A Semantic Sommelier as an Ontology-Powered Mobile Social Application and a Pedagogical Tool

Deborah L. McGuinness
Rensselaer Polytechnic Institute
110 8th Street
Troy, NY, USA, 12180
dlm @ cs.rpi.edu

Evan W. Patton
Rensselaer Polytechnic Institute
110 8th Street
Troy, NY, USA, 12180
pattoe @ cs.rpi.edu

ABSTRACT
Our semantic sommelier is the latest evolution in a long line of ontology-enabled wine advisors. This version utilizes mobile platforms (currently Apple's iOS and Google's Android) and some of their related capabilities such as global positioning information and portable nature to be more context-aware and ubiquitously available. It also accommodates input from social tools having been integrated with Facebook, Twitter, and wikis. Most recently it has been used to demonstrate the potential of semantic technology and has seen success with the press. We will briefly introduce the semantic sommelier, include connections to its lineage going back to the 1980s, discuss current directions, and include our perspective on its current success with journalists.

Keywords
Semantic Web, Wine Ontology, Wine Agent, Recommender System, Advisor, social media, mobile

1. INTRODUCTION
Numerous demonstration systems have been created by the author as well as others using some kind of semantic technology to give wine advice. The author created her first wine and food pairing system in an expert systems class in the 1980s. Since that time, she has encoded knowledge about wine and food and their pairings in multiple artificial intelligence and ontology languages including LISP, Prolog, DAML, DAML+OIL, and OWL and some instantiation of this knowledge has been integrated into prototypes for decades. These prototypes have been used to demonstrate how relatively simple semantic encodings can be used to power advising systems. They started as pedagogical tools and were often written about in overviews, tutorials, or best practice documents concerning semantic technologies (e.g., [1], [3], [4], [6], [8]). They have been used successfully for years as relatively stand-alone demonstration systems and relied originally on pre-encoded expert or semi-expert generated knowledge bases of pairing rules and wine and food instance data. As more information became accessible online, these demonstration systems began to leverage not just internal knowledge bases but also knowledge that was generated by other sources, for example the Stanford KSL wine agent [2] utilized its internal knowledge base of wines but also queried some online sources of wines such as wines.com to find matching content. Our next significant update expanded into collaborative settings including using Wikis for input [5]. Our latest update [7] does not limit itself to internal knowledge bases or simple querying of a few known sources for wines, but it now uses recommendation content that can be gathered from a user's social network along with content that can be scraped from the web concerning local menus and wine lists. Further it no longer requires a server or PC to run; it now can work on a user's mobile device. The basic task is the same: given a description of a meal (or meals), determine a description of a wine that would pair nicely, and then use that description as a query to some (perhaps many) wine lists to make a specific recommendation (or set of recommendations).

WebSci '11, June 14-17, 2011, Koblenz, Germany.

2. SYSTEM DESCRIPTION
In many ways, the Semantic Sommelier is not new in terms of its basic use of ontologies, recognition, and reasoning. What is new is its integration with new platforms and its connections to a wider range of applications and data. Now, instead of only using its internal recommendations for wine pairings, it can also use advice from a user's friends in their social networks. Additionally, it now has the option, using the phone's GPS, of knowing where the user is and thus which restaurant and wine store the user is near. It also functions on devices that can easily be used at a dinner setting, e.g., smart phones, so interactive exploration is more possible. Users no longer need to use it in advance or rely on a laptop with internet access in order to use it in a restaurant setting.

As with many efforts, taking even a well-used, long-lived demonstration system and updating it so that it can be used by a broader range of users with new platform issues generates a wealth of work. Now that we are integrating the agent with existing restaurant menus and wine listings, the ontology is being expanded. Extensions include more types of wine (e.g. new or esoteric varietals) and foods (including more preparation styles as well as more ingredients).

More significant issues result from the platform expansion onto smart phones. Interface issues arise because of limited real estate for displaying data. We have experienced much more work related to putting the agent on multiple platforms (iPhone and Android) and even more because of dealing with small devices. New issues of power consumption, limited and/or distributed processing, and memory constraints now all become issues. Connecting with social platforms also generates issues of trust, conflict resolution, and prioritization that did not exist in previous versions, or at least did not exist to this level of importance.

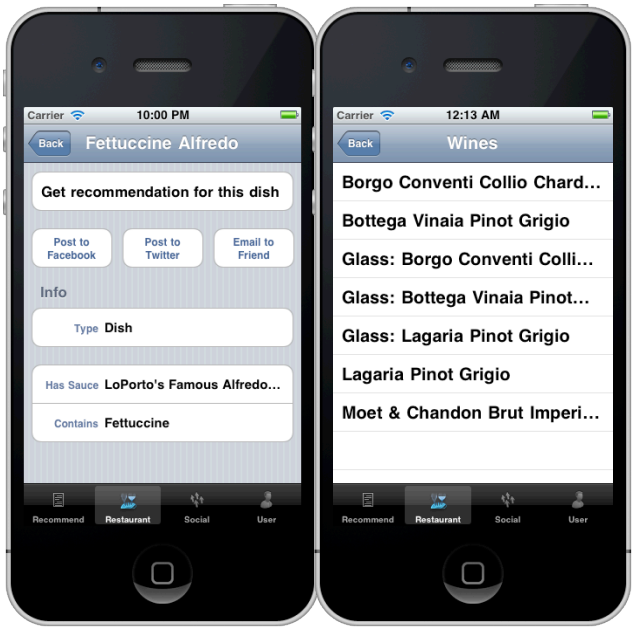

Figure 1. (a) Left, a description of Fettucine Alfredo as prepared by LoPorto's restaurant; (b) Right, a selection of wines from LoPorto's wine list that pair with the Fettucine based on the built-in recommendations.

We still are working on a number of the issues mentioned above but we do have a demonstration system that utilizes user positioning, local restaurant and wine listings, social network recommendations as well as our pre-programmed advice for parings. Using the GPS features of the device, it is possible to detect the user's location and identify nearby restaurants. We have converted a number of restaurant menu and wine lists into structured descriptions that are able to be parsed by the Wine Agent. Upon loading the menus, the user can choose to select a dish as shown in Figure 1a, where Fettucine Alfredo has been chosen. The agent will display information about the dish to the user and include options to share that dish and its suggested wine pairings and comments with friends via Facebook, Twitter, or e-mail. To get a recommendation using the default recommendation set, the user taps the "Get recommendation for this dish" entry, and the reasoner will realize, classify, and then identify recommendations that share common dimensions with the selected dish, and provide the user with a set of descriptions for wines (not shown). Tapping a recommendation will initiate a search on the wine list for wines that complete the pairing, and a list will be rendered (Figure 1b). From this point, users can continue to explore by selecting wines to learn more about them, or learn the source of the recommendation, or backtrack to find a different dish to select. For example, the user may choose one of the selections and see that the recommendation for a white dry wine came from one of his/her Facebook friends or it could be that a Twitter posting was from a friend who had enjoyed this particular wine with the particular dish recently at this restaurant and thought it was a good pairing.

3. News on the Web

We recently had an internal communications expert release an article about the Semantic Sommelier with a title that was still relatively technology-oriented (Wine Application Highlights the Power of Web 3.0)[1]. The article did, however, begin with a very approachable notion – the restaurant of the future where one might always have the perfect wine and food pairings and also know one's dining companions' wine preferences as well. We found it interesting that this article was re-published within hours of its initial publication in dozens of news locations ranging from Communications of the ACM to broad science reporting venues like Science Daily, to more specialized science settings like physorg.com, esciencenews, and the Web Science Trust, to wine reviews such as redwinereview, and even health publications such as HealthWD. It also caught the eye of international journalists in Spain and Canada.

Given the spread of this article across the web, we sought to collect and analyze information about the article and its references. Conveniently this article contained a phrase "semantic sommelier" that was unique to this work so tracking articles was relatively easy to do with web searches. For our analysis, we collected the first 500 results of a Google search for "Semantic Sommelier" and, using a few students, captured a number of different dimensions about each resulting page in the search: the source article being cited, the language of the surrounding website (the article was almost always republished in English), the date of the posting, the type of post, i.e. word-for-word repost, a short teaser paragraph, a blog posting discussing the article, post on an aggregation site (e.g. Digg, Reddit), tweets, spam pages, deadlinks (i.e. pages that Google crawled but yielded HTTP 404 errors or pages no longer mentioning the semantic sommelier), whether or not the post presented an original commentary on the article, how many comments were posted about the article if the site accepted comments, any re-tweets of the post, and whether or not the site publishing the content required payment to access the article.

The data were collected over a period of three days by capturing the URLs of articles in Google, writing them to a spreadsheet,

Figure 2. A day-to-day breakdown of posting numbers based on the publication of the original article. Over half of articles analyzed were posted within the first 24 hours of the original article, with a long tail of publications up to 67 days later.

and then analyzed by students to determine their relationship to

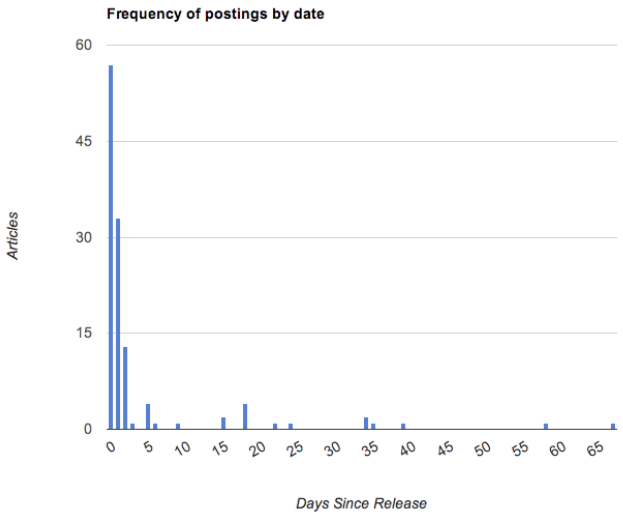

[1] http://news.rpi.edu/update.do?artcenterkey=2830

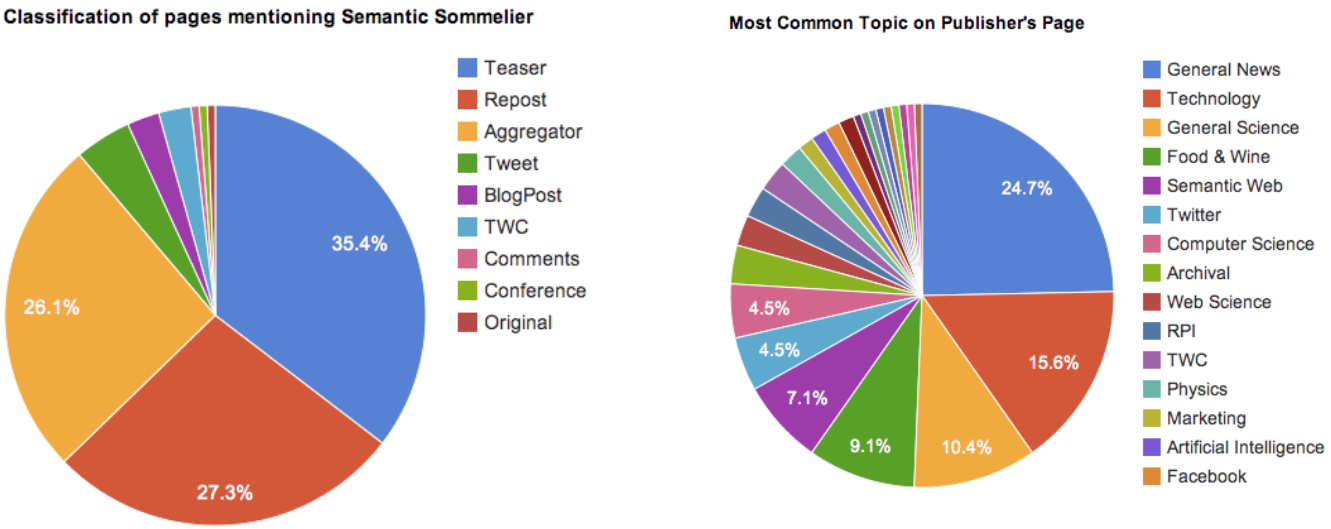

Figure 3. (a) Left: A breakdown of the different categories of postings. Short teaser paragraphs were most common and were accompanied with a link back to the original source, followed by complete re-posts of the article, usually with attribution although sometimes without, followed by short descriptions based on RSS feeds aggregated by sites. (b) Right: Over one third of posts were on sites aimed toward general areas, with a smaller number of sites being specifically focused on the areas of Food & Wine, Computer Science, and Artificial Intelligence, with 9.1%, 4.5%, and 1.3%, respectively. Note: categories are arranged largest to smallest.

the original article. The spreadsheet was then converted into an RDF graph and loaded into a triple store powered by the Joseki SPARQL endpoint. A number of SPARQL-powered visualizations were constructed using this endpoint in conjunction with the Google Chart Tools[2] (formerly Google Visualization Toolkit).

4. Observations

Three dimensions of particular interest were the types of posts generated by the original article and its derivatives, the amount and quality of discussions generated by those postings, and the frequency of reposting that occurred relative to the publication of the original article.

We examine first the re-publications as a function of time since the original article was released, as displayed in Figure 2[3]. Within the first day, there were close to 60 re-publications, most of which were word-for-word reproductions or short teasers linking back to a few key hubs. The subset of postings captured using this method demonstrated a power-law with respect to time.

Twenty different sites cited the original article posted by RPI, which was the second-most cited of the postings. ScienceDaily.com was the most cited source with 33 citations and cited the original article. Additional sites producing a large effect include EurekAlert, PhysOrg, NewsWise, and Communications of the ACM, with 15, 13, and 3, citations, respectively. Collectively, posts originating from outside RPI were cited 95 times whereas the RPI articles only accounted for 22 citations[4]. This indicated that getting the information into the right portals could have a substantially larger effect on the success of the project than a more conservative approach might have done.

We noticed that the article was being re-published in a broad range of areas so we attempted to understand the breadth and level of interest. We did an analysis of the topics on the different sites and assigned categories to the overall nature of the website in question. About one quarter of the sites where the article was posted were focused on general news and ten percent were generally focused on the sciences. Fifteen percent of sites were solely oriented in the technology space focusing on both hardware and software, and more focused categories such as computer science or artificial intelligence (4.5% and 1.3%, respectively) were less prevalent but still represented. We also wanted to see how much interest there might be for this technical application in the domain area of food and wine. We noticed that 10% of the overall postings were in food and wine publications.

5. Discussion

Different versions of the wine agent have provided semantic enhancements to the task of providing wine advice. We notice however a significant change now that the agent utilizes more real world data. Also, now that the semantic sommelier can leverage not just expert knowledge but the knowledge of a user's trusted social network, the agent can be more of a customizable social machine where the user can choose whose advice to include and whose to exclude. Further the user can interact with the agent to identify the source of the knowledge used.

Different versions of the wine agent have had some amount of press coverage. Wine and food, after all, is an approachable and

[2] Google Chart Tools

[3] All visualizations are available online at http://wineagent.tw.rpi.edu/vis/websci2011/index.html and will be further extended as analysis continues.

[4] These counts exclude articles where no source was mentioned, spam pages, and links that returned 404 errors.

appealing domain. (That was one of the reasons it was chosen initially as a demonstration domain). This effort is a bit more easily tracked than some since this article was the first place we used the phrase Semantic Sommelier so search results are more easily tracked and most appear to be us.

The interest level in this article and our work seems to be at a different level than we experienced previously. We speculate that the combination of the real world setting, use of current data, the setting of the mobile devices, and the integration with social platforms provides a sweet spot. We believe that the fact that the new system supports much more context-sensitive information is key. This system can use personal location along with local restaurant content and recommendations from an individual's social network thereby providing a much more personal experience. We plan to explore this hypothesis further and hope to report more news in a longer paper and at Web Science in few months. Further, we intend to field our system in some local restaurants and then collect usage data and comments about its features. One aspect of usage that we intend to evaluate is user reliance on advice from social sources vs. user reliance on advice from expert sources.

Acknowledgements: Thanks to Yu Chen, Philip Ng, Bharath Santosh, and Daniel Souza for help with mobile platform issues and article analysis. This work is partially supported by the following funding sources: a National Science Foundation Graduate Research Fellowship awarded to Evan Patton, the Tetherless World Constellation, and gifts from Lockheed Martin, Fujitsu Labs America, LGS, and Microsoft.

6. REFERENCES

[1] Brachman, R.J., McGuinness, D.L., Patel-Schneider, P.F., Resnick, L.A., & Borgida, A. Living with CLASSIC: When and how to use a KL-ONE-like language. Principles of Semantic Networks: Explorations in the Representation of Knowledge, 1990, 401-456.

[2] Hsu, E., and McGuinness, D.L. KSL Wine Agent: Semantic Web Testbed Application. In 16th Intl Workshop on Description Logics (Rome, Italy, September 2003).

[3] McGuinness, D.L., Abrahams, M.K., Resnick, L., Patel-Schneider, P.F., Thomason, R., Cavalli-Sforza, V., and Conati, C. ``CLASSIC Knowledge Representation System Tutorial." Artificial Intelligence Principles Research Department, AT&T Labs Research, 1994. http://www.bell-labs.com/project/classic/papers/ClassTut/ClassTut.html.

[4] McGuinness, D.L., and van Harmelen, F. OWL Web Ontology Language Overview, February 2004.

[5] Michaelis, J., Ding, L., and McGuinness D.L. The TW Wine Agent: A Social Semantic Web Demo. In Poster and Demonstration track of the 7th International Semantic Web Conference, (Karlsruhe, Germany, October 2008).

[6] Noy, N.F., and McGuinness, D.L. Ontology Development 101: A Guide to Creating Your First Ontology. Stanford Knowledge Systems Laboratory Technical Report KSL-01-05, March 2001.

[7] Patton, E.W., and McGuinness, D.L. The Mobile Wine Agent: Pairing Wine with the Social Semantic Web. In the 2nd Social Data on the Web workshop, (Washington, D.C., USA, October 2009).

[8] Smith, M.K., Welty,C. and McGuinness, D.L. OWL Web Ontology Language Guide. World Wide Web Consortium (W3C) Recommendation. February 10, 2004. Available from http://www.w3.org/TR/owl-guide/