Web Semantics: Science, Services and Agents on the World Wide Web 8 (2010) 110–124

Contents lists available at ScienceDirect

Web Semantics: Science, Services and Agents on the World Wide Web

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

Automatic construction of a large-scale situation ontology by mining how-to instructions from the web

Yuchul Jung, Jihee Ryu, Kyung-min Kim, Sung-Hyon Myaeng*

KAIST Computer Science Department, 335 Gwahak-ro, Yuseong-gu, Daejeon, South Korea

ARTICLE INFO

Article history:
Received 11 October 2009
Received in revised form 31 March 2010
Accepted 9 April 2010
Available online 18 April 2010

Keywords:
Automatic ontology construction
Situation ontology
Action mining
How-to instruction
Service recommendation
Automatic service composition

ABSTRACT

With the growing interests in semantic web services and context-aware computing, the importance of ontologies, which enable us to perform context-aware reasoning, has been accepted widely. While domain-specific and general-purpose ontologies have been developed, few attempts have been made for a situation ontology that can be employed directly to support activity-oriented context-aware services. In this paper, we propose an approach to automatically constructing a large-scale situation ontology by mining large-scale web resources, eHow and wikiHow, which contain an enormous amount of how-to instructions (e.g., "How to install a car amplifier"). The construction process is guided by a situation model derived from the procedural knowledge available in the web resources. Two major steps involved are: (1) action mining that extracts pairs of a verb and its ingredient (i.e., objects, location, and time) from individual instructional steps (e.g., <disconnect, ground cable>) and forms goal-oriented situation cases using the results and (2) normalization and integration of situation cases to form the situation ontology. For validation, we measure accuracy of the action mining method and show how our situation ontology compares in terms of coverage with existing large-scale ontology-like resources constructed manually. Furthermore, we show how it can be utilized for two applications: service recommendation and service composition.

1. Introduction

Ontological knowledge has become a main vehicle for semantically and conceptually oriented techniques and applications such as word sense disambiguation, searching, classification, question answering, entity resolution, and context/situation-aware reasoning for personalized services. However, currently available large-scale ontologies often fail to deal with diverse task situations that may arise in the real world because they lack in understanding the dynamic nature of daily lives of people and the associated activities. For example, automatically built ontologies like YAGO [12] driven by WordNet [40] and Wikipedia [39] do not have a sufficient coverage of contextual instances to reason about situations and activities arising from different domains. There is no consideration about such activities of daily living as shopping, driving, wedding, etc., for which the context variables like actions, location, and time should be made available. Without a situation ontology of this kind, it would not be possible to infer what activity the user is engaged in and what actions are likely to be taken from the current situation, which can be characterized with context variables like the current location, objects used, and time.

As a novel solution to the problem, we attempt to build a huge situation knowledge base of human activities by means of text mining techniques that exploit the structure of the how-to descriptions, which is essential for context/situation-aware services. Action level knowledge is extracted from eHow[1] and wikiHow[2], freely accessible websites currently storing more than one million articles on how to do things step by step, which collectively cover almost every domain of daily lives including business, cars, computers, education, health, travel, weddings, etc. An article can be converted into an instance of a situation ontology model that consists of a goal, action sequence, and contextual ingredient that includes location, time, and objects. To organize such knowledge, we have defined a situation ontology specification that includes six ontology classes, *topic*, *goal*, *action*, *object*, *time*, and *location*, and six types of semantic relations, *hasTopic*, *hasAction*, *hasNextAction*, *hasObject*, *hasTime*, and *hasLocation*, all of which are derived from the eHow articles, as in Fig. 2.

* Corresponding author. Tel.: +82 42 350 6149; fax: +82 42 350 6222.
E-mail addresses: enthusia77@kaist.ac.kr (Y. Jung), zzihee5@kaist.ac.kr (J. Ryu), kimdarwin@kaist.ac.kr (K.-m. Kim), myaeng@kaist.ac.kr (S.-H. Myaeng).

[1] http://www.ehow.com/.
[2] http://www.wikihow.com/.

1570-8268/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.websem.2010.04.006

We crawled the entire set of articles from the eHow and the wikiHow websites and applied natural language processing (NLP) techniques to obtain a highly refined situation ontology, which can help detecting the current situation of a user in a daily life and suggesting a solution suitable for the problem at hand if any. The task of the employed NLP techniques is to extract actions expressed in a verb form and associated contextual ingredient items from the goal and subsequent action sequences expressed in natural language in an article. In order to put the linguistic constituents in an ontological form,[3] we designed four additional steps: goal normalization, action normalization, action transition probability calculation, and ingredient resolution.

To assess the utility of the proposed method and its outcome, we measured accuracy and coverage of the automatically constructed ontology. Accuracy was measured by taking a random sample of the situation instances converted from the corresponding articles. We checked whether or not those instances were clear without ambiguity and well-formed. For coverage of the resulting ontology, it was compared for verbs against existing large-scale ontology-like resources: WordNet and OMICS [27].

In this paper, an automatic situation ontology construction based on action mining from the Web is presented to build a large-scale situation ontology that is required to reason about user intentions (or situations) and provide relevant recommendations in a given context. Its main contribution is to show that an automatic methodology can be employed to construct a large-scale situation ontology for the situation model with high precision. Given the dynamic nature of knowledge in people's daily activities, it is critical to devise an automatic method for constructing situation ontologies. Through the application scenarios, we also show that the ontology constructed as such can be of practical value for context-aware applications. We advocate that the high accuracy of the method and the sheer size and utility of the situation ontology lend themselves to further research and development in context-aware applications involving unconstrained daily lives.

Section 2 describes the main features and drawbacks of previous work concerning situation-awareness, situation ontology, and automatic ontology construction to set the stage for our work. In Section 3, we introduce our situation ontology model and the resources from which the current situation ontology is constructed. Section 4 explains the details of our situation ontology construction process focusing action mining and normalization. In Section 5, we present an evaluation of the constructed ontology for its accuracy and comparison to other ontology-like resources. Section 6 shows how the newly constructed situation ontology can be utilized in situation-aware recommendation and semantic web service composition. In Section 7, we give our conclusion and discuss future directions.

2. Related work

The notion of *context-awareness* in ubiquitous computing was proposed in 1990s to address the interaction between computer systems and environments [5]. *Situation-awareness* has also been used to refer to the same meaning [13]. The notion has received a great deal of attention because it is a basis for improving the quality of decisions in a heterogeneous, highly dynamic environment [26]. The meaning of information about the perceived objects can be correctly determined when the situation or context is taken into account.

[3] To focus on the construction process and utility of the ontology in this article, we define a situation model and the corresponding ontological structure, without representing them in a language like OWL.

Ontologies have been proposed to provide a support for situation awareness. In particular upper ontologies serve as a common vocabulary for collaborating agents and information sources [6,13,25,26,36]. Although those efforts succeeded in deriving design requirements [25,26] and developing domain specific models [13], there are few attempts to build a large-scale situation ontology that can handle diverse situations in daily lives. Context/situation-aware systems are still at an infant stage as far as knowledge support is concerned. To the best of our knowledge, there has been no explicit effort to automatically construct a large-scale knowledge base for understanding user situations and activities of daily living in various domains.

While not directly concerned about situation awareness, there have been some attempts to extract semantic relations among concepts or entities that exist in a corpus, which are often simply words and phrases. A study [21] used a set of lexico-syntactic patterns that occur frequently to extract a lexical relation of interest from a large text corpus. It is an early attempt to extend WordNet for hyponyms. Another study [10] introduced an automatic method to enrich WordNet by using the web. To overcome the limitations of WordNet, such as lack of relations between topically related concepts and proliferation of word senses, it linked document collections from the web to concepts in WordNet. It was shown that the Web resource could be practically used to collect lists of words that are topically related to the concept (called as topic signature) and to discriminate different word senses by clustering the concepts that lexicalize them.

In the same context, attribute extraction has been the subject of recent investigations within the information extraction community. Extracted attributes are considered relations between entities (or objects or values). Several studies [9,17,28,35] attempted to acquire attributes, possibly along with corresponding values, from Web documents. The method proposed [20] employs lexico-syntactic patterns for unstructured text in a small collection of Web documents. More recently, weakly supervised approaches [17,18,35] for unstructured Web documents were introduced to deal with a large number of classes and extract attributes which are not restricted to any pre-defined pattern types (e.g., X-of Y patterns). While related, they do not deal with situations in general nor actions and ingredients in particular.

There have been more explicit efforts to build an ontology automatically using a single resource for background knowledge (e.g., WordNet or Wikipedia) or multiple resources. For example, YAGO ontology [12], a large ontology extracted from Wikipedia and WordNet, is known to have high quality, over 95% accuracy. It comprises not only concepts in the style of WordNet, but also named entities like people, organizations, geographic locations, books, songs, products, etc. They are related among themselves with various relations such as what-is-located-where, who-was-born-when, who-has-won-which-prize, etc. Because of the high quality and large quantity of the data in the reusable RDF format, it can be used to enhance the performance of existing applications and to facilitate creations of new application for semantic web. Similar to YAGO, DBpedia [7] aims to extract structured information from Wikipedia (especially, semi-structured templates of Wikipedia) and uses RDF to represent the extracted information. It allows users to query relationships and properties with Wikipedia resources. In addition, the dataset is interlinked on RDF level with various other open datasets on the Web including: Freebase, OpenCyc, UMBEL, GeoNames, and DBLP. However, YAGO and DBpedia can hardly support reasoning about diverse cases people are situated in because they were obtained mainly by excavating characteristics (or properties) of existing named entities in Wikipedia.

Our work is unique in that we do not deal with an arbitrary natural language corpus that is difficult and inefficient to process using today's technology. Rather we deal with special-purpose web