Conceptual Enrichment of Locations Pointed Out by the User

Ana Alves[1], Raquel Hervás[2], Francisco C. Pereira[3], Pablo Gervás[2] and Carlos Bento[3]

[1] Instituto Superior de Engenharia de Coimbra,Portugal
aalves@isec.pt,
[2] Universidad Complutense de Madrid, Spain
raquelhb@fdi.ucm.es,pgervas@sip.ucm.es,
[3] Universidade de Coimbra, Portugal
camara@dei.uc.pt,bento@dei.uc.pt

Abstract. The growing availability of GPS and localization technologies in general opens up new challenges regarding the use of position information for Location Based Services. As a side effect, several research questions arise, one of which we consider in this paper: "What is the meaning of place?". Here, we present two ongoing projects that we believe can contribute to the answer. The first regards the semantic enrichment of place, starting from a simple description of the place towards a concept network; the second focuses the natural language description of this network, as a final output to the user. We present these two projects, focus on their intersection and give an example of the end results.

1 Introduction

A major question regarding places it is how we describe them. Considering as an example the British Library we have: "Position (LatitudeX, LongitudeY)", "Building at Euston Road, 96", "British Imperial Library", "A place for reading", "A library", etc. Clearly, given only the exact global position of a place (LatitudeX, LongitudeX), we do not necessarily know something of it or understand its meaning as a place. Humans do not think in longitude or latitude terms. Furthermore, the functional purpose or category of a place may be easily recognized by humans, but it can not be used by an automatic reasoning system without applying external sources of knowledge. This problem, nicely presented by Hightower [1], is being tackled by many (see Related Work - section 2).

Nowadays, freely available commonsense Ontologies, like WordNet[2], are helpful tools to deduce semantic meaning of several concepts, including places. This can lead to infer other concepts related to a place category through semantic relations between them. Normally, these concepts and their semantic relationships are attached to a generic perspective, thus not representing any instance of them. For example, the concept of "library" can be generically described as "a building that houses a collection of books and other materials" (in

WordNet [2]), but if we talk about a specific library (e.g. U.S. National Library of Medicine), further exploration is needed to grab a more precise meaning of that place, which is commonly not possible using commonsense Ontologies.

On the other hand, the World Wide Web (Web) contains an enormous amount of information with great redundancy. Following its growth, many Information Retrieval (IR) algorithms has been developed to perform efficient search and to index massive Web databases. Until now, the vast majority of this information is represented as texts and lacks structure and semantics. Taking advantage of these facts, we propose an approach to our first question considering work well known in Artificial Intelligence (AI) areas: Geographical IR, Natural Language Processing (NLP) and Ontologies. Geographical IR algorithms and NLP (see Related Work - section2) can help us to find, to extract relevant knowledge and to categorize specific information about a place name and geographical reference, mostly present in user written text on the Web; while Ontologies will allow the assessment of relatedness of new concepts given a place category.

All these techniques can provide a good conceptual framework, susceptible of being processed automatically in the context of particular applications, related to a given location. However, the very same characteristics that make such material machine-tractable tend to make it obscure to understand by everyday users. A traveller looking up a particular restaurant on his PDA while wondering along the street does not necessarily want an XML file that encodes all available information about it. Instead, he would probably be much happier if he could obtain a simple textual description of the most salient facts about it. If available information about his past preferences or experiences were available to the system, and it could be in some way applied to customise such a textual description appropriately, the user experience would be much better. This is where Natural Language Generation techniques can play a significant role in extracting the most benefit from the wealth of information provided by enrichment methods.

The present paper describes a two stage approach to the practical treatment of locations in service-oriented applications. During the first stage, available locations are enriched by resorting to various techniques. During the second stage, textual descriptions of the locations are automatically generated from the resulting enriched conceptual representation of the location.

2 Related Work

An inspiring paper [1] by Jeffrey Hightower argues that location must have more associated information than simply the absolute position in a global coordinate system. We may split this problem into two perspectives: **identify which places are "meaningful"**, [3, 4] which groups of GPS readings are build by a clustering algorithm and then attach its user-defined meaning; and **focus on a set of "meaningful locations" constructed by a community** [5, 6] and decide, for each context, which correct meaning to use. In the last case, which we are particulary more interested, the most common type of representation of these locations is by a tuple (LatitudeX, LongitudeX, NameX) also known as Point of Interest (POI). In such environments, ambiguitous names can be given to the same place (e.g vegetarian restaurant, macrobiotic food or Green Menu), or different locations may share an identical place name (e.g. "vito's pizza" in Coimbra, in Dublin, in Newcastle).

Common Sense Ontologies (such as WordNet, OpenCyc [7], ConceptNet [8] or others) are collections of trivial and semantic knowledge that allow the extension of the computational reasoning process. Their generic concepts and relationships that we specifically find for a known category of places (restaurant, museum, hospital, cinema, pharmacy, etc.) can be used to build a *CommonSense Place Ontology* comprising not only semantic related concepts to a given category but all concepts referred by descriptive definitions (or glossaries). This ontology will be discussed in more detail in section 3.

With the growth of the World Wide Web, some researchers have already considered the convergence of web-related technologies and ubiquitous computing. A good example of this was the CoolTown project [9]. Its main goal was to support web presence for people, places and things. They used Universal Resource Identifiers (URIs) for addressing, physical URI beaconing and sensing of URIs for discovery, and localized web servers for directories in order to create a location-aware ubiquitous system to support nomadic users. In the same line, Ubiquitous Web [10] was envisioned as a pervasive web infrastructure in which all physical objects are socially tagged and accessible by URIs, providing information and services that enrich users experiences in their physical context as the web does in the cyberspace.

It is widely accepted that the majority of on-line information is composed of unrestricted user-written texts. Information Extraction (IE) is a research subtopic in IR devoted to extract useful information from a body of text. Within this field, Term Extraction and Name Entity Recognition are Natural Language Processing (NLP) tasks.

The general process of text generation takes place in several stages, during which the conceptual input is progressively refined by adding information that will shape the final text [11]. During the initial stages the concepts and messages that will appear in the final content are decided (*content determination*) and these messages are organised into a specific order and structure (*discourse planning*), and particular ways of describing each concept where it appears in the discourse plan are selected (*referring expression generation*). This results in a version of the discourse plan where the contents, the structure of the discourse, and the level of detail of each concept are already fixed. The *lexicalization* stage that follows decides which specific words and phrases should be chosen to express the domain concepts and relations which appear in the messages. A final stage of *surface realization* assembles all the relevant pieces into linguistically and typographically correct text. These tasks can be grouped into three separate sets: *content planning*, involving the first two, *sentence planning*, involving the second two, and surface realization.

3 Semantic Enrichment of Place

The problem of "position to place" is a well known challenge within the area of Ubiquitous Computing and relates deeply with the connection humans have with places, their functionality and meaning. Attached to a tag name, even when a category is included, a place needs a richer semantic representation in our perspective. This knowledge can be used for whatever processes that demand semantics of place (e.g. understanding POIs while in navigation; searching for a place that has specific characteristics; route planning using locations with specific functionalities; inferring users activity, etc.). We name this process as Semantic Enrichment of Place and it consists of using available Common Sense Ontologies and Web information to build a collection of generic and instance facts about these places.

3.1 Generic Place Model Ontology

It represents a collection of commonsense and generic information about well-known place categories, like restaurants, cinemas, museums, hotels, hospitals, etc. At a first stage, this information is manually extracted from WordNet by Semantic Augmentation (see subsection 3.2 for more detailed description), but as the system is used, it is dynamically fed by new examples, and thus enriched by specific extracted information of these instances. Information extracted from WordNet has a different weight in instance classification than the information that comes from the Web, since the former is more generic and applied to all individuals pertaining to the same place category.

3.2 Categorization of New Places

Given collected POI from Web databases, we conceive a semantic enrichment system for places comprising by the following three sequential tasks:

1. *Geographical Web Search* is responsible for finding Web pages using only POI data as keywords: place name and geographical address. This search is presently made by the freely available Yahoo API. As we described in Related Work section 2, Geo-Annotated Web Pages are not available yet and are not the central goal of this work. For the moment, we are using a simple heuristic that uses the geographical reference as another keyword in the search. We think that works like Geotumba[12] will contribute to achive better precision for this process. At the end of this module the N more relevant pages are selected.

2. *Concept Extraction* from Web pages found earlier is made at this point. This processing includes POS tagging and word sense disambiguation using available NLP tools. On completion of these two tasks for each web page, we are able to extract the most relevant terms that will be used in classification task. Since WordNet does not cover widely proper nouns yet, we are only

considering common nouns as valid concepts. In our work, we attempt to propose a different approach to Location Geo-Recognition where we are not only interested in classifying a place name, but also capture the meaning of it (represented by its related concepts).

3. *Place Classification* will be obtained from the semantic distance between the concepts extracted before and Place Models, computing their semantic similarity measure [13]. This approach is not new [14] and we try to enhance its performance by doing off-line pre-processing of the related concept semantic similarities to perform faster categorization of new examples. We'll use the Generic Place Model Ontology to classify new examples. Thus, on one hand we have different place types and, on the other hand, we have distinct concept lists originated from the respective web pages. We believe that at this point we can conclude with some degree of confidence that the most scored category will eliminate ambiguity and noise from Geographical Web Search. Naturally, some problems will arise, like *to distinguish weighting between generic and instance concepts in the categorization task*, or *to have multiple possible categories to classify a new example*. We believe some well-known Machine Learning (ML) strategies can tackle to these questions.

4 Generating Textual Descriptions from the Enriched Conceptual Representation

The input for this stage consists in a enriched conceptual description of the POI to be described. This description is represented as a list of facts, where each facts contains one or two concepts and a relation between them. These relations are defined during the enrichment process of the POI, and they extend from the WordNet sense assigned to a concept to the relation between a concept and one of its hypernyms.

WN_sense(non-smoking, smokeless.1). *WN_does_not_contain(range_of_wines).* *WN_sense(coffee, coffee.1).* *WN_direct_hype(beverage.1,hot_chocolate.1).* *WN_direct_hype(beer.1,draught_beer.1).* *WN_direct_hype(beverage.1,beer.1).*	*"Boddington Arms serves coffees such as cappuccino and espresso, beers such as draught beer and Guinness, and tea and hot chocolate.* *It admits credit card and debit card.* *[...]* *More information: Range of wines, Wheelchair access,..."*

Table 1. Example of Input and Generated Text

As an example, a piece of the conceptual description of the POI corresponding to the pub Boddington Arms (coming from the previous stage, described in section 3.2) is shown in Table 1.

The process carried out over this representation follows the basic steps in NLG explained in Section 2. The generation of text is a knowledge-intensive process. Not only the information obtained from the enrichment is used, but also additional knowledge about the concepts that may be required for the generator at any stage of the process. In addition, a vocabulary for the concepts in the domain is necessary to express into words the description of the POI.

Content Determination From the whole conceptual information available for the POI only a part corresponds to the specific data that describes it: the initial words obtained for the POI before the enrichment. The rest is additional information obtained during the enrichment, from synonyms to related concepts, also necessary during the generation process but not specifically descriptive for the POI. This includes words that are found in WordNet and contribute to the enrichment - which appear in a *WN_sense* relation assigned to the corresponding WordNet sense -, but also words that not useful during the enrichment because there is no information related to them that can be used - they appear in a *WN_does_not_contain* relation. These words are not completely discarded, but attached at the end of the text as a list of related concepts. The results of this filtering process can be aggregated according to conceptual criteria based on the information obtained during the enrichment. Using relations as hypernymy concepts are organised into small taxonomies of related concepts. For instance, in the example above, a taxonomy of "beverages" is involved, with "coffee" and "beer" as subclasses, and are "draught_beer" as subclass of "beer".

Discourse Planning In order to determine the structure of the discourse needed to convey this information, we have considered that related concepts must be explained together in the text. As a first approximation, the discourse consists of a list of descriptions of sets of concepts, without relations between them. In the example, a possible plan for the discourse would be to describe the menu first, then the drinks, after that the accommodation facilities, and finally the payment options. Each of them would be a small taxonomy of related concepts. This basic solution could be refined by finding relations between the taxonomies of concepts in such a way that each one can introduce the next one. For the moment we have not explored this option.

Referring Expression Generation In the textual description of the POI, various elements would be mentioned one or more times: beverage, Boddington Arms, restaurant,... Different natural language expressions can be used to describe them at each occurrence in the text: "a drink", "the beverage", "it",... The referring expression generation stage decides at each occurrence how to refer to a given element, ensuring that enough information is given to be able to uniquely identify the intended referent in the current context, while avoiding redundant or otherwise unnecessary modifiers. For the moment only the name of the concepts and the corresponding pronouns have been used.

Lexicalization In this stage the structure of the text is already determined by the Discourse Planning, and the form of the concepts by the Referring Expression

Generation. The words to use for each concept or relation are chosen between the different options stored in the vocabulary. For each entry the vocabulary may contain more than one possible word, so some heuristic is needed to choose between different options. For instance, the concept "non-smoking" could be expressed using "smokeless" or "smoke-free". In this approximation there is only one word assigned to each concept or relation stored in the vocabulary.

Surface Realization Surface realization is a task very closely related to linguistic constraints, and reasonable remote from the specific restrictions of particular domains. For this reason, there are several available implementations for reusable modules that handle this task. We have developed a simple implementation that receives as input an abstraction corresponding to the content that needs to be realised as a sentence and unifies it with a grammar of the language in which the output is desired.

5 Conclusions

In this paper, we focussed on the question "What is the meaning of place?", which is becoming salient in many works on Location Based Services (LBS). We followed an approach that has some tradition in symbolic AI - to derive the "meaning" from a network of interrelated concepts. In this case, we showed an approach that is being developed to automatically build the network of concepts (that define a place), and finally described a method for generating a textual interpretation for this network.

The conceptual enrichment of an initial fragment of information can be greatly improved if the final goal for which the resulting information is intended is taken into account. The development described in this paper considers conceptual enrichment for two different tasks: information management - extending available POIs with enough information for the system to be capable of recognising their nature and improving the services provided to the user based on this knowledge; language generation - constructing a textual description of the POI that is at the same time useful and easy to read for the user. The results of the present work indicate that the enrichment required for enabling machine-processing of the POI - searching for associations with related concepts, as in query expansion during information retrieval - is not necessarily a positive contribution to the task of generating a textual description of the POI. In practical terms, it may be useful for the system to know that a POI corresponding to a restaurant has associations with concepts such as waiters, kitchens, bills,...; but to include such elementary information about the concept of restaurant in a textual description of the POI provided to a user would be redundant. In that case, the only information worth transmitting is whatever is exclusive to this particular POI in contrast with other POIs, and this should not include generic truths about restaurants (for instance).

On the other hand, whereas for simple identification of the nature of a POI considering a single level of hypernymy within the WordNet hierarchy might be enough - and in fact adding more might cloud the enriched result with too much clutter, when addressing tasks of discourse planning it may be more interesting to consider the distance to the closest common ancestor, sidestepping most of the actual information in the concepts found on the way.

The solution presented here, by operating in two separate stages, provides a good balance between enrichment for information management and enrichment for language generation. The language generation stage has been designed to perform an initial filtering to identify those items of information that are indeed most relevant as identifiers of a particular POI. To compensate for the possible loss of information resulting from this operation, this module has access to additional knowledge bases, either in the form of WordNet itself or in the form of OWL encoded ontologies, that it can search to find the information it requires to carry out its task in an informed manner. This procedure guarantees both the existence of a suitable enriched representation of POIs for internal treatment by the information system, and acceptable text outputs as descriptions of the POIs that may be both informative and easy to read by the user.

References

1. Hightower, J.: From position to place. In: Proc. of the Workshop on Location-Aware Computing. (2003) 10–12 part of the Ubiquitous Computing Conference.
2. Miller, G., Beckwith, R., Fellbaum, C., Gross, D., Miller, K.: Introduction to wordnet: An on-line lexical database. Int J Lexicography **3** (1990) 235–244
3. Ashbrook, D., Starner, T.: Learning significant locations and predicting user movement with gps. In: Proc. of 6 IEEE Intl. Symp. on Wearable Computers. (2002)
4. Zhou, C., Ludford, P., Frankowski, D., Terveen, L.: An experiment in discovering personally meaningful places from location data. In: extended abstracts on Human factors in Computing Systems, ACM Press (2005) 2029–2032
5. Alves, A., Figueiredo, N., Gomes, H., Pereira, F., Marcelino, M., Bento, C.: The ejaki project: a quality of service regulator for citizens. In: Proc. of the Intl. Conference on Portable Information Devices. (2007)
6. Tagzania: (http://www.tagzania.com/)
7. Opencyc: (www.opencyc.org)
8. Liu, H., Singh, P.: (Conceptnet: A practical commonsense reasoning toolkit) Avaiable at http://citeseer.ist.psu.edu/liu04conceptnet.html.
9. Kindberg., T., Barton, J., Morgan, J.e.a.: People, places, things: Web presence for the real world. In: Proc. of WMCSA2000. (2000)
10. Vazquez, J., J., A., D.L., I.: The ubiquitous web as a model to lead our environments to their full potential. In: W3C Workshop on the Ubiquitous Web. Position paper. (2006)
11. Reiter, E., Dale, R.: Building Natural Language Generation Systems. Cambridge University Press (2000)
12. Cardoso, N., Martins, B., Andrade, L., Silva, M.J.: The xldb group at geoclef 2005. (Technical report)
13. Seco, N., Veale, T., Hayes, J.: An intrinsic information content metric for semantic similarity in wordnet. In: ECAI. (2004) 1089–1090
14. Corley, C., Mihalcea, R.: Measuring the semantic similarity of texts. In: Proc. of the ACL Workshop on Empirical Modeling of Semantic Equivalence and Entailment, Association for Computational Linguistics (2005) 13–18