Expert Systems with Applications 40 (2013) 399–407

Contents lists available at SciVerse ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

A general framework for time-aware decision support systems

V. Milea [a,*], F. Frasincar [a], U. Kaymak [b]

[a] Econometric Institute, Erasmus School of Economics, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands
[b] Information Systems Department, School of Industrial Engineering, Eindhoven University of Technology, The Netherlands

ARTICLE INFO

Keywords:
Automated trading
Market recommendations
tOWL
Decision support system
Temporal knowledge
Temporal reasoning

ABSTRACT

In this paper we present a general framework for time-aware decision support systems. The framework uses the state-of-the-art tOWL language for the representation of temporal knowledge and enables temporal reasoning over the information that is represented in a knowledge base. Our approach uses state-of-the-art Semantic Web technology for handling temporal data. Through such an approach, the designer of a system can focus on the application intelligence rather than enforcing/checking data related restrictions manually. Also, there is an increased support for reuse of temporal reasoning tools across applications. We illustrate the applicability of our framework by building a market recommendations aggregation system. This system automatically collects market recommendations from online sources and, based on the past performance of the analysts that issued a recommendation, generates an aggregated recommendation in the form of a buy, hold, or sell advice. We illustrate the flexibility of our proposed system by implementing multiple methods for the aggregation of market recommendations.

1. Introduction

Decision systems often rely on historical information for the formulation of a best course of action. Storing, retrieving and checking the large volumes of data and information for consistency represents one of the main challenges in building decision support systems that use historical data. Although flat representations of data used in, for example, business intelligence, provide intelligent storage and retrieval of data (Batini, Ceri, & Navathe, 1991), automated inference, as needed in consistency checks, is limited in approaches based on such formalisms due to the rather inexpressive semantics of the underlying structures. Modern knowledge representation approaches provide for more finely grained semantics and additional expressiveness from a semantic perspective. From these approaches, Semantic Web (Lee et al., 2001; Shadbolt et al., 2006) languages such as Resource Description Framework (RDF) and RDF Schema (RDFS) (Brickley & Guha, 2004; Klyne & Carroll, 2004), and the Web Ontology Language (OWL) (Motik et al., 2009; Patel-Schneider, Hayes, & Horrocks, 2004) provide the most expressive choices when the problem of automated inference is considered. When historical data is used, some time-enabled formalism is also required. In this context, an approach such as the tOWL language (Milea, Frasincar, Kaymak, 2012; Milea, Frasincar, Kaymak, Houben, 2012), a temporal Web Ontology Language based on OWL, provides an appropriate formalism for the representation of time-varying knowledge.

These properties of the Semantic Web languages make them suitable for use in decision support systems, where some of the intelligence of the system is already incorporated in the representation language and does not need to be explicitly accounted for in the main system. For example, in a system where future product prices are predicted based on past prices, the restriction that a product may only have one price at any point in time can be enforced generically (through a temporal cardinality restriction), at the level of the whole knowledge base, and outside the main application. This can be achieved by using a temporal reasoner that is able to check the knowledge base for consistency, thus eliminating the need to incorporate such checks in the main system. In this way, the designer of the system can focus on the application intelligence rather than enforcing/checking data related restrictions manually. Also, there is an increased support for reuse of temporal reasoning tools across applications.

In this paper we propose a framework for designing semantic, time-aware decision support systems, based on the state-of-the-art tOWL language. The systems that we propose provide the means to efficiently store and retrieve data, and allow (temporal) inference on the represented data. Restrictions on this data can be represented in a generic way at the level of the temporal language. Different (temporal) properties of the entities in the knowledge base can also be represented, both at abstract as well as at concrete level. In this way, data-related operations are separated from the main intelligence component of the decision support system.

* Corresponding author.
E-mail addresses: milea@ese.eur.nl (V. Milea), frasincar@ese.eur.nl (F. Frasincar), u.kaymak@ieee.org (U. Kaymak).

0957-4174/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.eswa.2012.08.001

The framework that we introduce is deployed in a practical context. We illustrate how such semantic, time-enabled decision support systems can be used by means of an example from the financial domain. The finance area has received attention in the development of expert systems, both from a theoretical perspective, as for example in Blue and Andoh-Baidoo (2010) and Matsatsinis, Doumpos, and Zopounidis (1997), as well as in more practical contexts, such as (Falavigna, 2012; Nedovic & Devedzic, 2002; Rada, 2008). We choose the aggregation of market recommendations as a proof-of-concept. Market recommendations are advices, in the form of indicated courses of action, issued by financial analysts, regarding the stock of a certain company. These recommendations most often materialize in *buy*, *hold*, or *sell* advices, and are issued at different times. Since multiple analysts can issue such recommendations, more often than not, at a specific point in time, a company may have recommendations issued by different analysts. When these recommendations diverge, in the sense that there is no consensus within the analyst group whether an asset should be bought, held, or sold, choosing the appropriate course of action might not be obvious. The system that we present investigates which aggregation method for market recommendations gives the best results, given the evidence from the past.

In investigating which approach provides the best results for the aggregation of market recommendations, we consider two different alternatives: a *majority voting* approach, in which we choose the recommendation to which most analysts concur, and an approach that takes into account the *analysts' past performance* when deciding the course of action with the highest expected performance. For measuring the analysts' past performance, we rely on the Sharpe ratio (Sharpe, 1966).

The example that we present contains several features that make it interesting to consider. First, the recommendations issued by analysts either have a limited validity in time, or hold until a new recommendation is issued. By relying on the tOWL language for the representation of recommendations, we can define default durations for advices, and also set as ending valid time for a recommendations the time when a new recommendation is issued by the same analyst, regarding the same company. Determining which recommendations hold at any point in time can also be achieved generically due to the timeslices representation used by the tOWL language. This allows us to determine, at any point in time, which recommendations have been issued by an analyst in the past, which in turn allows for determining the past performance of the analysts.

The outline of the paper is as follows. In Section 2 we provide an overview of the research related to the subject of this paper. The tOWL language, the basis of the framework presented in the paper, is presented in Section 3. The semantic, time-aware framework that we propose is presented in Section 4. The application, as well as the methodology we use for aggregating market recommendations is presented in Section 5. Our results and a discussion of the results are presented in Section 6. Finally, we conclude in Section 7.

2. The temporal dimension in decision support systems

A thorough approach towards the design and evaluation of temporal expert systems is presented in Chinn and Madey (1997). The approach starts by evaluating specific characteristics of expert systems and temporal applications separately, and then formulating a framework that brings both of them together. The application area that is considered for this framework consists of business problems. One of the temporal requirements that the authors formulate relates to support for a time-line view of events as well as the ability to maintain a historical repository of events, requirements that are both supported by the tOWL language. Other requirements formulated in Chinn and Madey (1997) relate to being able to define and use time in different knowledge base constructs and the ability to represent temporal relationships. By relying on the tOWL language for the representation of temporal knowledge, these requirements are fulfilled due to the ability to represent temporal intervals and Allen's interval relationships (Allen, 1983) in the language, as well as timeslices and fluents for representing what is changing.

The importance of a temporal dimension in knowledge bases is also identified in Juarez, Campos, Palma, and Marin (2008) for context-dependent temporal diagnosis. The authors present an integration of Model-Based Reasoning and ontologies. Based on this framework, the authors successfully develop a medical diagnosis system, where the domain knowledge is described in a medical ontology.

The temporal dimension in medical expert systems is also discussed in Kohane (1987). This work relies on the Temporal Utility Package (TUP), which provides, to a limited extent, some of the temporal abilities of modern, temporal Semantic Web approaches. The proof-of-concept application consists of an expert system illustrating temporal reasoning in different phases of the medical diagnostic process.

Temporal reasoning in expert systems in a more general sense is discussed in Perkins and Austin (1990). The authors develop an architecure for a temporal expert system where attribute values can be associated with time tags. The proof-of-concept consists of an expert system used for diagnosing a specific set of problems of the Hubble space telescope.

Similarly, the temporal dimension is also considered for recommender systems in Arroyo-Figueroa, Sucar, and Villavicencio (1998). The model presented is an extension of a standard Bayesian network. This extension introduces temporal nodes for the representation of states or events, and arcs for the representation of causaltemporal relationships between nodes.

Although the temporal dimension in decision support systems has been investigated to a certain extent in the literature, approaches relying on Semantic Web technologies have not yet been considered. Semantic Web languages, such as the state-of-the-art tOWL language, fulfil the domain-independent requirements formulated until now in different studies. Hence, the Semantic Web approach that we propose provides a generic mechanism for dealing with domain-independent aspects of temporal knowledge representation and reasoning in time-aware decision support systems.

3. The tOWL language

The tOWL language (Frasincar, Milea, & Kaymak, 2010; Milea, Frasincar, Kaymak, & di Noia, 2007, 2008, 2012, 2012) is a temporal Web Ontology Language based on the $\mathcal{SHIN}(\mathcal{D})$ description logic, which is an expressive subfragment of OWL-DL (McGuinness & Harmelen, 2004). tOWL is built on top of OWL-DL, which is the current state-of-the-art ontology language and W3C standard. An overview of the different layers introduced by the tOWL language on top of OWL-DL is provided in Fig. 1. The language enables the

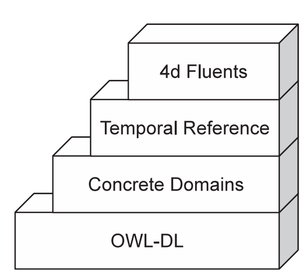

Fig. 1. tOWL layer cake.