till providing efficient actions based on predictions. Just as short example imagine a SME that can produce small series of specialized chips and can discover (using specific predictions calculated from data streams) that in a segment of the market there would be, in a shorter period of time, either problems in the supply or new demand for specific types of chips (small series).

Obviously, this posts new requirements for processing streams of data[2] in the future, which the following two are the most important of: a) generating streams of interesting information out of streams of data (interestingness is defined through the statistics-based notion of unusuality) and b) generating different trend-based analytics (incl. relevant predictions) out of streams of interesting information. We argue that these two processing steps define a basic pipeline for the so-called *proactive processing of streams* that will enable the creation of many business models around various derivatives of streams of data.

3.2 State of the art/current situation

The need for predicting some parameters of the business is well recognized and mainly supported by predictive analytics, which encompasses a variety of techniques from statistics, data mining and game theory that analyze current and historical facts to make predictions about future events. However, predictive analytics is usually used for defining some operative data (like what revenue can be expected and from what channels - direct sales, partners, etc.) and it is based on developing very formal mathematical models (e.g. a customer profile can include more than 200 descriptive variables, such as income, age group, gender, occupation, and education level). As such it doesn't exploit the potential of data streams for listening to the rumors and finding weak signals in the marketplace that can be used for the competitive advantage.

On the other hand, many portals have been developed in last two years around available real-time data streams (like twitter) with the goal of using them for building different (real-time) opinion maps, usually as a part of analyzing effects of some marketing activities. We have analyzed more than 230 solutions (mainly from USA and Germany) in the area of social media monitoring. Social media monitoring tools are used to observe user-generated-content in order to gain marketing relevant insight. Radia6 (acquired by Salesforce recently), Sysomos, Lithum, Attensity360 are the most popular solutions in this area. Almost more than the half of them analyzes content only from blogs, Twitter and Facebook. The remaining systems are built either for Twitter or Facebook. Usage of other available sources, like crawling discussion forums and web pages is missing in these systems. None of the systems consider mining techniques in order to gain intelligent insights. Prediction based on collected data is also not in the focus of these systems. There is no existing solution

---

[2] Note that we consider as streams of data all data which is coming with a time stamp from a data source, independently of the frequency of changes. For example, publishing changes/updates in a web page represents such a stream of data. Other examples are updates from discussion forums, twitter and facebook (wall). Note that we are talking mainly about community-driven data.

which deals with an advanced analysis of the data (like providing unusual information) or working towards predicting some trends for the future. To the best of our knowledge there is only one approach which deals scientifically with the predictions based on the data from social media [10]. This approach uses the chatter from Twitter.com to forecast box-office revenues for movies and outperform other prediction methods. However, this approach, like above mentioned predictive analytics, deals with predicting some operative data and the goal approach aims to discover/predict completely new/unexpected business threats and opportunities, that couldn't be detected otherwise.

3.3 Research challenges and Beyond State of the Art

New research work should be mainly focused on developing efficient and scalable solutions for finding unusuality (changes) in the data flow: the value of data streams is determined not only by the data itself (i.e. data value) but rather by the context the data values have appeared in. For example, if there are two complaints coming from the gold-customer about X in a short period of time, one will react on that differently depending on that if that customer complains very rarely or quite often including multiple complaints. This will enable finding interesting information in data streams (see previous section). This will require methods for the detection of unusual situations in textual data streams.

Discovering unusual situations have been studied widely through many different fields of science and economy. In [11] the authors isolate outliers to assure high quality of stored data. In [12] and [13] a feature-based clustering algorithm is presented. They mark that previous approaches considers outliers as binary properties and introduced LOF (local outlier actor) to measure the degree of an outlier.

From the data mining point of view, very important is the work of [14], who provides a framework for effective examining of multidimensional data streams with the use of "velocity density estimation". It allows both a fast computation and to derive a visual perspective of the changes. The idea behind it is to estimate the changing rate of the data based on some user-defined temporal window.

As a short conclusion, most approaches in this field are using unsupervised learning techniques to detect unusualness. The large number of such approaches is due to the possibility for fast computation. However, they are only adjusted to some specific domains. In unpredictable situations it is not desirable to readjust the underlying parameter. Further, none of the presented approaches, aside Data Stream Mining solutions, has real-time capabilities.

4 Conclusions

Although CEP is in the nutshell about efficient management of the pattern detection process in the huge and dynamic data streams, it can be also viewed as a technology for detecting trends in the flow of data that leads to predicting some future situations. In this paper we give an overview of the application of complex event processing for the real-time situational awareness from these two different viewpoints: better detection of emerging complex situations (complex real-time-situation awareness) and prediction of future situations (future-situation awareness). In order to illustrate these viewpoints we consider two application areas: a) activity recognition from the video content, and b) social media observation, respectively. We presented the most important research challenges for these areas and gave guidelines for the further research in these directions.

Acknowledgments

Alexander Artikis is supported by the EU-funded PRONTO project (FP7-ICT 231738). Nenad Stojanovic is partially financed by EU in the following FP7 projects: ALERT (ICT-258098) and PLAY (ICT-258659)

References

1. K. Thirunarayan, C. Henson, and A. Sheth, Situation Awareness via Abductive Reasoning for Semantic Sensor Data: A Preliminary Report, in Proc. 2009 Int. Symposium on Collaborative Technologies and Systems (CTS 2009), Baltimore, MD, May 18-22, 2009.
2. Artikis A., Porter F., Skarlatidis A. and Paliouras G. Logic-based representation, reasoning and machine learning for event recognition. International Conference on Distributed Event-Based Systems (DEBS), 282-293, ACM, 2010.
3. Biswas, R., Thrun, S., and Fujimura, K. 2007. Recognizing activities with multiple cues. In Workshop on Human Motion. LNCS 4814. Springer, 255–270.
4. Domingos, P. and Lowd, D. 2009. Markov Logic: An Interface Layer for Artificial Intelligence. Morgan & Claypool Publishers.
5. Hakeem, A. and Shah, M. 2007. Learning, detection and representation of multi-agent events in videos. Artificial Intelligence 171, 8–9, 586–605.
6. Muggleton, S. 1991. Inductive logic programming. New Generation Computing 8, 4, 295–318.
7. Shet, V., Neumann, J., Ramesh, V., and Davis, L. 2007. Bilattice-based logical reasoning for human detection. In Proceedings of International Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, 1–8.
8. Thonnat, M. 2008. Semantic activity recognition. In Proceedings of European Conference on Artificial Intelligence (ECAI). 3–7.
9. Tran, S. D. and Davis, L. S. 2008. Event modeling and recognition using markov logic networks. In Proceedings of Computer Vision Conference. 610–623.
10. Asur S., Huberman, B. A. "Predicting the Future with Social Media," wi-iat, vol. 1, pp.492-499, 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, 2010
11. Last M, Kandel A. Automated Detection of Outliers in Real-World Data. In Proc. of the Second International Conference on Intelligent Technologies, 2001, S. 292-301.
12. Minh L. M., Nguyen Q., Omiecinski, E.. A Fast Feature-Based Method to Detect Unusual Patterns in Multidimensional Datasets. Springer. 2009.
13. Breunig M. M., Kriegel H.-P., Ng R. T., Sander J. LOF: identifying density-based local outliers. SIGMOD Rec. 29(2), 2000, S. 93-104.
14. Aggarwal C. A Framework for Diagnosing Changes in Evolving Data Streams. In ACM SIGMOD Conference, 2003, S. 575-586.