

Understanding the Types of Information Humans Associate with Geographic Objects

Ahmet Aker
University of Sheffield
a.aker@dcs.shef.ac.uk

Robert Gaizauskas
University of Sheffield
r.gaizauskas@sheffield.ac.uk

ABSTRACT

In this paper we investigate what sorts of information humans request about geographical objects of the same type. For example, *Edinburgh Castle* and the *Bodiam Castle* are two objects of the same type – *castle*. The question is whether specific information is requested for the object type *castle* and how this information differs for objects of other types, e.g. *church*, *museum* or *lake*. We aim to answer this question using an online survey. In the survey we showed 184 participants 200 images pertaining to urban and rural objects and asked them to write questions for which they would like to know the answers when seeing those objects. Our analysis of 7644 questions collected in the survey shows that humans have shared ideas of what to ask about geographical objects. When the object types resemble each other (e.g. *church*, *temple*) the requested information is similar for the objects of these types. Otherwise, the information is specific to an object type. Our results can guide tasks involving automatic generation of templates for image descriptions, and their assessment, as well as image indexing and organization.

Categories and Subject Descriptors

IR [Language technologies for IR]: [NLP, IE, Summarization, QA]; IR [Evaluation, Test collections, Crowdsourcing for IR]: [Mechanical Turk]

General Terms

NLP, Summarization, Image Description Generation

Keywords

Geographical object description

1. INTRODUCTION

In every day life we see and categorize things in our built or natural environment. For instance, if we look at different churches, we see that each of them has different style, different look, different size, some of them are newer than the

others, etc., but still we are able to categorize all of them as churches. For categorization purposes visual attributes such as the style and the size of a church might be enough; however, to report about each church separately we need more attributes whose values make each church distinguishable from others.

Knowing what sets of attributes people use to describe geographical objects has several applications. Text summarization employed to automatically generate descriptions of geographical objects can be improved by incorporating human preferences into the output summaries. Template based summarization methods, for example, have been used successfully for news event summarization [2, 3]. The same strategy can be applied to the task of generation of descriptions for geographical objects, e.g. for purposes of generating descriptions for images of these objects. Templates can be derived from the set of attributes relevant to humans to bias the summarization system towards the text units containing the values for these attributes. Furthermore, the attributes could be used in a guided summarization task as organized by the Document Understanding Conferences (DUC)¹ and the Text Analysis Conference (TAC).² Additionally, the attributes can be used to evaluate the output of an automatic summarization system by testing whether highly relevant attributes are covered in the output summary. Furthermore, the attributes supply a guide for what to index for images pertaining to locations. If information highly relevant for humans is used for indexing, this could lead to better retrieval and organization of those images.

For such applications it is relevant to know whether there is a general set of attributes that people use to describe any geographical object. Or, given that humans categorize objects into types (e.g. church, museum, etc.), are the sets of attributes specific to single object types, or perhaps shared between types that are subtypes of a common higher level type (e.g. church and temple are both religious buildings)?

In this paper we aim to address these questions. Specifically, we ask: For what attributes related to geographic objects would people like to know the values when seeing such an object? And: Is this set of attributes specific to a particular object type (e.g. church) or shared between different object types?

We aim to answer these questions using an online survey conducted on Mechanical Turk. In the experiment we show the participants different images of objects from around the world. Our set of images contains only images of static fea-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CIKM'11, October 24–28, 2011, Glasgow, Scotland, UK.
Copyright 2011 ACM 978-1-4503-0717-8/11/10 ...\$10.00.

¹<http://duc.nist.gov/>

²<http://www.nist.gov/tac/>

Table 1: Objects types used in the experiment.

Urban types	Rural types
abbey, aquarium, avenue, basilica, boulevard, building, cathedral, cemetery, church, gallery, house, monument, museum, opera house, palace, parliament, prison, railway, railway station, residence, square, stadium, temple, tower, university, zoo	beach, canal, cave, garden, glacier, island, lake, mountain, park, peak, river, ski resort, village, volcano

Table 2: Example attributes with related questions.

location	where is garwood glacier?, where exactly is edmonton?, where it's located?
foundationyear	when was it build?, which year was this zoo opened?, when it was established?

tures of the built or natural landscape, i.e. objects with persistent geo-coordinates, such as buildings and mountains, and not images of objects which move about in such landscapes, e.g. people, cars, etc. We ask the participants to write questions for which they would like to know the answers when seeing each image. We collected 7644 questions from 184 participants. The results of our analysis reveal that people share ideas as to what to ask about geographical objects in general. When the object types resemble each other, the requested information is similar. Otherwise, this information is specific to a specific type.

In this paper we first present the experimental setting of our survey and outline the preprocessing process of the data (Section 2). In Section 3 we present our analysis and report and discuss the results. We conclude the paper in Section 4.

2. EXPERIMENTAL SETTING

2.1 Experiment

In the experiment we showed the participants an image pertaining to a particular object or place. We also presented the participants with the name of the place (e.g. *Eiffel Tower*) and its object type (e.g. *tower*). The participants were asked to take the role of a tourist and provide ten questions for which they would like to know the answers when they see the place shown in the image. We showed images picturing places of 40 different object types randomly chosen from our entire set of 107 object types. From these 40 object types 25 are urban and 15 rural types (see Table 1). For each object type five different places were shown. For example, for the object type *tower* the images of *Eiffel Tower*, *Flag Tower of Hanoi*, *BT Tower*, *Munntoren*, *Bettisons Folly* were shown. These places (towers) were manually selected from Wikipedia. Each image was shown to five different participants. We ran the experiment for four weeks. In total we collected 7644 questions for 187 different places. The questions came from 184 different participants. The expected number of questions was 10.000 (40 types \times 5 objects \times 5 participants \times 10 questions). However, there are some objects for which we only have questions from two or three workers. Most objects have 40 questions. The number of questions for the urban types is 4815 and for the rural ones 2829. We used Mechanical Turk for our experiment.

2.2 Question Preprocessing

We manually analyzed all questions in order to assess quality. Approximately 2% of the questions were empty because not all the workers wrote 10 questions for each object.

Moreover, some questions are only related to the image itself rather than to the place shown in the image (e.g. “*when the picture is taken?*”, “*how many flowers you found in the image?*”, “*is there a bus in the picture?*”). In addition to these, some questions present non-resolved references so it is impossible to know what object they refer to (e.g. “*what language do they speak?*”). Finally, there are questions which bear no relation at all with the object in the image (e.g. “*how is the manager?*”). These questions do not address the task, which is to ask questions about the *object* shown in the image, not about the image itself or related information. Therefore, we categorized all these questions as noise which makes up 19% (1479 out of 7644) of the entire question set.

We categorized the remaining 81% of the questions (6169 out of 7644) by the *attribute* the worker was seeking the value for with his/her question. An attribute is an abstract grouping of similar questions. We regard two or more questions as similar if their answers refer to the same information type. For instance, we regard the questions “*where is garwood glacier?*” and “*where exactly is edmonton?*” as similar because both aim for answers related to the information type *location*. We name the attribute according to the information type it refers to (e.g. *location*). Table 2 shows examples attributes with questions.

In total we identified 146 attributes, however, 95 of them contain less than five questions, so we ignore these attributes in further analysis. We analyse the remaining set of 51 attributes (see Table 3), each of which has at least five questions related to.

3. ANALYSIS

3.1 Is there a set of attributes that people generally associate with geographic objects?

In Table 3 the 51 attributes are given. The left column of the table shows the very frequent addressed attributes (*top ten*). More than 65% of the questions can be categorized by these ten attributes. This means that people do share ideas as to what types of information are required about a place, and the set of *top ten* attributes captures the majority of these information types. However, although these attributes occur for many object types, their *popularity* is not the same for all the object types with which they are associated.

We define the popularity of an attribute for an object type as the number of questions categorized under this attribute for that particular object type. Attribute popularity indicates how important the type of information represented by the attribute is for the particular object type.

For instance, the attribute *visiting* is the most popular attribute (has the maximum number of 52 questions) in the object type *museum* (see Table 4). This indicates that it is most important for people to know how much the entry to the museum costs or when the museum opens. This information is more relevant than, e.g. knowing when the museum was built. However, if we look at the object type *house*, we can see that the same attribute *visiting* is not the most popular one. It occurs at position five with only 11 questions. For this object type people seem to be interested most in knowing the *design* of the house and less in information related to visiting.

The remaining 35% of the questions are spread over the remaining 41 attributes as shown in the right column (*below top ten*) of Table 3. From this list we can observe that there

Table 3: Attributes used to categorize the questions.

Top Ten	Below Top Ten
visiting, location, foundationyear, surrounding, features, history, size, design, mostattraction, naming	purpose, istouristattraction, height, visitors, founder, owner, events, type, status, habitants, designer, constructioninfo, temperature, length, eruptioninfo, comparison, capacity, depth, maintainance, arts, studentsinfo, subjectsoffered, gravesinfo, preacher, firstdiscovery, parkinginfo, religioninfo, destination, origin, workers, width, travellers, travelcost, waterinfo, travelling, firstclimber, tributaries, prayertime, personinmonument, snowinformation, first-founder

Table 4: List of attributes (eight most frequent) for example object types. The numbers within the brackets attached to each attribute indicate how many questions were categorized by that attribute.

Object Type	Attributes
cathedral	location(30), visiting(23), foundationyear(16), history(12), founder(9), size(8), design(8), preacher(6)
house	design(27), location(22), foundationyear(16), owner(13), visiting(11), history(10), istouristattraction(9), founder(9)
museum	visiting(52), location(22), foundationyear(22), mostattraction(10), features(9), history(8), arts(5), owner(5)
opera house	visiting(33), foundationyear(20), location(17), events(17), mostattraction(10), surrounding(10), designer(9), owner(8)
volcano	eruptioninfo(42), visiting(26), location(21), surrounding(20), height(15), status(14), foundationyear(7), naming(6)
zoo	visiting(41), features(34), location(19), foundationyear(14), size(10), mostattraction(8), surrounding(5), owner(5)

are some specific attributes which are likely to be present only in a specific object type or in very few object types. For instance, as shown in Table 4 the attribute *eruptioninfo* is associated only with the object type *volcano*. This attribute contains questions related to the eruption of different volcanos. From Table 4 we can see that this attribute is also the most popular for the *volcano* object type, while information related to *visiting*, *location*, etc., which is generally most frequently asked for, comes after *eruptioninfo*. From this we can conclude that even though object types share the *top ten* attributes, the differences in attribute popularity indicate that these are not equally important for all object types. Therefore, the question arises as to whether each object type has an associated specific set of attributes, or whether the attributes relevant to it are shared between several object types. This is our second research question.

3.2 Is each set of attributes specific to a particular object type or shared between several object types?

To address this question we compare different object types and investigate the degree of similarity between them. We assume, if two object types are similar, then they will not only share a set of attributes, but the attributes within this set will also have similar popularity ranking.

We use Kendall’s Tau rank correlation coefficient as a metric which indicates similarity between object types, while considering the attribute popularity ranking in their attribute sets. Kendall’s Tau correlation coefficient is equal or close to 1 if two events are highly correlated in rankings and close to 0 if there exists little or no correlation between the ranks of the two events.

The attributes are ranked according to their popularity, i.e. the number of questions contained under each attribute. If attributes of two different object types have similar rankings, this is an indication that these attributes are of similar importance for both object types. Kendall’s Tau for these object types will return a high correlation for the pair, and we will refer to such object types as similar. A difference in attribute rankings between different object types indicates that these attribute sets are more specific to the one object

type and not shared by the other(s). In this case Kendall’s Tau will return a low correlation, and we will refer to such object types as dissimilar.

In our analysis we report comparisons based on three different sets of attributes: *all*, *top ten* and *below top ten*. *All* is the entire set of attributes shown in Table 3 (attributes from columns *top ten* and *below top ten*). The *top ten* attributes include the top ten attributes which are shared between most of the object types and thus more likely to render similarity. The *below top ten* set contains all remaining attributes which are as we assume above rather object type specific.

3.2.1 Results

If *all* attributes are taken, there is in general a correlation between object types in both urban and rural categories. On average the urban object types correlate with *0.55* (median *0.57*) and the rural types with *0.53* (median *0.53*). The mean and median correlation coefficient lie close to each other, which indicates that the distribution of highly correlating object types is similar to that of object types with low correlation. To analyse this result further, we aim to identify the object types whose correlation coefficient is higher than the mean correlation and those which are correlated with a coefficient lower than the mean. We also aim to understand whether there are object types which are only highly correlated when top ten most frequent attributes are considered and whose correlation potentially drops for the remaining set of attributes.

Table 5 shows three groups of object type pairs in rural and urban areas. In the first column the object type pairs whose correlation coefficient is higher than the mean for *all* attributes are shown. The second column shows the object type pairs whose correlation is higher than the mean in the *top ten* attributes, but drops below the mean for the *below top ten* attributes. Finally, in the third column the pairs which are correlated with the coefficient lower than the mean for *all* attributes are presented.

From Table 5 we can see that high correlation in attributes is always present when the object types have e.g., the same look, design or the same purpose. For the urban areas we have object types such as *church*, *basilica*, *abbey*, *cathedral*

Table 5: Change of correlation between object types when different sets of attributes are taken for comparison.

Always High Correlation	High Correlation in Top Ten and Low Correlation in Below Top Ten Attributes	Always low correlation
church;temple, abbey;temple, building;tower, building;residence, abbey;basilica, house;residence, abbey;church, basilica;temple, basilica;cathedral, museum;operahouse, house;palace, gallery;museum	house;prison, museum;railway, basilica;railway, monument;prison, university;cathedral, prison;palace, prison;cathedral, cemetery;aquarium, railwaystation;cathedral, cemetery;parliament	prison;aquarium, monument;railway, railway;cathedral, university;railway, museum;prison, temple;railway, railway;operahouse, parliament;railway, abbey;railway, tower;railway
mountain;volcano, mountain;skiresort, mountain;peak, peak;volcano, peak;glacier, mountain;glacier, glacier;skiresort, park;garden, lake;canal, canal;river	village;river, lake;island, canal;volcano, island;river, mountain;river, peak;lake, island;garden, canal;glacier, skiresort;garden, mountain;canal	lake;volcano, canal;island, lake;glacier, glacier;river, canal;village, lake;garden, village;garden, glacier;village, island;glacier, park;cave

and *temple* which correlate with a coefficient higher than the mean of 0.55 when *all* attributes are considered. A similar picture can be drawn for other urban object types such as *house-residence*, *building-residence* or *museum-operahouse*. In rural areas, object types related to mountainous areas (e.g. *glacier*, *mountain*, *peak*, *volcano*, *ski resort* etc.) are correlated above the mean of 0.53. The same is valid for water bodies like *canal*, *lake*, *river*, etc.

All these object types clearly share features like look, design, purpose, etc. The correlation coefficient is always low for object types which do not share these aspects such as the ones shown in the third column of Table 5. This indicates that aspects used for categorizing objects into object types also play a role in deciding which object types are similar for purposes of describing them.

However, the second column of the tables also highlights the importance of shared ideas of what information is relevant for describing geographic objects for object type similarity (Section 3.1). The object types shown in the second column of the tables are correlated above the mean for *top ten* attributes and their correlation drops below the mean when *below top ten* attributes are used.

3.3 Discussion

Our analyses have shown that people taking the role of a tourist do share ideas as to what information is relevant for describing geographic objects, and we identified a set of information types that reflects these ideas. However, although we could say that there are clearly several information types which people like to know when seeing any geographic object, it was also clear that not every type of information is equally relevant for each object type. These findings prompted the question, whether each object type has a specific set of information associated with it, or whether object types can be grouped according to how they are described by people.

We found that some object types are similar and not only share the information types associated with them, but also the importance ranking of these types. Such similar objects were mostly objects of similar purpose, look or design (e.g. churches and temples, rivers and lakes, etc.). This indicates that sets of features which people use to categorize objects also plays a role in deciding which information should be used to describe these objects. However, some object types which do not share these features are still similar in terms of which information is required to describe them. In these cases people will refer to the shared set of frequently requested attributes we identified.

4. CONCLUSION

In this paper we investigated which information types (attributes) humans associate with geographic objects from urban and rural landscapes. We identified a set of attributes which were relevant for any geographic object type, but also found that an appreciable proportion of attributes is object type specific. Even in the set of shared information types, not all information types were equally important for each object type. Based on the importance ranking of information types for each object type, we were able to identify similar object types.

In our earlier work we extracted the information types which humans associate with geographic objects from existing text resources describing locations around the world and used them to automatically generate image descriptions [1]. One of our future works will be to compare the information types we identified in this paper with the ones extracted from these existing text resources. If these correlate with the set of attributes obtained through our user survey, we could regard the text descriptions containing these attributes as conceptual models of geographic objects of the same type.

5. REFERENCES

- [1] A. Aker and R. Gaizauskas. Generating image descriptions using dependency relational patterns. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1250–1258. Association for Computational Linguistics, 2010.
- [2] K. McKeown and D. Radev. Generating summaries of multiple news articles. *Proc. of the 18th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 74–82, 1995.
- [3] D. Radev and K. McKeown. Generating natural language summaries from multiple on-line sources. *Computational Linguistics*, 24(3):470–500, 1998.