Suggestive Geo-Tagging Assistance for Geo-Collaboration Tools

Marius Austerschulte and Carsten Keßler

Institute for Geoinformatics, University of Münster, Germany, m.austerschulte@gmail.com, carsten.kessler@uni-muenster.de

Abstract

An argumentation map is an online discussion forum for spatially related topics that combines the forum with an interactive map. The utility of an argumentation mapping tool highly depends on the accuracy and quantity of the geo-tags that link the discussion contributions to geographic locations. These geo-tags can be created manually by the users of the argumentation map or automatically by a geo-parsing application. However, in the case of manual geo-tagging users often do not create geo-tags as extensively as desired. In contrast, automatic geo-parsers have difficulties with the informal language often used in user-generated content and with resolving small-scale features.

This paper proposes a hybrid approach for geo-tagging user-generated content which involves the users in the process but supports them with an automatic geo-parser which suggests locations. The implementation of a prototype as well as a human participants test are presented in order to analyze the geo-tagging performance of this approach. It turns out that it is possible to reduce the number of geo-tagging errors but keep the recall rate approximately constant, compared to automatic geo-parsers.

1 Introduction

With the increasing importance of the web as a medium for global communication, new applications for asynchronous discussions have been developed. Rinner (2001) describes the concept of *argumentation maps* to support map-based discussions in online planning. Argumentation maps combine an online discussion forum (ODF) with an interactive geographic mapping component. Their web-based design implicates several advantages compared to traditional offline means of discussion, such as a lower inhibition threshold for participation in a discussion (Kingston et al. 1999), the ability to attend a discussion from any computer with online access, or the possibility to share information among many attendees (Laurini 2004).

While conventional forums without a mapping component only allow the structuring and retrieval of contributions by keywords, argumentation maps in principle also provide the possibility to retrieve messages by the geographic area they are referring to. Discussion participants are instantly aware of what geographic area a specific discussion is about. Locations which would be difficult or laborious to explain in words (e.g., locations in uninhabited areas) can easily be marked in a map and ambiguous place names can be disambiguated with little effort.

Although the general concept of argumentation maps has been around for several years now, practical implementations are still relatively rare. Keßler (2004) created a prototypical discussion forum for spatial decision-making and geo-collaboration called "Argumap". Its successor is called "ArgooMap". Both combine an online discussion forum with an interactive map. Forum contributions can be *geo-tagged*, i.e., linked to geographic locations by clicking locations in the map.

A usability test was carried out by Sidlar and Rinner (2007). They attested the *Argumap* prototype a generally high usability but revealed several issues that could be improved. According to Rinner et al. (2008) users often do not create references to every location that appear in their contributions, or do not even set any references at all.

Another important limitation of both Argumap and ArgooMap is that entire contributions instead of single geographic names mentioned in the texts are linked to geographic locations (Rinner et al. 2008). This means that the match between place name and location is lost. The resulting complexity could lead to insufficient geo-tagging activity since the discussion participants cannot see a benefit from adding more references to the map.

In this paper we will present an argumentation mapping prototype that allows users to assign geo-tags to single terms instead of entire contributions.
With the term *user-generated text content* we denote text resources that are created on publicly accessible web sites by end-users and are not direct subjects to an editorial authority. Such text resources are sometimes written in colloquial or not well-authored language and may contain slang words and misspellings, especially in terms of upper and lower case. Important examples for web sites that are characterized by user-generated text content are online discussion forums, wikis, or blogs. Geo-tagging of text is often done *a posteriori*, i.e., it is not done by the authors of the texts themselves (which is *a priori* geo-tagging). A human or, in most cases, a software program tries to infer the intended geographic locations from the place names mentioned in the text and their context. This is in principle inexact in many cases, especially when performed by software programs that cannot understand the context in which a place name is mentioned. Automatic geo-tagging software[1], so-called "geo-parsers" or "geo-taggers", already achieves relatively good recognition rates on corpora containing well-authored texts, such as news stories (Silva et al. 2006). However, these geo-parsers have problems with resolving ambiguous place names. Locations that are not listed in the geo-parser's gazetteer, such as single buildings or locations in open land, cannot be geo-tagged at all[2]. User-generated text as encountered in discussion forums is often written in an informal style, contains typos and slang expressions and thus makes it even harder for geo-parsers to achieve high recognition rates. Therefore, a further motivation for this work is to find a method that allows effective geo-tagging of user-generated content. It should outperform automatic geo-parsers in this task.

2 Related Work

Online discussion forums[3] (ODF) are a well-established and popular application on the web. Besides the use of ODFs for personal informal discussions and information exchange, the concept also got into the focus of researchers who are investigating its benefits for professional

1 For example MetaCarta, http://www.metacarta.com

2 Except by entering the according geographical coordinates directly. However, this is cumbersome and not practically feasible for a large number of users.

3 One of the first online discussion forums, namely UBB.classic, emerged in 1996 (according to http://www.tomrell.com/

applications. The idea of discussing things by using ODFs is gaining increasing attention in areas like e-Learning (Wu 2004), public participation[4], and the theory of deliberative democracy (Wright 2007). Laurini (2004) depicts the advantages of online discussion forums in public participation. He states that web-based discussions do not need fixed appointments and are more convenient and relaxed than conventional participation procedures. For people who feel uncomfortable when speaking in front of large groups, online discussion forums are a good way of making themselves heard (Kingston 1999).

Argumentation mapping tools are based on the combination of an ODF and an online mapping component. Rinner (2001) defines the argumentation maps concept as an object-based model for geographically referenced discussions. Argumentation maps aim at "supporting any argumentative process that has a spatial component and can benefit from explicit links between arguments and the corresponding places they refer to" (Rinner et al. 2008, p. 6).

In the argumentation map model, Rinner (2006) distinguishes argumentation elements, geographic reference objects, and graphic reference objects (see Fig. 1). Argumentation elements are the formal representations of arguments expressed by the participants of a discussion. In a spatially related discussion these argumentation elements potentially refer to one or more geographic reference objects, which are part of the map. At the same time they refer to one or more graphic reference objects, which are created by the discussion participants. Graphic reference objects are markers in the map that highlight a point or an area. Between all three kinds of elements and objects several kinds of relations are defined that represent the structure of the corresponding discussion. For instance, multiple argumentation elements may be linked to each other through logical relations, which may indicate a response to a certain argument.

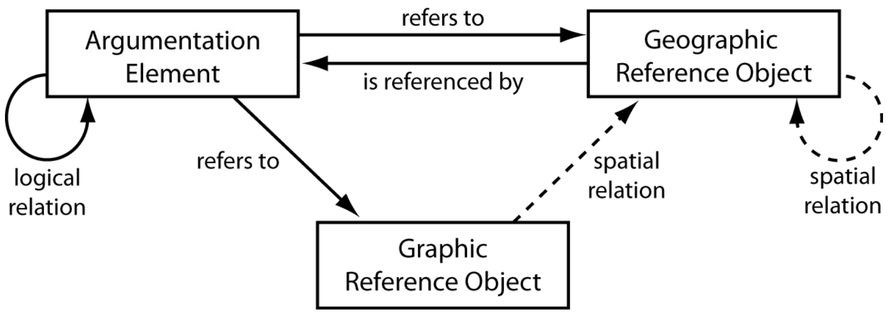

Fig. 1. Conceptual model for argumentation maps. Source: Rinner (2006).

[4] http://www.e-participation.net/taxonomy/term/32

3 Suggestive Geo-Tagging

The usefulness of an argumentation mapping tool highly depends on the accuracy and quantity of the geo-tags that link the discussion contributions to geographic locations. Generally we can assume the following: The more tedious it is for a user to create geo-tags and the less the user can see a benefit from these geo-tags, the less geo-tags will be created by her. Geo-tagging resources such as photographs and video is a common task on the web which can be completed easily by clicking a point in a map. In contrast, geo-tagging of texts generally requires more effort since text resources, unlike photographs, often refer to many different locations. Thus making this process easy and convenient is a key aspect of improving the overall quality of the geo-tags.

Two specific factors are crucial for a geo-parser's recognition rate on text content from online discussion forums: the authoring quality, including the grammatical and orthographic correctness of the given text, and the general prominence of the geographic features mentioned in this text.

Authoring quality of a text. Geo-parsers make use of Natural Language Processing and Named Entity Recognition to spot possible geographic identifiers and therefore require well-authored texts with only few grammatical mistakes and typos to achieve good results. In contrast, discussion forum contributions contain user-generated text content that is written spontaneously in many cases. Often, no great store is being set on grammatically and orthographically correct writing. As shown by Amitay et al. (2004), this negatively influences the recognition rate of geo-parsers to a great extent.

Prominence of a feature. The prominence of a feature is the general importance of a feature. Features with a high prominence, e.g., cities with a high population, are usually more likely to be mentioned than features with a low prominence. Geo-parsers utilize this assumption to disambiguate between equally-named features. However, spatial topics are often discussed at large scales (Rinner et al. 2008). Features at this scale, e.g., single buildings, that discussion participants might refer to, normally have a low prominence and are difficult for a geo-parser to recognize and resolve correctly.

To shift the geo-tagging process from an a posteriori to an a priori approach, the contribution authors have to be involved in the geo-tagging process to some extent. In the following we present a prototype of an argumentation mapping tool that supports the authors by suggesting locations based on the content of the contributions but leaves the final decision whether to create a new geo-tag in the hands of the authors (*suggestive geo-tagging*).

Geo-Parsing and Suggestive Geo-Tagging

A geo-parser tries to find terms in a text that denote location names and maps found entity names to the geographic coordinates of the intended real-world counterpart:

`Washington` → (38.895, -77.037)

Geo-parsing and suggestive geo-tagging are technically for the most part identical. However, while geo-parsing aims at geo-tagging a given text mainly automatically, a suggestive geo-tagger suggests features for names in a text. It generates a list of feature candidates, ranks them according to their computed relevance, and lets a human make the final decision of linking a geographic location to a name in the text or not. The key objective here is to support the users with a sufficient number of suggestions. We assume that a higher number of false positives (i.e. terms that are falsely regarded as geographic names) delivered by the geo-parser is tolerable, since the final geo-tagging decision is up to the users. However, too many false positives might confuse users and may therefore result in a lower number of geo-tags created by the users.

Disambiguation of Geographic Names

By far the biggest challenge that a geo-parser is confronted with is *ambiguity*: In many cases one word not only has exactly one meaning but is the name of different geographic features and denotes other completely different non-geographic things. Amitay et al. (2004) distinguish two kinds of ambiguities:

geo/non-geo ambiguity occurs if a word or a sequence of words is the name of a geographic place, but also has a different meaning that is not referring to any geographic place. For example, *Turkey* is the name of a country but also denotes a bird species.

geo/geo ambiguity occurs if there are several geographic places with the same name, e.g., *Paris* (France) and *Paris* (Texas).

The prototype presented in this paper utilizes its own geo-parser that has been implemented for the task of suggestive geo-tagging. This geo-parser is based on the Geonames gazetteer[5], whose database contains several types of features like cities, countries, mountains, lakes, parks, etc. In the following, the geo-parsers' disambiguation techniques are introduced.

In most cases humans can easily resolve ambiguity from linguistic and extra-linguistic context (Leidner et al. 2003). In contrast, geo-parsers do not truly understand a text. They try to resolve ambiguity by combining several different methods, rules, and heuristics (Overell et al. 2006; Amitay et al. 2004; Leidner 2004; Rauch et al. 2003; Blessing et al. 2007; Leidner et al. 2003). To make use of implicit context information, two special *minimality heuristics* can be applied (Gardent and Webber 2001) to disambiguate place names:

One sense per discourse. If one geographic place name is mentioned several times, it is assumed that it refers to the same location throughout the text (Gale et al. 1992).

Minimal spanning region defines interpretation. If there are more than one geographic names occurring in a text, the location candidates (interpretations) which span the smallest region are chosen (Leidner et al. 2003). For example, if the cities *Bedford* and *Everett* are mentioned, it is assumed that Bedford and Everett, Pennsylvania, are intended (since they only lie 10 km apart) and not Everett, Pennsylvania, and Bedford, UK.

The geo-parser makes use of several other methods used to disambiguate place names. These include Part-of-Speech tagging (information about the grammatical structure of the user-generated text), population data (places with higher population are more likely to be mentioned), and the "focus score". The focus score method calculates a score based on the map area that is currently visible in the prototype and an inverse distance weighting. The map area can be adjusted by the users independently and should specify the approximate geographic scope of the discussion. Feature suggestions lying inside this area are assigned a higher score than those lying outside.

4 Design and Implementation of the Prototype

In order to overcome the shortcomings of the previous versions of Argoomap the prototype has been redesigned, away from a contribution-based towards a word-based geo-tagging mechanism. This should clarify

5 http://www.geonames.org/

the relations between the words in the contributions and the map and increase the benefit the users gain from adding new geo-tags.
The new prototype is called “ArgooMap 2”. It is a web page based on standard web technologies (HTML, CSS, JavaScript) that can be used with any common web browser. No additional plug-ins or software installations are required to run it.

The prototype layout is horizontally divided into two panels (see Fig. 2). The left panel contains the textual content where forum contributions can be displayed and new topics and replies can be written. The right panel is dedicated to an interactive map that can be dynamically panned and zoomed. Users can choose between four different map types: a street map, satellite imagery, satellite imagery with an overlaying street map, or OpenStreetMap data. ArgooMap 2 makes use of the Google Maps API[6]. Due to its prominence, many users may already be familiar with the map controls.

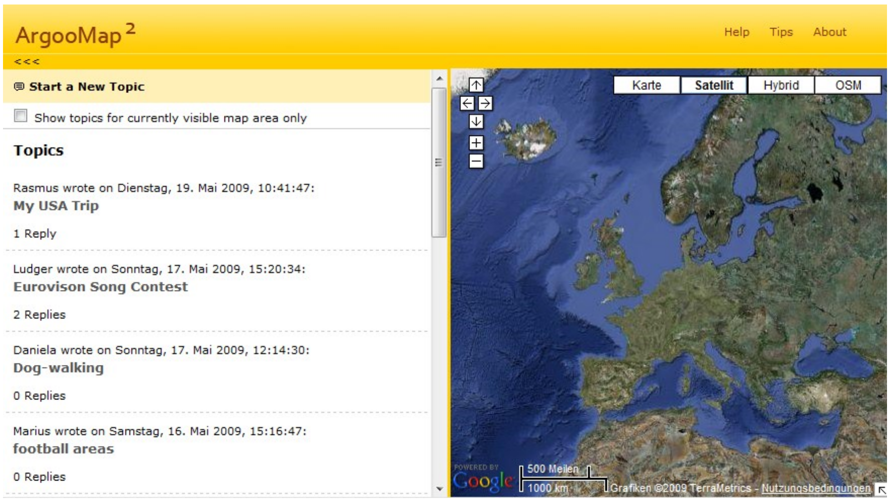

Fig. 2. Front page of the ArgooMap 2 prototype website.

The left panel lists the posted forum topics in chronological order. As the map is panned and zoomed to a different geographic area the list of topics is dynamically updated with posts referring to the current map extent.

Clicking the header of a topic brings up the according discussion including all the replies that have been posted. It automatically adjusts the map pane so that all geographic locations mentioned in the discussion are shown. Words that have been linked to geographic locations in the map are

6 See http://code.google.com/apis/maps/documentation/reference.html

highlighted in blue. The associated points are flagged with blue markers in the map. When the mouse is moved over a highlighted word in the text, the corresponding markers in the map change their color to red. The other way around, when a marker in the map is clicked, an info window opens that lists all words in the discussion that refer to this marker. Moving the mouse over an entry in the info window immediately highlights the associated word in the discussion.

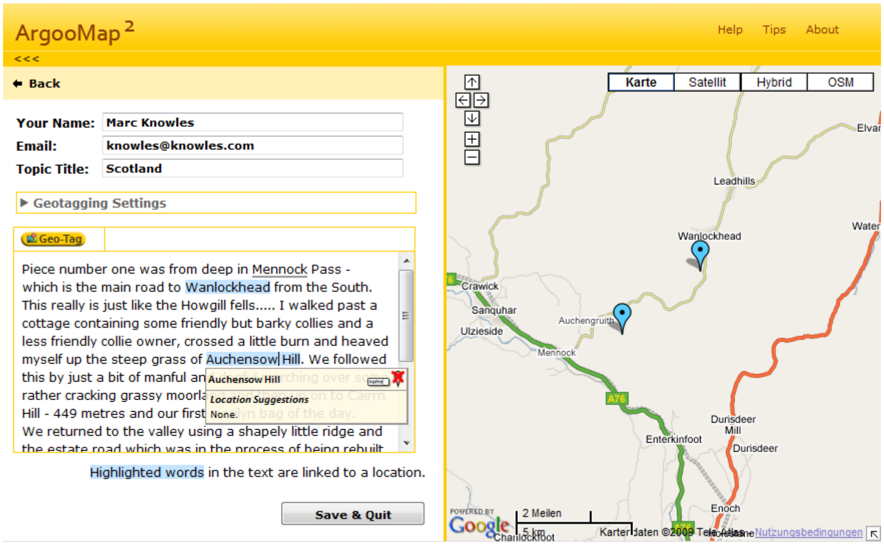

Fig. 3. Composing a new contribution. The two terms in the text highlighted in blue are each linked to one of the two blue markers in the map.

There are two different possibilities to link words or terms in the contribution text to one or more geographic locations: automatic geo-tagging and manual geo-tagging. The easiest way is to start the automatic geo-tagging process by clicking the “Geo-Tag“ button. Words that are recognized as geographic names are underlined. Clicking such a word brings up the location suggestions that the author can choose from.

If none of the automatically retrieved suggestions matches the intended location, or if a word has not been recognized as a geographic name at all it can still be geo-tagged manually. This is done by simply clicking the according location in the map. Until the contribution has been saved the authors can create, edit, and remove references at any time.

5 Human Participants Test

To analyze the applicability of the prototype for geo-tagging user-generated content, a group of people was asked to take part in an experimental online discussion by using ArgooMap 2. The achieved geo-tagging effectiveness of ArgooMap 2 has been compared to the performance of the automatic geo-parsers Yahoo! Placemaker and MetaCarta GeoTagger.

Preparation of the Test

The performance of automatic geo-parsers is usually evaluated on large annotated corpora taken from newspapers, Wikipedia (Overell et al. 2006), or different web pages (Amitay et al. 2004). Such corpora only have to be annotated once and can then be reused as a basis for automated testing of a geo-parser. However, this is only partially possible with suggestive geo-tagging since the content to geo-tag is created by the participants during the test[7]. Therefore, each time after carrying out a test, the newly generated content has to be searched for geographic references manually and then annotated with the corresponding geographic locations. This can be very time-consuming.

The test discussion was not fixed to a certain topic. Due to diverse thematic interests of the participants, restricting the discussion to a specific topic would have excluded a number of people from taking part. Instead, some initial questions were posted in the forum acting as conversation starters.

Execution of the Test

For the test, 41 people were invited via email to participate. The email explained the objective of the test and the capabilities of the prototype but did not *prompt* the readers to make use of the geo-tagging function. The participants should not have the feeling of having to place geo-tags extensively for the purpose of this test. In contrast, it was left up to them to find out the advantages of well geo-tagged content.
When using the prototype the first time the participants were presented an introductory video of approximately four minutes length. This video demonstrated the basic concepts and the operation of the program.

---

7 Nevertheless, the geo-parser integrated in ArgooMap 2 could be trained on an annotated corpus containing user-generated text content.

Although the discussion started sluggishly it became more vital with an increasing number of contributions. Travel reports emerged as a favorite topic. To keep the discussion going, additional contributions and replies were occasionally posted by the moderator.

Evaluation

Table 1 depicts the participation statistics for the entire test. Until the end of the testing period, 20 of the 41 invited persons (49%) had written at least one contribution. The total number of contributions (excluding those that were created by the author) was 33 which corresponds to an average number of 1.7 contributions per participant. There were 19 threads with a total number of 17 replies, i.e., approx. 1 reply per thread on average.

Table 1. Discussion participation statistics. Note: Contributions posted by the researchers are not included.

Number of invitations	41
Number of participants	20
Participation rate	49%
Number of contributions	33
Average number of contributions per person	1.7
Number of threads	19
Total number of replies	17

To determine the geo-tagging performance, all contributions had to be scanned for geographic references manually and then annotated with the corresponding geographic locations. There were some cases where the correct annotation was not completely clear. For instance, one thread was about getting from Münster to Coventry and several airplane routes were discussed. The participants provided advice such as "fly from Dortmund to Stansted". In these cases it was defined that the locations of the corresponding airports were referenced and not the cities of Dortmund and Stansted themselves. Other ambiguities frequently encountered referred to the correct tagging of multi-word units like "Münster central station" or "Cologne cathedral". Here often only the city name ("Münster") was geo-tagged. In this case it was defined that the correct tagging applies to the full multi-word unit "Münster central station". Phrases like "Dresden, Germany" (where "Germany" acts as a specifier) were defined to be treated as references to one single location ("Dresden") and not separately ("Dresden" and "Germany").

The annotation results are summarized in Table 2. 192 geographic references have been made in the contributions which refer to 152 distinct places. This means that there were 4.7 distinct places mentioned per contribution on average. This number is relatively high compared to the two discussions from the case study that was carried by Rinner et al. (2008) (2.94 and 3.13, respectively).

Table 2. Geographic references statistics.

Total number of geographic references	192
Number of referenced places	156
Contributions without geographic references	5 (15%)
Contributions with one geographic reference	4 (12%)
Contributions with multiple geographic references	24 (73%)
Average number of geographic references per contribution	5.8
Average number of referenced places per contribution	4.7

Evaluation Measures

In information retrieval, the most frequently used measures for the effectiveness of a system are recall and precision (Manning et al. 2008). These measures can easily be adapted for evaluating a geo-tagging application:

$$R = \frac{\#(\text{correct geo-tags})}{\#(\text{geographic references})}$$

$$P = \frac{\#(\text{correct geo-tags})}{\#(\text{created geo-tags})}$$

where R and P denote recall and precision, respectively.
Depending on the application, one of the two measures might be more important than the other. As a trade off, the *F measure* is applied (Manning et al. 2008):

$$F = \frac{1}{\alpha \frac{1}{P} + (1-\alpha)\frac{1}{R}} = \frac{(\beta^2+1)PR}{\beta^2 P + R} \text{ where } \beta^2 = \frac{1-\alpha}{\alpha}$$

where $\alpha \in [0,1]$ and $\beta \in [0,\infty]$

With $\beta = 1$ (written $F_{\beta = 1}$), recall and precision are equally weighted. If $\beta < 1$ precision is emphasized whereas values of $\beta > 1$ emphasize recall. In the case of maps more emphasis should be put on precision, since incorrect geo-tags might easily confuse people during a discussion and might adulterate the geographic scope. A value of 0.5 for β is chosen here. It is considered to accentuate precision and decrease the influence of recall in a reasonable magnitude. Accordingly, we will also look at $F_{\beta = 0.5}$, besides recall and precision.

Geo-Parsers

The ArgooMap 2 prototype is evaluated against the two automatic geo-parsers Yahoo! Placemaker[8] and MetaCarta GeoTagger[9]. These are freely available geo-parsing web services which are accessed via HTTP POST. After posting the text content the services return an XML document containing the found places and the corresponding geographic coordinates as well as information on where in the text the geographic names have been found.

Test Results

To determine recall and precision of ArgooMap 2, Yahoo! Placemaker, and MetaCarta GeoTagger, for each contribution the correct geo-tags, the geo/geo tagging errors, the geo/non-geo tagging errors, and the geographic references that were not geo-tagged were counted. Table 3 summarizes the results achieved by each of the three applications.

From the 192 geographic references mentioned in all contributions, 91 where geo-tagged correctly with ArgooMap 2. There were 13 geo/geo tagging errors and 88 geographic references were not geo-tagged. Geo/non-geo tagging errors were not encountered. About one third of the geo-tags was created manually by clicking a location in the map and two thirds were inserted after selecting a feature suggestion. It is remarkable that all of the 13 geo/geo tagging errors occurred after selecting a suggestion but none after marking a location manually in the map. Although ArgooMap 2 provides the possibility to reference multiple locations to single geo-tags, this function was not utilized by any of the participants.

---

8 http://developer.yahoo.com/geo/placemaker

9 http://ondemand.metacarta.com/?method=GeoTagger

The number of geographic references tagged correctly by Yahoo! Placemaker was 90, and therefore slightly lower than the value achieved by ArgooMap 2. However, the number of tagging errors was more than twice as high (27), most of them geo/geo errors and only 2 geo/non-geo errors.

The MetaCarta geo-parser achieved the highest number of correct geo-tags (105) but also made the highest number of mistakes (34). Thus, with 54.7% the MetaCarta geo-parser attained by far the highest recall value, followed by ArgooMap 2 with 47.4% and Yahoo! Placemaker with 46.7%. However, this is achieved at the price of the highest error rate and the lowest precision. ArgooMap 2 clearly gained the highest precision value, namely 87.5%. Yahoo! Placemaker's precision is 10.6 percentage points lower (76.9%), followed by MetaCarta having a precision of 75.5%.

Table 3. Geo-tagging effectiveness of ArgooMap 2 compared to Yahoo! Placemaker and MetaCarta.

	ArgooMap 2	Yahoo!	MetaCarta
Correct	91	90	105
Geo/geo errors	13	25	31
Geo/non-geo errors	0	2	3
Not geo-tagged	88	77	56
Recall	47.4%	46.9%	54.7%
Precision	87.5%	76.9%	75.5%
$F_{\beta=0.5}$	0.75	0.68	0.70

Table 4. ArgooMap 2 specific geo-tagging effectiveness statistics.

Suggested geo-tags	71 (68.3%)
Manual geo-tags	33 (31.7%)
Geo/geo errors in suggested tags	13 (100%)
Geo/geo errors in manual tags	0 (0%)
Geo-tags referring to multiple locations	0

Table 4 shows some additional ArgooMap 2 specific statistics. More than two thirds of the geo-tags created in ArgooMap 2 were created by selecting an automatically generated location suggestion. Manual geo-tags were mainly created for places for which no appropriate suggestion was available. The corresponding error rates show that all geo-tagging errors occurred together with suggested tags. In contrast, the manually created geo-tags were correct in all cases.

6 Discussion

The experimental discussion has shown that the recall rate of the implemented prototype was slightly higher than the one achieved by Yahoo! Placemaker, but did not reach the recall rate of the MetaCarta geo-parser. On the other hand only every eighth geographic reference was tagged falsely by the users of the ArgooMap 2 prototype, while MetaCarta and Yahoo! Placemaker tagged about one quarter of the geographic references incorrectly. Hence, the geo-tagging precision of ArgooMap 2 was clearly the highest among all geo-taggers. However, the prototype's recall rate is not completely satisfying. Some issues became apparent during the human participants test that negatively influence the recall rate. These issues as well as possible solutions are discussed in the following.

The task of geo-tagging was well understood by the participants. There were only 3 contributions that did contain geographic references, but no geo-tags. In such cases, the reason was mostly that the geo-parser did not provide any suggestions for the mentioned geographic names. Although it was possible to geo-tag these names manually, the participants did not, either to avoid extra work or because they were not aware of this possibility. As the feature of manually creating geo-tags is a key functionality of the prototype which also separates it from fully automatic geo-parsers, its usage should be made clearer in the user interface. This could be achieved, by adding a prominent button that explicitly provides a simple option to manually create geo-tags. Additionally, the introduction presented to first time users of ArgooMap 2 should put more emphasis on this feature.

The human participants test has shown that the majority (68%) of the referenced geographic names were geo-tagged by confirming a suggestion made by the built-in geo-parser. This illustrates the importance of the geo-parser's performance for achieving a high recall. However, especially two shortcomings of the geo-parser had a negative impact on the recall during the test. First, generating location suggestions for a given text may take, depending on the text length, up to 20 seconds. This is clearly too long and it may discourage people to utilize this functionality. Hence, this fact very likely leads to a decrease in recall. Special focus should therefore be put on accelerating this process. The second drawback concerns the different names of features in various languages. The ArgooMap 2 geo-parser gazetteer only contains the local names of geographic features, e.g., it contains *Köln* but not *Cologne*, and it contains *Roma* but not *Rome*. As a result, mentions of *Cologne* were not recognized by the geo-parser as geographic names, not highlighted, and for this reason also not geo-tagged by some users.

In many contributions, specific geographic features were mentioned several times. Even if there are correct suggestions available for each instance of this name, users clearly tended to reference only one instance of this name. For example, one user mentioned “Coventry“ four times but only geo-tagged the first occurrence of it, presumably for convenience. Here the *one sense per discourse* heuristic could be applied to relieve the user from the work of confirming the same suggestion for every single geographic name instance. When one instance of such a name is geo-tagged by the user, he should be asked whether all remaining occurrences of this name should be automatically referenced as well.

A problem that was encountered several times was that participants forgot to adjust the map area of interest before initiating the geo-parsing process. Hence, places lying further outside the map area were not recognized. Since the adjustment of the map area by the participants turned out to be an important means for disambiguation, it should be preserved as such. However, a simple solution to this problem might be to just remind the users of adjusting the map area each time they start the automatic geo-parsing process. A balance has to be struck here in order to not annoy the users unnecessarily.

Transfer to Other Application Areas

The concept of suggestive geo-tagging is not limited to the field of argumentation maps and online discussion forums. Potential other application areas include all kinds of public platforms that deal with user-generated text content, such as wikis (Wikipedia) or blogs. Figure 4 shows a workflow diagram that abstracts the general principle of suggestive geo-tagging.

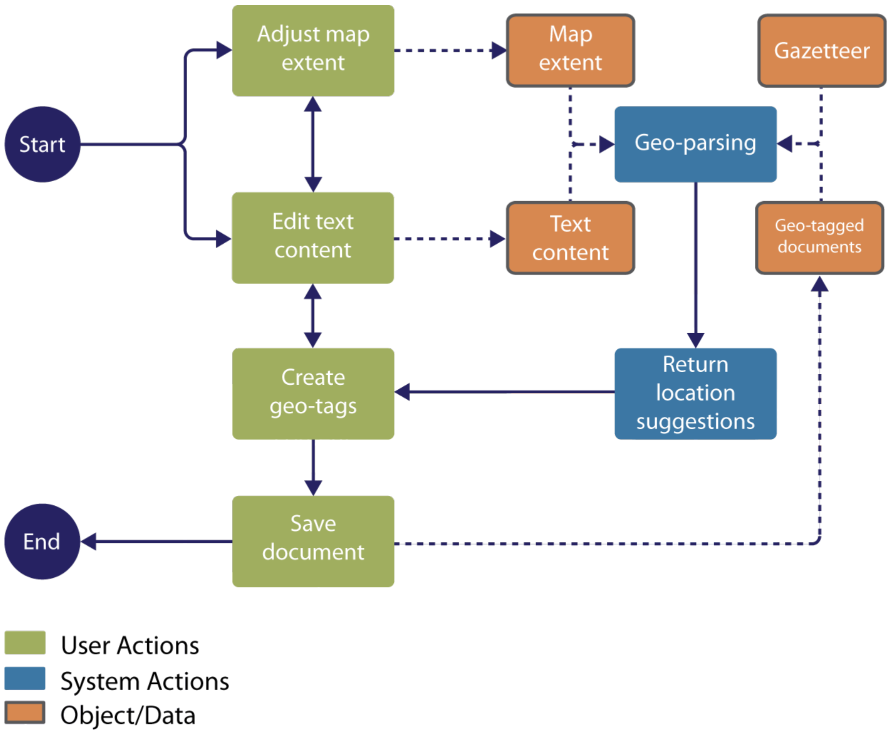

Fig. 4. Suggestive geo-tagging workflow diagram.

7 Conclusion

This work has shown that geo-tagging user-generated text content can be done with acceptable recall and with an especially high precision if the geo-tagging is delegated to the users, i.e., is done a priori. An important prerequisite is a geo-tagging software that actively suggests locations. It should support the users in the geo-tagging work as much as possible to achieve a better recall rate. On the other hand it should still give them full control over what is being geo-tagged in order to keep the precision rate high.

Future Work

In the medium term it should be considered whether the geo-parser implemented for ArgooMap 2 can be replaced by the MetaCarta

GeoTagger web service, due to its good geo-tagging performance and its high speed. The service can be used as a suggestive geo-parser, too, since it returns not only one location for a recognized geo-term (like, for instance, Yahoo! Placemaker does), but several suggestions, if available. These are weighted by calculated *confidence* values, analogously to the score values in ArgooMap 2. However, the reason why this has not been done yet is the fact that the MetaCarta results do not inhere information about the actual names of the suggested features and their according administrative regions, but solely information on their feature types and their locations. Therefore, it can be difficult to figure out which suggested location actually refers to the intended feature.

Currently both Yahoo! Placemaker and MetaCarta GeoTagger use proprietary formats to access their geo-parsing services. It would be desirable to have a standardized way of initiating such services, preferably through an OGC compliant Web Processing Service[10] (WPS) interface. In this case the suggestive geo-tagging application could be implemented independently of the associated geo-parser.

Geographic Scope

One of the most important advantages of geo-tagged documents is the possibility to retrieve them according to geographic criteria. However, not all referenced geographic locations can be considered equally important for the subject matter of the document. Therefore, to figure out the most relevant geo-tags and thus the most relevant documents for a specific region, the *geographic scope* of a document is attempted to be determined. The geographic scope is defined, if it exists, as the region where more people than average would find that document relevant (Silva et al. 2006). There are different approaches of how to calculate the geographic scope (Martins et al. 2005; Silva et al. 2006). Amitay et al. (2004) propose the *focus scoring algorithm* to calculate a geographic scope (they call it *page focus*) based on the existing geo-tags in the document. Silva et al. (2006) describe a number of heuristics that they rely on for their approach to compute a geographic scope for web pages.

In the context of suggestive geo-tagging it should be investigated how the geographic scope can be inferred from user-generated content that has been geo-tagged with the help of the ArgooMap 2 prototype. The human participants test posed the assumption that, albeit the recall rate is moderate, users tended to geo-tag mostly those geographic names that they thought were most important for the subject matter of their contribution.

10 OGC Web Processing Service, http://www.opengeospatial.org/standards/wps

Furthermore, the map extent that was set at the time of geo-tagging might prove useful at this point.

Acknowledgments

This work has been funded by the SimCat project granted by the German Research Foundation (DFG Ra1062/2-1 and Ja1709/2-2, see http://sim-dl.sourceforge.net). Special thanks go to Claus Rinner and Werner Kuhn for valuable comments.

References

Amitay E, Har'el N, Sivan R, Soffer A (2004). Web-a-where: Geotagging web content. In: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval, ACM Press, pp 273-280

Blessing A, Kuntz R, Schütze H (2007). Towards a context model driven German geo-tagging system. In: GIR '07: Proceedings of the 4th ACM workshop on Geographical information retrieval, ACM, pp 25-30

Chaves M, Silva M J, Martins B (2005), A Geographic Knowledge Base for Semantic Web Applications. In: Proceedings of SBBD-05, the 20th Brazilian Symposium on Databases, UFU, pp 40-54

Gale W, Church K, Yarowsky D (1992). One sense per discourse. In: Proceedings of the workshop on Speech and Natural Language, HLT '91, Association for Computational Linguistics, pp 233-237

Gardent C, Webber B (2001). Towards the Use of Automated Reasoning in Discourse Disambiguation. In: Journal of Logic, Lang. and Inf. 10(4): 487-509, Kluwer Academic Publishers, Hingham, USA

Keßler C (2004), Design and Implementation of Argumentation Maps, Diploma thesis, Westfälische Wilhelms-Universität Münster, Germany.

Kingston R, Carver S, Evans A, Turton I (1999), A GIS for the Public: Enhancing Participation in Local Decision Making. GIS Research UK.

Laurini R (2004), Computer Systems for Public Participation. Laboratoire d'Ingénierie des Systèmes d'Information, University of Lyon, France. http://www.gisig.it/VPC_sommet/CD_Sommet/ws3/articololaurini.pdf

Leidner J, Sinclair G, Webber B (2003). Grounding spatial named entities for information extraction and question answering. In: Proceedings of the HLT-NAACL 2003 workshop on Analysis of geographic references, Association for Computational Linguistics, pp 31-38

Leidner J (2004). Toponym Resolution in Text: Which Sheffield is it? In: Proceedings of the SIGIR 2004 conference on Research and Development in Information Retrieval, Sheffield, UK, July 25th -29th 2004. ACM Press, New York, USA

Mahemoff M (2006), Ajax Design Patterns, O'Reilly Media, Sebastopol, USA

Manning C, Raghavan P, Schütze H (2008). Introduction to Information Retrieval, Cambridge University Press, New York, USA

Martins B, Chaves M, Silva M J (2005), Assigning Geographical Scopes To Web Pages. In: Advances in Information Retrieval, Springer, pp 564-567

Overell S, Magalhaes J, Ruger S (2006). Place disambiguation with co-occurrence models. In: A. Nardi, C. Peters, and J. L. Vicedo, editors, CLEF 2006 Workshop, Working notes, September 2006

Rauch E, Bukatin M, Baker K (2003). A confidence-based framework for disambiguating geographic terms. In: Proceedings of the HLT-NAACL 2003 workshop on Analysis of geographic references, Association for Computational Linguistics, pp 50-54

Rinner C (2001), Argumentation maps: GIS-based discussion support for on-line planning. Environment and Planning B: Planning and Design 30(6): 847-863

Rinner C (2006). Argumentation Mapping in Collaborative Spatial Decision Making. In: S. Dragicevic, S. Balram, Collaborative GIS, Idea Group Publishing, pp 85-102.

Rinner C, Keßler C, Andrulis S (2008). The use of Web 2.0 concepts to support deliberation in spatial decision-making. In: Computers, Environment and Urban Systems 32(5): 386-395

Sidlar C, Rinner C (2007), Analyzing the Usability of an Argumentation Map as a Participatory Spatial Decision Support Tool. URISA Journal 19(1): 47-55

Silva M J, Martins B, Chaves M, Afonso A P, Cardoso N (2006). Adding Geographic Scopes to Web Resources. In: Computers, Environment and Urban Systems 30(4): 378-399

Wright S, Street J (2007). Democracy, deliberation and design: the case of online discussion forums. In: New Media Society 9(5): 849-869

Wu D, Hiltz S R (2004). Predicting Learning from Asynchronous Online Discussions. In: JALN 8(2), April 2004