



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
Main Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2007

A theoretical grounding for semantic descriptions of place

Edwardes, Alistair ; Purves, Ross S

Abstract: This paper is motivated by the problem of how to provide better access to ever enlarging collections of digital images. The paper opens by examining the concept of place in geographic theory and suggests how it might provide a way of providing keywords suitable for indexing of images. The paper then focuses on the specific challenge of how places are described in natural language, drawing on previous research from the literature that has looked at eliciting geographical concepts related to so-called basic levels. The authors describe their own approach to eliciting such terms. This employs techniques to mine a database of descriptions of images that document the geography of the United Kingdom and compares the results with those found in the literature. The most four most frequent basic levels encountered in the database, based on fifty possible terms derived from the literature, were road, hill, river and village. In co-occurrence experiments terms describing the elements and qualities of basic levels showed a qualitative accordance with expectations for example, terms describing element of beaches included shingle and sand and, those for beach qualities included being soft and deserted.

Posted at the Zurich Open Repository and Archive, University of Zurich
ZORA URL: <https://doi.org/10.5167/uzh-77800>

Book Section

Originally published at:

Edwardes, Alistair; Purves, Ross S (2007). A theoretical grounding for semantic descriptions of place. In: Ware, J Mark; Taylor, George E. Web and Wireless Geographical Information Systems 7th International Symposium, W2GIS 2007, Cardiff, UK, November 28-29, 2007. Proceedings. Berlin: Springer, 106-120.

A theoretical grounding for semantic descriptions of *place*

Alistair J. Edwardes and Ross S. Purves

Department of Geography, University of Zurich
Winterthurerstrasse 190, 8057 Zurich (Switzerland)
{ajelrsp}@geo.unizh.ch

Abstract. This paper is motivated by the problem of how to provide better access to ever enlarging collections of digital images. The paper opens by examining the concept of *place* in geographic theory and suggests how it might provide a way of providing keywords suitable for indexing of images. The paper then focuses on the specific challenge of how places are described in natural language, drawing on previous research from the literature that has looked at eliciting geographical concepts related to so-called basic levels. The authors describe their own approach to eliciting such terms. This employs techniques to mine a database of descriptions of images that document the geography of the United Kingdom and compares the results with those found in the literature. The most four most frequent basic levels encountered in the database, based on fifty possible terms derived from the literature, were road, hill, river and village. In co-occurrence experiments terms describing the elements and qualities of basic levels showed a qualitative accordance with expectations for example, terms describing element of beaches included shingle and sand and, those for beach qualities included being soft and deserted.

Keywords. place, semantics, image retrieval, basic levels, concept ontology, natural language description

1 Introduction

Needs for geographic information increasingly encompass sources of data that do not fit easily with conventional systems for spatial data handling. In addition, technologies supporting these needs such as geographic information retrieval, collaborative web-maps and location-based services suppose quite different modes of use to more conventional Geographic Information Systems (GISystems). In particular, they seek to bring access to geographic information much closer to the ways in which people engage with the world and each other in everyday life [13][20]. This move from the universal and objective to the more individual and subjective necessitates a change in how geography is presented in GISystems. One consideration in achieving this shift is in how the idea of *place* might be articulated in GISystems [17][18][11]. Fisher and Unwin [14] note of this problem, “GI theory articulates the idea of absolute Euclidean spaces quite well, but the socially-produced and continuously

changing notion of *place* has to date proved elusive to digital description except, perhaps, through photography and film.” (p. 6).

Curiously, this notion of *place* as being encapsulated by images has received relatively little attention in the field of image retrieval where indexing and retrieving a photograph is predicated on the ability to access a description of its contents in some form. Often such a description will have a strong geographical basis related to aspects of *where* the photograph was taken [2][3][26] such as by naming locations and features, events and activities, identifying types of geographical forms or environmental scene, or annotating particular characteristics of the physical or socio-economic landscape depicted.

In the field of image retrieval, which is generally recognised to lag behind the current state of the art in text retrieval, two main techniques exist to index the contents of a photograph, text-based and content-based [15].

Text-based image retrieval uses textual descriptions associated with image. These might be drawn from:

1. Descriptive terms or phrases chosen freely by an indexer [1].
2. Classification codes or controlled keyword lists manually assigned by an indexer [5].
3. The name of the image file or text associated with the photograph, e.g. the contents of a web-page [23].

The main issue with the first method is that terms selected can be highly subjective, prone to misspellings and inter-annotator agreement is often low [22]. One method to overcome such shortcomings has been through the deployment of annotation games where players gain points when they provide the same terms for a given image [1]. A possible criticism of this approach however is that it might bias annotation towards the use of terms that are essentially the ‘lowest common denominators’, for example objects and dominant colours, and so miss out on more complex semantics. The second method can ameliorate such factors by using a controlled list of terms and a more systematic approach to indexing. However, such word lists will usually be extensive, for example the Getty Images collection is indexed using 12,000 keywords with 45,000 synonyms [5], and are usually designed for specialised purposes e.g. Medicine or Art History. This means it is a time-consuming and labour intensive process to index images. In addition retrieval may be dependent on how familiar a user is with the indexing system and its terminology. The third method has the problem that there is often only a limited semantic overlap [23] between the contents of an image and the related text, which can be misleading to image retrieval. For example, Balasubramanian *et al.* [3] note that for most document genres on the web the text does not describe the image, but rather the image is used as support for the text.

Content-based image retrieval uses image processing techniques to try to match images to primitive features, for example based on colour, shape and texture [10]. A number of systems have been made commercially available; however their take-up

has been minimal owing to a mismatch between the needs of users for higher level semantics and the capabilities of such systems that operate at a more primitive level. This difference is generally termed *the semantic gap*. Smeulders *et al.* [27] define this as:

“The semantic gap is the lack of coincidence between the information that one can extract from the visual data and the interpretation that the same data have for a user in a given situation.” (p.1353)

In attempting to bridge the semantic gap, content based techniques will often classify primitive level features against higher level semantic concepts (e.g. ‘car’ or ‘mountain’) defined in a controlled vocabulary (for example LSCOM <http://www.lscom.org/>). There is therefore a need in both text and content-based retrieval techniques for a structured set of concepts which can be used to index photographs.

The work described here focuses on text-based image retrieval and deals with the question of how people describe photographs in natural language, and in particular photographs whose contents are particularly geographic such as landscapes. The motivation of this work is the development of new technologies for indexing and retrieval of photographs using geographical aspects related to the setting in which they were taken. We hypothesise that, by describing *where* a picture was taken, we will also be able to describe *what* is in the picture on many occasions. Furthermore, we assume that by mining geographic data at a given location we will be able to formulate a useful description of *place*. Ultimately, the aim of this work will be to develop a *concept ontology* that structures a set of geographical descriptors that can be used to describe photographs and that can be derived from spatial data sources.

In this paper we firstly set out a theoretically solid foundation for exploring ways of describing *place* by reviewing literature on how the problem has been dealt with in different domains. Secondly, we examine a rich collection of georeferenced images to empirically explore how, in a given context, *place* is described. We then discuss how these results can be employed to help develop a *concept ontology* based on the use of basic levels [25].

2 Describing images through *place*

2.1 Image description

Shatford [26] considered the general issue of how people classify images, their contents and meaning. Shatford simplified the terminology of art historian Panofsky into three levels, *Generic Of*, *Specific Of* and, *About*. In relation to these she specified

four facets, *who*, *what*, *where* and *when*, to generate the Panofsky-Shatford facet matrix (see also [2]) (Table 1).

Table 1. The Pansofsky-Shatford facet matrix (Shatford, [26], p. 49)

Facets	Specific Of	Generic Of	About
Who?	Individually named persons, animals, things	Kinds of persons, animals, things	Mythical beings, abstraction manifested or symbolised by objects or beings
What?	Individually named events	Actions, conditions	Emotions, Abstractions manifested by actions
Where?	Individually named geographic locations	Kind of place geographic or architectural	Places symbolised, abstractions manifest by locale
When?	Linear time; dates or periods	Cyclical time; seasons, time of day	Emotions or abstraction symbolised by or manifest by

The ‘Where?’ facet in Table 1 is particularly interesting for geographers. Here, “*Specific Of*” can be seen as referring to toponyms and entries in a gazetteer and “*Generic Of*” represents different types of topographic and architectural features. “*About*” then essentially introduces the concept of *place* and associates meaning with geographic locations. Hence, it is interesting to investigate the extent to which *place* related *About* concepts are used by people when describing images.

2.2 Place

It is often stated that Geography is the study of *Space* and *Place* [14]. These seemingly simple categories underlie many of the differences among geographers both in terms of what they study and how they go about it. The two concepts are often examined with respect to a continuum of viewpoints that range from the particular and the experiential at one end, to the abstract and the universal at the other [7]. *Place* is most distinct at the start, relating geography to human existence, experiences and interaction [29]. At the other end is the more detached, abstract and objective view of *space*, ultimately represented as geometry, which provides a means to think about, describe and encode the world in a logical way. As such, *Space* and *Place* are often emphasised as a ways of understanding (epistemology) rather than as ontological categories such as a geographic location [9].

A number of researchers have sought to characterise the concept of *place*. Relph [24] highlighted elements of the physical setting (including location and physical appearance), the activities performed in places, and the meanings places have for people. Canter [6] identified four facets of place, *functional differentiation* (activities), *place objectives* (individual, social and cultural aspects of place), *scale of interaction* (e.g. room, home, neighbourhood city), and *aspects of design* (physical characteristics

of a place). Gustafson [16] developed a somewhat different model encompassing, *self* (e.g. personal meanings and self-identification), *others* (social relations and norms), and *the environment* (the physical characteristics and natural conditions).

2.3 Natural language descriptions of *place*

Work has also sought to understand how places can be described using natural language. Such studies have tended to take one of three approaches:

1. Collection of words from a dictionary validated by experts or experiment participants e.g. [12]
2. Collection of terms from documents relevant to a domain validated by experts or experiment participants e.g. [19]
3. Elicitation of terms directly from participants in an experiment

Results from the third approach are particularly relevant to this work. For example, Craik [8] sought to characterise the personality ‘traits’ of a landscape by eliciting adjectival qualities from participants in response to photographic stimuli of natural landscapes. Tversky and Hemenway [30], looked at how people categorise environmental scenes using *basic levels* [25]. They identified different types of basic-level scene by considering the informativeness of a word for a scene in respect to the different characteristics, attributes (qualities), types of related activities and, component parts. Their resultant basic level scene types have since been employed by researchers interested in content based image retrieval. For example, Vogel and Schiele [32] perform statistical analysis of local detail in images (amounts of *sky, water, grass, trunks, foliage, field, rocks, flowers* and *sand*) to categorise images into basic-level scene categories and transitional states between them. The work of Smith and Mark [28] also falls into this type of approach. They drew on the methodology of Battig and Montague [4] for eliciting category norms (similar to basic levels) from participants, and sought to identify ‘folk’ categories related to geography. They showed that, for example, that “mountain” was the most common norm identified for the category “a geographic feature.” Most recently, Van Overschelde *et al.* [31] have also repeated Battig and Montague’s work, also finding “mountain” as the most common norm for “a natural earth formation.”

3 Experiments in data mining descriptions of places

A common goal in the work described in the previous section has been to identify vocabularies of concepts related to *place*. An obstacle for many of these studies is that they require human subject experimentation which is both costly and time consuming to perform. In this work we attempt to overcome this problem by analysing a database of textual descriptions of photographs which have a strong geographic component.

In order to explore the potential of this resource for developing such a vocabulary (which we will term a *concept ontology*), we present two experiments that employ

simple data mining techniques to identify significant concepts. In the first we consider relative term frequencies (as an indicator of basic level categories) in order to compare the database as a resource with human subject testing reported in the literature. In the second we investigate how the resource can be employed to obtain sets of concepts that help characterise the semantics of basic level types of places

3.1 Data

In these experiments we used data obtained from Geograph (www.geograph.org.uk). Geograph is a project with the aim to collect “geographically representative photographs and information for every square kilometre of the UK and the Republic of Ireland.”



Fig. 1. Typical Geograph image – the associated caption reads: **Gatliff Trust Hostel on Berneray:** Picture taken from the beach on Berneray of the historic Gatliff Trust Hostel. Visited in the 1990s, shortly before the causeway linking Berneray to North Uist was built.

The project allows contributors to submit photographs representing individual 1km grid squares, and after moderation these images are uploaded together with descriptive captions to a publicly available web site. Figure 1 shows a representative example. Since a key aim of the project is to “...show at close range one of the main geographical features within the square” and to produce an associated description, we contend that the dataset may well be a digital representation of *place*. The database underlying Geograph is available under a Creative Commons licence. We obtained a snap shot taken on the 24th of February 2007, consisting of 346270 images taken by a total of 3659 individuals.

3.2 Methods

We performed two experiments using the Geograph data. In the first of these we looked at counts of terms identified as being basic levels or category norms in the experiments of Battig and Montague, Smith and Mark and Van Overschelde *et al.* (Table 2). Since we expected that some of these terms may vary according to the domain in which the original experiments were carried out (the US) we added appropriate synonyms for a UK context, (see [21] for a discussion). For example, we included the terms *loch* (Gaelic for lake) and *glen* (Gaelic for valley). We counted the total number of occurrences of each of these terms in the image descriptions in the Geograph database.

Table 2. Set of seed terms used for Experiment 1, sources are 1: Battig and Montague, 2: Smith and Mark, 3: Van Overschelde *et al.* and, 4: Added by the authors.

<i>Term</i>	<i>Source</i>	<i>Term</i>	<i>Source</i>	<i>Term</i>	<i>Source</i>
Atlas	2	Glacier	3	Plain	1,3
Beach	3	Glen	4	Plateau	1,3
Building	2	Globe	1	River	1,2,3
Canyon	1, 3	Gorge	4	Road	1
Cave	1, 3	Grass	3	Rock	1,2,3
City	2	Highway	1	Sea	1
Cliff	1, 3	Hill	1,2,3	Sky	4
Clouds	4	Island	1,3	Stream	1,3
Coast	4	Lake	1,2,3	Street	1
Compass	1	Land	1	Town	1
Continent	1	Loch	4	Tree	4
Country	1	Map	1	Valley	1,2,3
County	1	Motorway	4	Village	4
Delta	1	Mountain	1,2,3	Volcano	1,3
Desert	2,3	Ocean	1,2,3	Waterfall	3
Elevation	1	Park	1	Woodland	4
Field	4	Peninsula	1	Woods	4
Forest	1				

In the second experiment we wished to explore how basic levels can be characterised. We chose four *basic levels* identified in the first experiment and explored the co-occurrence of these basic levels with two wordlists containing terms about: *elements* and *adjectives*. The list of elements attempted to describe the component parts of a scene as suggested by Tversky and Hemenway [30] and consisted of a list of 540 terms obtained from dictionary analysis using WordNet (<http://wordnet.princeton.edu/>). The list of adjectives employed the Landscape Adjective Check List of Craik and was aimed at capturing more abstract, *About* type concepts relating to *place*. In this experiment we searched for descriptions in which both a basic level and a term co-occurred. A description was counted only once however many times a pair of terms co-occurred within it, although in practice most descriptions contained only one or two sentences. We retrieved not only the count of co-occurrences of terms with basic levels, but also the count of occurrences of both

the term and the basic level, in order to explore whether the frequency of a co-occurrence was simply explained by a term being very common.

4 Results

Table 3 compares the top rankings of basic level terms described in the literature with the ranking from the Geograph database using the seed terms given in Table 2. For the Battig and Montague and Van Overschelde *et al.* experiments participants had been asked to name ‘natural earth formations’ for the Smith and Mark experiment we unioned all terms elicited by all their experimental categories (e.g. ‘a kind of geographical feature’, ‘something geographic’ etc.) and ranked them according to the total number of responses.

Table 3. Top 10 terms from Battig and Montague (B&M), Top 20 from Smith and Mark (S&M), Van Overschelde *et al.*, (vanO) and this work (Geograph).

B&M	VanO	S&M	Geograph
Mountain	Mountain	Mountain	Road
Hill	River	River	Hill
Valley	Ocean	Lake	River
River	Volcano	Ocean	Village
Rock	Lake	Hill	Building
Lake	Valley	Country	Park
Canyon	Hill	Sea	Street
Cliff	Rock	City	Valley
Ocean	Canyon	Continent	Field
Cave	Plateau	Valley	Loch
	Tree	Plain	Land
	Plain	Plateau	Town
	Cave	Map	Forest
	Glacier	Road	Map
	Grand Canyon	Island	Sea
	Island	Desert	Woodland
	Stream	Peninsula	Tree
	Cliff	State	Beach
	Desert	Volcano	Country
	Beach	Forest	Glen

In Table 4, the top ten terms of Battig and Montague are used to consider the rank obtained from the other experiments. The Geograph data are first ranked according to the raw occurrences of terms and secondly after substituting the terms sea for ocean, loch for lake and, gorge for canyon, since these were considered to be synonyms for these terms in British English. Correlations between the rankings are also shown with the ranking of Smith and Mark being significantly correlated with that of van ($p < 0.05$) and the Geograph ranking (without substitution) significantly correlated with that of Battig and Montague ($p < 0.05$).

Table 4. Correlation between the relative rankings of each study using only the terms found by Battig and Montague.

	B&M	S&M	VanO	Geograph	with subs
Mountain	1	1	1	6	7
Hill	2	6	6	1	1
Valley	3	5	5	3	3
River	4	2	2	2	2
Rock	5	7	7	5	6
Lake	6	3	4	4	5
Canyon	7	8.33	8	10	10
Cliff	8.5	8.33	10	7	8
Ocean	8.5	4	3	9	4
Cave	10	8.33	9	8	9
Spearman's Rank Correlation					
	B&M	S&M	VanO	Geograph	with subs
B&M	1	0.625	0.597	0.7	0.555
S&M		1	0.968	0.463	0.555
vanO			1	0.345	0.539
Geograph				1	0.818
With subs					1
significance level at $P_{0.05} = 0.648$					

Tables 5 and 6 show the total number of counts of elements and adjectival terms co-occurring with four basic level types of place (beach, village, hill and mountain). The terms were ranked by the frequency of the co-occurrence with respect to the occurrence of the term (e.g. shingle appears 254 times in total, of which 43% of occurrences are found in conjunction with beach).

Table 5. Element terms associated with basic-levels. *N* is the total number of occurrences of element term, *n* the no. occurrences of the basic level scene term and *Freq* the number of co-occurrences/*N*.

Elements					
<i>Term</i>	<i>Freq.</i>	<i>N</i>	<i>Term</i>	<i>Freq.</i>	<i>N</i>
Beach		n=2824	Village		n=12707
Shingle	0.43	254	Pub	0.15	3205
Sand	0.17	1037	Shop	0.13	1464
Cliff	0.10	1124	Inn	0.12	1892
Headland	0.10	524	Church	0.10	16157
Bay	0.08	2770	Housing	0.07	2820
Sea	0.08	3725	Edge	0.07	6541
Rock	0.06	1914	Cottage	0.06	2966
Coast	0.04	2548	Main Road	0.05	6735
Shore	0.04	1298	Village green	0.05	12707
Island	0.03	2645	Stone	0.05	4817

Hill			Mountain		
		n=16232			n=1256
Fort	0.30	1124	Peak	0.03	823
Top	0.21	6086	Summit	0.03	3427
Summit	0.18	3427	Ridge	0.02	3263
Horizon	0.14	2047	Moorland	0.02	1484
Ridge	0.13	3263	Quarry	0.01	2273
Sheep	0.10	2659	Stream	0.01	2478
Valley	0.08	8563	Sheep	0.01	2659
Side	0.08	11369	Forest	0.01	4542
Trees	0.08	7753	Top	0.01	6086
Track	0.08	8331	Path	0.01	6477

Table 6. Adjectival terms associated with basic-levels. N is the total number of occurrences of adjectival term, n the no. occurrences of the basic level scene term and $Freq$ the number of co-occurrences/ N .

<u>Adjectives</u>					
<i>Term</i>	<i>Freq.</i>	<i>N</i>	<i>Term</i>	<i>Freq.</i>	<i>N</i>
Beach			Village		
		n=2824			n=12707
Sandy	0.26	553	Deserted	0.31	233
Deserted	0.10	233	Pretty	0.15	616
Eroded	0.09	288	Green	0.14	4993
Soft	0.08	181	Quiet	0.12	1005
Rocky	0.06	756	Lovely	0.10	919
Warm	0.06	241	Pleasant	0.10	781
Glacial	0.06	290	Beautiful	0.08	1177
Low	0.05	2866	Remote	0.07	510
Beautiful	0.05	1177	Unusual	0.07	1175
Lovely	0.02	919	Large	0.06	5273
Hill			Mountain		
		n=16232			n=1256
Steep	0.19	2131	Distant	0.03	810
Distant	0.18	810	Black	0.03	1893
Wooded	0.16	1053	Remote	0.02	510
Black	0.12	1893	Rocky	0.02	756
Rough	0.12	1367	Grassy	0.01	622
Grassy	0.11	622	Steep	0.01	2131
Round	0.11	1605	Natural	0.01	702
Big	0.10	1337	Dark	0.01	537
White	0.10	2642	Broad	0.01	655
Broad	0.09	655	Running	0.01	2480

5 Discussion

Experiment 1 looked at the relative occurrence of particular terms in relation to studies on geographic basic levels reported in the literature. The aim of this experiment was to investigate how comparable the basic level terms identified in the Geograph database were with those resulting from direct human subjects testing.

The results show a number of interesting differences between the Geograph data and the other work. One obvious difference is the presence of highly ranked anthropogenic features in the Geograph list. This is not surprising in the case of the Battig and Montague and Van Overschelde *et al.* lists since participants in these experiments were asked for ‘natural earth formations’. However, Smith and Mark’s experiments did elicit anthropogenic objects, but these were not found in their top ten. By contrast, such terms appear frequently in the Geograph result. In addition, those that are in the Smith and Mark’s list tend to be on have a larger spatial extent (e.g. city) than those of the Geograph list (e.g. building, town and village).

A second observation is that certain concepts in the Geograph list appear to have been substituted. Hill appears to be used instead of mountain, sea rather than ocean, and loch rather than lake. On the one hand, these differences can be seen as relating to the variations in experimental methods. The literature studies asked participants to think of concepts in response to abstract questions whereas the Geograph data comes from people describing concepts that can be seen in a picture. For example, there are few actual examples of concepts such as canyons or volcanoes in Britain, which is likely why these were not frequently found. On the other hand, this substitution could suggest that the use of terms is not insensitive to geographical context, as found in previous empirical work [21]. In Britain the term hill appears to be much more widely used than mountain, and large bodies of open water which would be described as lakes in England occur more frequently in Scotland where the Gaelic appellation loch is almost exclusively used. Interestingly, lochs may be open to the sea and thus may describe a concept which overlaps with, but is not identical to lake, defined by WordNet as “a body of (usually fresh) water surrounded by land.” In addition, whilst Britain is bordered by both oceans and seas, it would appear that the term sea is more conventionally used.

Table 4 compared the relative rankings of the most frequently used ‘natural earth formations’ of Battig and Montague. For better consistency in terminology, the terms and rankings for the question “a kind of geographic feature” were used in the case of Smith and Mark. It is interesting to note that the term rankings of Van Overschelde *et al.* and Smith and Mark using students from the US correlate significantly at the $p < 0.05$ level, but that neither of these lists correlate significantly with either those of Battig and Montague or derived from Geograph. Again, this might suggest that the conceptualisations are sensitive to locale. Of note is that the substitutions (e.g. sea for ocean) increased correlations though not sufficiently to be significant. One possible suggestion is therefore that the difference is not only one of terminology but also of

environmental familiarity (for example British landscapes consist less of lakes and mountains and more of hills and valleys). The Geograph ranking was significantly correlated with that of Battig and Montague, though only where no terminology substitutions were applied. However caution is required to avoid over interpreting these results since the significance of the Spearman's rank correlation is very sensitive when comparing so few categories.

Experiment 2 aimed to investigate the ability of the Geograph dataset to provide characteristic sets of concepts describing semantics for particular examples of basic level places. It used word lists of elements (parts) and adjectives in relation to two basic levels (village and hill) identified in the previous experiment and two (mountain and beach) suggested by Tversky and Hemenway. The selection of these levels was to both consider places that were typical to the area covered by the Geograph dataset and ones that would allow comparison with Tversky and Hemenway's results in terms of contents. In addition, it was intended to examine the ability of the technique to provide sets of terms that are both sensible and distinct enough to discriminate between levels even, in the case of hill and mountain, when they could be seen as somewhat synonymous.

It is interesting to note (Table 5 and 6) that there are strong differences in the terminology used for each of the basic levels. In terms of elements, beach was identified with coastal landforms and materials like rock, sand and sea. This compares well with the results of Tversky and Hemenway whose list contained water and sand. The elements of village were instead very functional and architectural and typical of what might be expected in typical British villages e.g. cottages, a village green, a shop, a pub and a church. There are some similarities between hill and mountain, particularly in the use of perceptual terms such as top, summit, and ridge. However, differences can also be seen both in terms of the language used for similar concepts, for example path and forest were found for mountains and track and trees for hills, and in terms of composition, for example valley and hill sides for hills and moorland and streams for mountains. Similarities for mountain were also found with the work of Tversky and Hemenway, for example animal, trail, and stream and; sheep, path, and stream. One interesting difference is where they found grass and tree, our results identified moorland and forest, which are more typical of the mountain landscapes found in Britain.

The results for adjectives were also quite distinct for each of the levels. Beaches, for example, tended to be described using process (eroded, glacial) and affective (soft, beautiful, warm, lovely) terms. Villages have some similarities with the use of affective terms, but tend to invoke more of a sense of an idyllic place (pretty, lovely, pleasant, beautiful), in addition there seems to be an emphasis on remoteness from other settlements (deserted, quiet, remote). The adjectives for hill and mountain are in many ways similar to the description of their elements describing perceptual characteristics and land cover. To an extent these are shared (steep, distant, broad, black, grassy) but there are also differences. Hills are described as being big, round and wooded whereas mountains are described in terms that are more desolate such as remote, natural, rocky and, dark.

Overall the results are very promising. The sets of concepts found for the basic levels fit well with the authors' conceptions of the types of places they describe, at least in a British context. They are also diverse enough to allow discrimination between levels. For mountain and hill there are a number of commonalities (which is expected for basic levels, see [30] p.123). However these tend to reflect the perceptual characteristics of the landforms which are inherently similar. Other characteristics such as land cover, infrastructure and remoteness were much more differentiated suggesting that the two concepts cannot be seen as synonymous. Nonetheless, a degree of caution is required as, as can be seen in Tables 5 and 6, the occurrence values (n) for mountain concepts are generally an order of magnitude lower than for hill. This reflects the scarcity of mountains within the dataset in general and makes the level sensitive to idiosyncratic descriptions.

The experiments presented thus far are preliminary work and have a number of important limitations, related to the dataset used, the set of assumptions and, the experimental methods themselves. Firstly, we assume that the Geograph dataset will help us to build a concept ontology of how people describe *place*. The dataset itself is very large (346270 images taken by a total of 3659 individuals) which we hope gives some universality to the results. However, as demonstrated in Experiment 1 the results (for example the use of the term hill instead of mountain) appear sensitive to locale, and the participants in the Geograph experiment are likely to be a self-selecting group (they have access to the internet and digital cameras and an interest in documenting the geography of the UK), which implies in turn that the results may be biased. Secondly, the terms found are limited to those we searched for, and although we based our search terms on theory we cannot rule out synonyms having much higher frequencies. Third, we did not treat toponyms in a special way – we assume that if hill is mentioned the picture itself and thus the place is a hill – in other words that a toponym is an instantiation of a type of place. More importantly, terms which are found in toponyms (e.g. “Black Hill”) are not being used to describe a location and thus our co-occurrence experiments should be controlled for such effects. However, in most cases it appears unlikely that the adjectives found are likely to be parts of toponyms. In general we did not control for either word sense (i.e. does running refer to an activity or describe a property of a stream) or context (i.e. is the term really being used in conjunction with the basic level in question). Nonetheless, we believe that the results presented demonstrate the potential richness of new sources of data as possible methods of exploring the creation of, for example, landscape concept ontologies, which have until now been mostly limited to human subject experiments with correspondingly much smaller data volumes.

6 Conclusions and further work

The research reported here has argued that as geographic information technologies expand into new areas and new aspects of life they need to be supported by more extensive ideas of geography, better reflecting how people experience and think about

the world in everyday life. Here, theory on *place* has been proposed to provide support in exploring the problem of image description and retrieval.

The work has demonstrated a novel approach to eliciting concepts of *place* as they are expressed in natural language through the analysis of a large database of images and their captions from a collaborative web project (Geograph). This resource has been shown to compare favourably to approaches in the literature that have been based on human subject testing, though caveats have been acknowledged in relation to sources of bias. The use of online resources has the advantage that it can help to obviate the need for costly and time-consuming participant experiments.

The potential of the resource to develop an ontology of *place* concepts for use in indexing and search of images has also been explored. This was based on the idea of basic level scene types and characteristic terms (elements and adjectives) that commonly co-occur with these. In particular, it was found that adjectival terms (less studied previously) are important descriptors and discriminators of places.

These experiments will be expanded in future work to populate a more comprehensive concept ontology, considering both a wider variety of basic levels and a number of other place related aspects (e.g. activities). The influence of toponyms and ambiguity on our results also be examined, as well as the issue of how to detect such concepts from analysis of spatial and other data.

Acknowledgements. We would like to gratefully acknowledge contributors to Geograph British Isles, see <http://www.geograph.org.uk/credits/2007-02-24>, whose work is made available under the following Creative Commons Attribution-ShareAlike 2.5 Licence (<http://creativecommons.org/licenses/by-sa/2.5/>). The research reported in this paper is part of the project *TRIPOD* supported by the European Commission under contract 045335.

References

1. Ahn, L. von and Dabbish L.: Labelling Images with a Computer Game, CHI 2004, April 24–29, 2004, Vienna, Austria. (2004)
2. Armitage, L.H. and Enser, P.G.B.: Analysis of user need in image archives. *Journal of Information Science*, 23(4), 287–299 (1997)
3. Balasubramanian, N., Diekema, A.R., and Goodrum, A.A: Analysis of User Image descriptions and Automatic Image Indexing Vocabularies: An Exploratory Study. International Workshop on Multidisciplinary Image, Video, and Audio Retrieval and Mining. Sherbrooke, Quebec, Canada, October 25-26, (2004)
4. Battig, W. F., and Montague, W. E.: Category norms for verbal items in 56 categories: a replication and extension of the Connecticut Norms. *J. of Experimental Psychology*, 80(2), 1–46 (1969)
5. Bjarnestam, A.: Text-based Hierarchical Image Classification and Retrieval of Stock Photography, The Challenge of Image Retrieval Conference. Newcastle UK. (1999)
6. Canter, D.: *The Psychology of Place*. Architectural Press, London (1977)
7. Couclelis H.: Location, place, region, and space. In: Abler R.F., Marcus M.G. and Olson J.M. (eds.), *Geography's Inner Worlds*, Rutgers University Press, 215–233 (1992)

8. Craik, K.H.: Appraising the Objectivity of Landscape Dimensions, In: *Natural Environments: Studies in Theoretical and Applied Analysis*. John V. Krutilla (ed). Resources for the Future, Baltimore. pp. 292–346 (1971)
9. Curry M.R.: Toward a Geography of a World Without Maps: Lessons from Ptolemy and Postal Codes, *Annals of the Association of American Geographers*, 95(3), 680–691 (2005)
10. Eakins, J.P. and Graham, M.E.: Content-based image retrieval, *JISC Technology Applications Programme Report 39*, October (1999)
11. Edwardes, A.J.: *Re-placing Location: Geographic Perspectives in Location Based Services*, Ph.D Thesis, University of Zurich. (2007)
12. Edwards, J. A and Templeton, A.: The structure of perceived qualities of situations, *European J. of Social Psychology*, 35, 705–723 (2005)
13. Egenhofer, M. J., and Mark, D. M.: Naive geography. In: *Spatial Information Theory: A Theoretical Basis for GIS*, A.U. Frank and W. Kuhn (eds), LNCS 988 Berlin: Springer-Verlag (1995)
14. Fisher P. and Unwin D.: Re-presenting Geographical Information Systems. In: *Re-presenting GIS*, Unwin D.J. and Fisher P (eds) Wiley&Sons, London, pp. 1–17. (2005)
15. Goodrum, A. A.: Image Information Retrieval: An Overview of Current Research, *Informing Science*, 3(2), 64–67 (2000)
16. Gustafson P.: Meanings of place: everyday experience and the theoretical conceptualizations, *J. of Environmental Psychology*, 21, 5–16 (2001)
17. Harrison S. and Dourish P.: Re-place-ing space: The roles of place and space in collaborative systems, *Proc. Computer Supported Cooperative Work (CSCW '96)*, Boston MA, ACM Press, New York, pp. 67–76 (1996)
18. Jordan T., Raubal M., Gartrell B. and Egenhofer M.: An affordance-based model of place in GIS, *Proc. of 8th Intl. Symp. on Spatial Data Handling, SDH'98*, Vancouver, Canada: International Geographic Union, pp. 98–109 (1998)
19. Kasmar, J.V.: The Development of a Usable Lexicon of Environmental Descriptors, *Environment and Behavior*, 2, 153–169 (1970)
20. Llobera M.: The nature of everyday experience: examples on the study of visual space. In: *Re-presenting GIS*, Unwin D.J. and Fisher P (eds.), Wiley & Sons, London, 171–195 (2005)
21. Mark, D.M and Turk, A.G.: Landscape Categories in Yindjibarndi: Ontology, Environment, and Language. *COSIT 2003*, 28–45 (2003)
22. Markey, K.: Interindexer Consistency Tests: A Literature Review and Report of a Test of Consistency in Indexing Visual Materials., *Library and Information Science Research*, 6, 155–77 (1984)
23. Marsh, E.E., White, M.D.: A taxonomy of relationships between images and text. *J. of Documentation*, 59, 647–672 (2003)
24. Relph E.: *Place and Placelessness*, Pion, London. (1976)
25. Rosch, E.: Principles of categorization. In: *Cognition and Categorization*, E. Rosch and B. B. Lloyd (eds.), Erlbaum, Hillsdale, NJ (1978)
26. Shatford, S.: Analyzing the subject of a picture: a theoretical approach. *Cataloguing and Classification Quarterly*, 6, 39–62 (1986)
27. Smeulders A.W.M., Worring, M., Santini, S., Gupta, A. and Jain, R. Content-Based Image Retrieval at the End of the Early Years, *IEEE Transactions on Pattern Analysis And Machine Intelligence*, 22(12), 1349–1380 (2000)
28. Smith, B. and Mark, D.M.: Geographical categories: an ontological investigation, *Int. J. Geographical Information Science*, 15(7), 59–612 (2001)
29. Tuan Y-F.: *Space and Place: The Perspective of Experience*, U of Minnesota Press. (1977)
30. Tversky, B. and Hemenway, K.: Categories of Environmental Scenes, *Cognitive Psychology*, 15, 121–149 (1983)

31. Van Overschelde, J.P., Rawson, K.A and Dunlosky, J.: Category norms: An updated and expanded version of the Battig and Montague (1969) norms, *J. of Memory and Language*, 50, 289–335 (2004)
32. Vogel, J. and B. Schiele, B.: A Semantic Typicality Measure for Natural Scene Categorization. German Symposium on Pattern Recognition DAGM 2004, Tuebingen, Germany, LNCS 3175. Springer-Verlag, Berlin (2004)