Web Semantics: Science, Services and Agents on the World Wide Web 16 (2012) 33–41

Contents lists available at SciVerse ScienceDirect

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

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

BOTTARI: An augmented reality mobile application to deliver personalized and location-based recommendations by continuous analysis of social media streams

Marco Balduini [a], Irene Celino [b], Daniele Dell'Aglio [b], Emanuele Della Valle [a,b,*], Yi Huang [c], Tony Lee [d], Seon-Ho Kim [d], Volker Tresp [c]

[a] Dip. Elettronica e Informazione – Politecnico di Milano, via Ponzio 34/5, 20133, Milano, Italy
[b] CEFRIEL – ICT Institute, Politecnico di Milano, via Fucini 2, 20133, Milano, Italy
[c] Siemens AG, Corporate Technology, Otto-Hahn-Ring 6, 81739 München, Germany
[d] Saltlux, 7F, Deokil Building, 967, Daechi-dong, Gangnam-gu, Seoul 135-848, Republic of Korea

ARTICLE INFO

Article history:
Received 23 January 2012
Received in revised form
10 May 2012
Accepted 12 June 2012
Available online 26 June 2012

Keywords:
Social media analysis
Mobile app
Personalized recommendation
Location-based recommendation
Stream reasoning

ABSTRACT

In 2011, an average of three million tweets per day was posted in Seoul. Hundreds of thousands of tweets carry the live opinion of some tens of thousands of users about restaurants, bars, cafes, and many other semi-public points of interest (POIs) in the city. Trusting this collective opinion to be a solid base for novel commercial and social services, we conceived BOTTARI: an augmented reality application that offers personalized and localized recommendation of POIs based on the temporally weighted opinions of the social media community. In this paper, we present the design of BOTTARI, the potentialities of semantic technologies such as inductive and deductive stream reasoning, and the lessons learnt in experimentally deploying BOTTARI in Insadong – a popular tourist area in Seoul – for which we have been collecting tweets for three years to rate the hundreds of restaurants in the district. The results of our study demonstrate the feasibility of BOTTARI and encourage its commercial spread.

1. Introduction

Imagine that you are a tourist in Seoul. You would like to dine out. You prefer to avoid *tourist traps* and dine where the locals do. You have been told that Insadong district would be the perfect place; it offers a choice of more than a hundred restaurants in two square kilometres, and most of the district is reserved for pedestrians.

When you reach Insadong-gil (the main street of the district), you find yourself surrounded by hundreds of restaurant advertisements (see Fig. 1). You know that you can still open your guide book and choose one of the few restaurants listed there, but you definitely want a place where the locals go. You take out your mobile and check various apps that recommend restaurants based on users' reviews. The number of user-rated restaurants is smaller than you expected: only ten restaurants are rated more than ten times. This is probably because you are in Seoul, one of the cities world-wide in which people *tweet* a lot.[1] You wish that a service existed that continuously analysed the social media streams and that could show you how the locals have been rating Insadong's restaurants over the last few months.

This is exactly what we designed BOTTARI for. BOTTARI is an augmented reality application for personalized and localized restaurant recommendations, experimentally deployed in the Insadong district of Seoul. At first glance, it may appear like other mobile apps that recommend restaurants, but BOTTARI is different: BOTTARI uses inductive and deductive stream reasoning [1] to continuously analyse social media streams (specifically Twitter) to understand how the social media users collectively perceive the points of interest (POIs) in a given area, e.g., Insadong's restaurants.

In this paper, we describe the choices we made in designing BOTTARI and the lessons we learned by experimentally deploying it in Insadong. The paper is organized as follows. Section 2 introduces relevant background. Section 3 illustrates the BOTTARI mobile app from the user's point of view, i.e., the main task being pursued. Section 4 explains how to understand the data used in experimentally deploying BOTTARI in Insadong district. Section 5

* Corresponding author at: Dip. Elettronica e Informazione – Politecnico di Milano, via Ponzio 34/5, 20133, Milano, Italy.
E-mail addresses: marco.balduini@polimi.it (M. Balduini), irene.celino@cefriel.it (I. Celino), daniele.dellaglio@cefriel.it (D. Dell'Aglio), emanuele.dellavalle@polimi.it (E. Della Valle), yihuang@siemens.com (Y. Huang), tony@saltlux.com (T. Lee), shkim@saltlux.com (S.-H. Kim), volker.tresp@siemens.com (V. Tresp).

[1] An average of three million tweets were posted each day in Seoul in 2011.

1570-8268/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.websem.2012.06.004

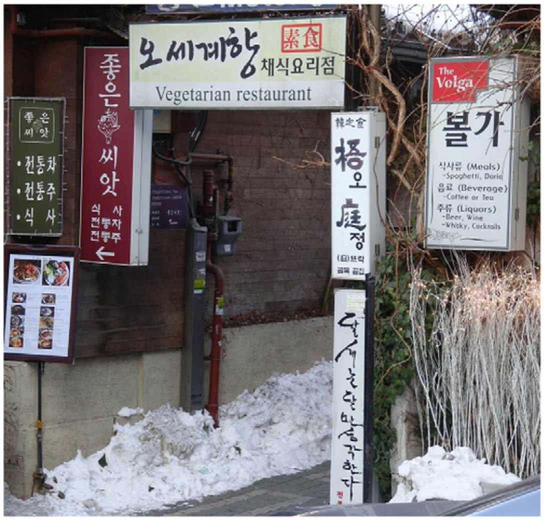

Fig. 1. A picture of Insadong district: the density of restaurants is very high (see also Section 4).

briefly illustrates the ontology at the core of BOTTARI that was used to integrate the available information. Section 6 presents the BOTTARI back-end. Sections 7 and 8 report our evaluation results both in terms of quality of recommendations and scalability of the BOTTARI back-end. Finally, Section 9 concludes the paper by discussing the lessons we learnt and sketches our future works.

2. Background work

In this section, we briefly illustrate the context in which the idea of BOTTARI was conceived and the technological *ingredients* we used in implementing it.

2.1. How tourists locate points of interest in the 2010s

In the first half of the 2000s, the popularity of tourist guide books faded out. The rise of Web 2.0 impacted the tourist market: just as virtual travel agencies replaced physical ones, so collaborative tourist guides such as TripAdvisor,[2] Yelp,[3] or Qype[4] eroded the market of guide books. Today's tourists commonly refer to those Web sites when planning a travel.

In the second half of the 2000s, the search engine industry intercepted this trend by introducing the concept of *local search*: search facilities to find POIs described on the Web while on the go [2–4]. In local searches, users' preferred POIs are shown: those located close to the user's current position [5–7], those for which ratings are available [8,7], and those that match user preferences [9,10]. Nowadays, a large number of tourists search the Web for POIs while on the go.

In the 2010s, the increasing availability of GPS-enabled smart phones allowed the wide-spread of location-based services (LBSs), i.e., mobile and Web applications that ask users to *check in* and, in response, provide context-dependent information and services. Examples of LBSs are foursquare,[5] Gowalla,[6] and Facebook Places.[7] LBSs can almost *listen to the pulse of the city*; therefore many tourists, when choosing which museum to visit or where to dine out, look up the POIs with the highest number of check-ins on LBSs.

2.2. Stream reasoning

In 2008, Della Valle et al., in [11], asked the Semantic Web community for techniques able to *reason upon rapidly changing information*. When reasoning on massive data streams, such as those characterizing BOTTARI, well-known artificial intelligence techniques have the right level of expressivity, but their throughput is not high enough to keep pace with the stream (e.g., belief revision [12]). The only technological solutions with the right throughput are Data Stream Management Systems (DSMSs) [13] and Complex Event Processing [14]; however, they are not sufficiently expressive. A new type of inference engines was thus needed to reason on streams. Della Valle et al. named them *stream reasoners*.

In the following years, a number of stream reasoning approaches were developed [15–18]. They share three main concepts: (a) they logically model the information flow as a resource description framework (RDF) stream, i.e. a sequence of RDF triples annotated with one or more non-decreasing timestamps, (b) they process the RDF streams "on the fly", often by rewriting queries to the raw data streams, and (c) they exploit the temporal order of the streaming data to optimize the computation.

BOTTARI uses both a deductive and an inductive stream reasoner. The *deductive stream reasoner* is based on Continuous SPARQL (C-SPARQL) [15]—an extension of SPARQL that continuously processes RDF streams observed through windows (as done in DSMSs). The syntax and semantics of C-SPARQL are described in [19]. The C-SPARQL execution engine and its optimization techniques are illustrated in [20]. The optimization needed for high-throughput RDFS++ reasoning are described in [21]. The approach to publish an RDF stream as linked data (namely streaming linked data) is based on [22].

The *inductive stream reasoner* is based on the SUNS (statistical unit node set) approach [23,24]—a scalable machine learning framework for predicting unknown but potentially true statements by exploiting the regularities in structured data. SUNS employs a modular regularized multivariate learning approach able to deal with very high-dimensional data [25] and to integrate temporal information using a Markov decomposition [26].

2.3. The LarKC platform

The *Large Knowledge Collider* (LarKC) is an EU FP7 Integrated Project [27] that aimed at massive distributed incomplete reasoning. The LarKC platform [28] is one of the main results of the project. It is a pluggable Semantic Web framework that can be deployed on a high-performance computing cluster. It allows for orchestrating multiple heterogeneous units for data processing and reasoning (named plug-ins), and for exposing their aggregated capabilities as a SPARQL endpoint.

3. The BOTTARI mobile app

As shown in Fig. 2, BOTTARI is an Android application (for smart phones and tablets) in augmented reality (AR) that directs the users' attention to restaurants and dining places in the neighbourhood of their position.

In the Korean language, "bottari" is a cloth bundle that carries a person's belongings while travelling. BOTTARI carries the collective perceptions of social media users about POIs in an area and uses them to recommend POIs. As shown in the upper-left corner of the screenshot in Fig. 2(a), BOTTARI users can search POIs in their proximity using four buttons.

2 See http://www.tripadvisor.com/.
3 See http://www.yelp.com/.
4 See http://www.qype.com/.
5 See http://foursquare.com/.
6 See http://gowalla.com/.
7 See http://www.facebook.com/about/location.