Gender-based Models of Location from Flickr

Neil O'Hare
Yahoo! Research, Barcelona, Spain
nohare@yahoo-inc.com

Vanessa Murdock*
Microsoft
vanessa.murdock@yahoo.com

ABSTRACT

Geo-tagged content from social media platforms such as Flickr provide large amounts of data about any given location, which can be used to create models of the language used to describe locations. To date, models of location have ignored the differences between users. This paper focuses on one aspect of demographics, namely gender, and explores the relationship between gender and location in a large-scale corpus of geo-tagged Flickr images. We find that male users are much more likely to geo-tag their photos than female users, and that the geo-tagged photos of male users have wider geographic coverage than those of females.

We create gender-based language models of location using the Flickr tags describing geo-tagged photos, and find that Flickr tags created by male users contain more geographic information than those created by female users, and that they can be located based on their tags far more accurately. Further, models created exclusively with data from male users are more accurate than those created from female users' data. Although our results suggest that there is some benefit from using gender-specific models, this benefit is quite minor, and is overwhelmed by the richer location information in the male data. The results also show that gender-based differences in location models are more important at the hyper-local level.

Categories and Subject Descriptors

H.3.1 [**Content Analysis and Indexing**]: Miscellaneous

Keywords

Geo-tagging, Flickr, gender

1. INTRODUCTION

Geo-tagged content from social media platforms such as Flickr[1] and Twitter[2] provide a huge amount of data describing what people are doing, thinking or photographing at a particular time and place. People tag photos to indicate the content and the context of the images, often including the locations where they were taken. A geo-tagged tweet describes what people are doing, thinking about, or planning to do in a particular place and time.

Existing work has investigated the use of language models of locations, based on these resources [13, 5]. That work, however, does not take into account the differences between users, and treats all users equally. The photos taken, or the moments tweeted about, by a twenty year old male walking along La Rambla in Barcelona, for example, may be quite different from those of the thirty year old mother of a new baby. If we can exploit profile information about users such as age and gender, we can build models that capture the diversity of language used by different groups of people. We propose that this gives insight into what people in specific demographic groups do and talk about in a given place.

This paper represents a first step in exploiting user profile information to create models of location for particular groups. We focus on gender, as one of many potentially useful types of demographic information, as it is ready available as self-reported information in Flickr profiles. We build gender-specific language models of location, and explore the degree to which gender-specific representations of location differ from one another, and the degree to which gender-specific models can improve over gender-agnostic models.

The rest of this paper is organised as follows: in the next Section we review related work, followed by a description of gender-based models of location in Section 3. We describe our Flickr dataset in Section 4, followed by an empirical evaluation in Section 5, before concluding in Section 6.

2. RELATED WORK

In this Section we summarise related work, on modelling location with user generated content, and on exploiting gender in creating models from user generated content.

Modeling Location with User Generated Content. There has been a lot of interest in using geo-tagged Flickr photos for modeling locations. Crandall et al. [6] investigated the use of visual, textual and temporal features to infer the location of Flickr photos. They perform a 10-way classification over 10 popular landmarks in 100 America cities, and found

*Work done while at Yahoo! Research, Barcelona.

[1]http://www.flickr.com
[2]http://www.twitter.com

GeoMM'12, October 29, 2012, Nara, Japan.

that using image content in addition to textual features improved performance. Hays and Efros [7] estimate the location of Flickr photos purely based on the visual content of photos using K-Nearest Neighbour matching based on low level image features, and report 16% accuracy in placing test photos within 200km of the true location.

The work of Serdyukov et al. [13] places no restrictions on the locations that are modeled, only requiring that photos have descriptive text in the form of tags. They quantise the globe into a grid at various resolutions (100km, 10km, 1km) and build a language models for each of these cells, ranking candidate locations for test photos by the probability that their language model created the photo's tags. The current work is quite similar to that work in that we also create language models over a quantized representation of the globe. Our models are created with a much larger dataset, however, and we improve the models by estimating term probabilities with user frequency. Most importantly, the focus of this work is on gender variations in the language used to describe locations, whereas they were primarily interested in the task of locating Flickr photos. Cheng et al.[5], in similar work, created models of users home location based on their tweets, building models of major cities in the US, and report that they can accurately predict the location of almost 51% of users to within 100km of their declared home city

Other work has used query logs to model the geographic scope of queries. Backstrom et al.[2] propose an approach that estimates the geographic centre and dispersion of a query, and also tracks geographically shifting topics over time. Other work has leveraged the declared locations of a user's Facebook friends to determine their location [3].

Gender-based models from User Generated Content. There is relatively little work on using gender information when building models using data from Flickr or other social media platforms. Popescu & Grefenstette [12] explore gender differences in Flickr tags. Their work is different from the current work in that they are not interested in the intersection of gender and location, but instead they explore two separate tasks: classifying users as male or female, and estimating the home location of a user. In subsequent work they explored gender differences in image search [11].

Weber et al [14] use click data from query logs to explore demographic variations in web search, and show that search results and query suggestions can be improved through the use of demographic data. Gender-based differences in writing styles in written documents are explored by Argamon et al [1]. Since the focus of that study is on free-text (as opposed to Flickr tags) not all of their conclusions are applicable to this work. They report differences in the use of pronouns and noun modifers, and they also report that female writing tends to be more 'involved' or personal and male writing tends to be more 'informational'.

3. GENDER-BASED LOCATION MODELS

In this Section we describe how we model locations based on the tags used to describe them, and how we extend the approach to create gender-specific models.

3.1 Representing Locations

We represent the globe as a grid where the grid is created by quantizing the latitude/longitude values to create cells that are 100km or 1km by latitude. Each cell will contain

	All	Geo	Geo/All %
Known Gender	2,465,832	199,343	8.08
Male	1,431,758	151,363	10.57
Female	1,034,074	47,980	4.64
Female/Male	0.72	0.32	-

Table 2: User Statistics for Flickr Dataset

	Users	Photos	Photos Per User
Male	151,363	6,330,719	41.83
Female	47,980	1,835,973	38.27
Known Gender	199,343	8,166,692	40.97
Unknown Gender	122,217	1,934,432	15.83

Table 3: Geo-tagged Corpus Statistics after Bulk Upload Filtering

a number of geo-tagged photos, taken within the cell. For each cell we have a set of photos, all of their textual metadata, including tags, title and description, along with user profile information for the owners of a subset of the photos. In this work we use the tags, which are created with the intention of describing the content and the context of the photos. Although the tags will occasionally explicitly refer to a place-name (e.g. 'barcelona'), this is not always the case, and tags may be used which are more indirectly linked to the location (e.g. 'tapas', 'gaudi'). Other approaches to creating models of location have used the tags as normalised by Flickr [13] (whitespace and special characters removed).

3.2 Language Models of Location

Having created a textual representation for each location in the grid, we create a language model for each of these grid cells, using the standard language model approach to Information Retrieval, as proposed by Ponte & Croft [10]. We build a language model for each location. Given some arbitrary text, T, we want to rank candidate locations, L, by their probability, given that text, $P(L|T)$. Using Bayesian inversion, this gives:

$$P(L|T) = \frac{P(T|\theta_L)P(L)}{P(T)} \tag{1}$$

where θ_L is the model for a location, and $P(T|\theta_L)$ is the probability of the text, given that location model. Assuming independence between terms, this can be calculated as:

$$P(T|\theta_L) = \prod_{i=0}^{|T|} P(t_i|\theta_L) \tag{2}$$

In Equation 1 $P(L)$ is normally assumed to be uniform, and $P(T)$ can be ignored since it is a constant for all locations and does not affect the ranking. This leads to a model where locations are ranked solely by $P(T|\theta_L)$, the probability that the location model created the query text.

Equation 2 requires that we can estimate the likelihood of an individual term, given the language model of a location. The simplest way to estimate this to use maximum likelihood estimate (MLE), which maximizes the observed

	Male Subset	Female	Combined	All
Unique Terms	568,852	416,295	495,371	1,799,551
Locations	254,380	192,517	226,963	886,685
Max Terms	12,695	8,545	9,969	58,282
Male Test Photos Not in Model	36.60%	40.90%	38.07%	20.96%
Female Test Photos Not in Model	39.58%	40.36%	39.75%	22.22%

Table 1: Statistics for the 1km Location Models

likelihood given the data:

$$P_{mle}(t|\theta_L) = \frac{c(t, L)}{|L|} \quad (3)$$

where $c(t, L)$ is the term frequency of the term t in location L, and $|L|$ is the total number of terms in the location. This estimate can be improved using Dirichlet smoothing, which performs a linear interpolation with an estimate calculated from a global background model:

$$P_{dir}(t|\theta_L) = \frac{|L|}{|L| + \mu} P_{mle}(t|\theta_L) + \frac{\mu}{|L| + \mu} P_{mle}(t|\theta_G) \quad (4)$$

$|L|$ is the size of the location (number of terms), θ_G is a background model based on the entire globe, and μ is the Dirichlet smoothing parameter.

Equation 3 estimates the probability of a term based on the term frequency term as a proportion of the term frequency of the location. A problem with this approach is that individual users who tag a lot of photos in a location can come to dominate the textual representation of that location. An alternative is to base the estimate on user frequency, the number of unique users who use the term:

$$P_{user\_mle}(t|\theta_L) = \frac{c_{user}(t, L)}{\sum_{t_i \in L} c_{user}(t_i, L)} \quad (5)$$

where $c_{user}(t, L)$ is the number of unique users who use the term in the location. We no longer divide by the total number of terms in the location, but by the sum of the user frequency of all terms in the location. Estimating term probabilities based on user frequencies in this way should alleviate any bias caused by users who create a disproportionate number of tags in one location. We have shown in previous work that user frequency is significantly more accurate than term frequency in modeling locations [8].

3.2.1 Gender-based Language Models

Since we have gender information for many users, we make use of this information to create gender-specific models of location, allowing us to investigate gender-based differences in describing location. Instead of $P(L|T)$, we calculate $P(L|T, G)$, the probability of a location given the text and a gender, meaning that we replace Equation 5 above by the following:

$$P_{user\_mle}(t|\theta_{L,G}) = \frac{c_{user}(t, L, G)}{\sum_{t_i \in L} c_{user}(t_i, L, G)} \quad (6)$$

where $c_{user}(t, L, G)$ is the number of unique users of a given gender who use the term in the location. To infer the location of a string of text T, written by a user of gender G, we rank by $P(L|T, G)$.

4. GEO-TAGGED FLICKR DATA

We evaluate on a large scale set of photos, randomly selected from Flickr, containing almost 2.5 million users who have declared their gender. From Table 2, we see that male users are much more likely than female users to geo-tag their photos: over 10% of males in this dataset have geo-tagged photos, compared with fewer than 5% of females. This also leads to a large gender bias in the geo-tagged photo collection: while there are approximately 38% more male users than females in the entire collection, there are more than 3 times more known males than known females with geo-tagged photos. Overall, our geo-tagged collection is made up of photos from over 320K users, 199K of whom have declared their gender, with 151K men and almost 48K women.

After applying a bulk upload filter, our Flickr dataset has over 10 million photos after applying the bulk upload filter. Table 3 shows that, after bulk upload filtering, we have an average of 41.8 photos for each male user and 38.3 photo for each female user. Users who do not declare their gender have an average of 15.8 photos each after bulk upload filtering, indicating that users who take the effort to declare their gender are likely to be much more active those who do not.

For evaluation, the dataset was partitioned into three distinct sets: a training set for building the models (80%, 8M), a tuning set for optimising model parameters (10%, 1M, and a test set for evaluation (10%, 1M). The corpus was partitioned by user, so that users whose photos were used for training the models were not used for testing.

5. EVALUATION

Although we are interested in building models of location that can be used in a number of applications, our evaluation focuses on the task of placing photos. Given our set of test photos, the locations of which are known, we use the models to predict the location of the photos based on the tags, comparing the predicted location with the ground truth.

5.1 Test Indexes

We use the Terrier search engine [9] to index the photo tags from the training set and create their language models, without stopword removal or stemming. To explore the effect of gender-specific language models, different subsets of the training set were used to create the models:

- **Full Model.** A gender-agnostic model is created using all of the data in the training set, a total of over 257K users and over 8 million photos. Photo tags evaluated against this model are ranked by $P(L|T)$ as in Equation 5, with gender information effectively ignored.

- **Female Model.** This gender-specific model is created using data from all of the female users in the training set, a total 38,940 users and over 1.3 million photos.

Method	Acc	MRR	Acc@1	Acc@2	Acc@3	3-hit	5-hit
100km							
Male	0.6178	0.682	0.7017	0.7291	0.7479	0.7261	0.7544
Female	0.5384	0.6019	0.6172	0.6449	0.6697	0.6408	0.6725
Unknown	0.5718	0.6386	0.6626	0.6934	0.714	0.6833	0.7147
Micro Average	0.5962	0.6607	0.6804	0.7084	0.7285	0.704	0.7334
Macro Average	0.576	0.6408	0.6605	0.6892	0.7105	0.6834	0.7139
1km							
Male	0.1842	0.2543	0.3268	0.3814	0.4153	0.2887	0.3328
Female	0.1474	0.2086	0.2778	0.3274	0.3572	0.2385	0.2775
Unknown	0.1549	0.2188	0.2901	0.3409	0.3732	0.248	0.2917
Micro Average	0.1728	0.2403	0.312	0.3651	0.398	0.273	0.3162
Macro Average	0.1622	0.2272	0.2982	0.3499	0.3819	0.2584	0.3007

Table 4: 100km and 1km results for the Full Model

Photo tags tested against this model are ranked by $P(L|T, G = female)$, as in Equation 6.

- **Male Model.** To create the male gender-specific model, we take a random subsample of the male users, of the same size as the female sample (38,940 users, over 1.4 million photos). This model ranks by $P(L|T, G = male)$, as in Equation 6.
- **Combined Model.** The full model will be expected to significantly outperform both the female model and the male data model, as it is built using much more data. For this reason, we create a combined model using data from the same number of users as the Male and Female Models (19,470 of each). Like the Full Model, this model ranks by $P(L|T)$, but without the implicit bias towards male users.

We always use the entire parameter tuning and testing sets. To avoid the inherent bias stemming from the fact that we have many more test photos from men than women, when reporting overall results we use the macro-average in addition to the micro-average. The *macro-average* first calculates an average score for all photos in a category (male, female or unknown) and then takes an average of these, whereas the *micro-average* simply averages over all photos, meaning that the results are dominated by the male users.

Table 1 summarises some statistics for the location models created from 1km and 100km grids. For 100km cells, the total number of unique locations represented in the model is over 10% more for men than women, which is consistent with the 9.6% more photos on average for male users (Table 3). In the 1km models, however, the difference in the number of unique locations is more interesting, with 32% more unique locations in the male model than the female model, and there are many more unique terms (36.6%) used by the males.

Table 1 also shows the number of photos from the test set whose locations are not represented in the training data for the 1km model. The model cannot possibly locate these photos correctly, because they are not represented in the model. The percentage of photos represented tends to be higher for male test photos for all of the models (except the female model, where the coverage is similar across genders), and this is most evident for the male model, where 36.6% of male test photo locations are not represented in the model, compared to 39.58% females photos. When all of the data in the training set is used, over 20% of the test photos are not represented in the 1km model.

5.2 Evaluation Measures

For our main evaluation measure, we use accuracy (*Acc*), the percentage of photos that are located in the correct grid cell. We also report results on a number of other metrics:

- *Mean Reciprocal Rank (MRR).* The mean reciprocal rank is favored over mean rank because it can be interpreted without knowing the number of documents (it always lies between 0 and 1), and is it not severely influenced by target documents retrieved at lower ranks.
- *Accuracy within K Cells (Acc@K).* Accuracy within K cells measures the ability of the model to predict the correct location within K cells of the correct location.
- *H-Hit Rate.* H-Hit rate measures the percentage of photos correctly placed in the top H results in the list ($h = 1$ is equivalent to accuracy)[4].

We tune the model parameters with a grid search on the tuning set, and optimising the accuracy evaluation measure.

5.3 Results

Table 4 shows the results from the Full Model, a gender-agnostic model built using the entire training set. These results show that the photos of male users are consistently located more accurately than those of female users. Male photos can be correctly located within the 100km cell almost 62% of the time (14.7% better than female photos), and 18.4% of the time for 1km cells (25% better than female photos). These results represent models based on all of the data, but it could be argued that the better performance of the male photos is due to the inherent male bias in the data, since there are more about 3 times as many men than women in the training set for this model.

Table 5 show the results for each of the index types for the 1km models (results for the 100km models are omitted, due to lack of space). For each index type, we report the male, female and unknown gender results for that model. So, for example, using the male index and evaluating female photos, we are ranking female photos by the probability of the location given the text and that the gender is male, allowing us to evaluate the relative accuracy of each of the models independently of whether it is used with the 'correct' gender. The gender-specific index evaluates all photos against models of its own gender: male photos against the male model, female photos against the female model, and photos of unknown gender against the combined model. The

Method	Acc	MRR	Acc@1	Acc@2	Acc@3	3-hit	5-hit
Combined Model							
Male	0.1346	0.1916	0.2602	0.3104	0.3421	0.2177	0.2571
Female	0.1152	0.1656	0.2283	0.2752	0.3025	0.1883	0.2223
Unknown	0.1132	0.1652	0.2297	0.2792	0.312	0.1889	0.2243
Micro Average	0.1276	0.1826	0.2491	0.299	0.3301	0.2077	0.2455
Macro Average	0.121	0.1741	0.2394	0.2882	0.3189	0.1983	0.2346
Male Model							
Male	0.14	0.1986	0.2653	0.3157	0.348	0.2272	0.2652
Female	0.1123	0.1621	0.2251	0.2696	0.2982	0.1847	0.2178
Unknown	0.1173	0.1697	0.2346	0.2836	0.3152	0.1934	0.2278
Micro Average	0.1313	0.1874	0.2531	0.3022	0.3338	0.2141	0.2506
Macro Average	0.1232	0.1768	0.2417	0.2896	0.3205	0.2018	0.2369
Female Model							
Male	0.1227	0.1799	0.2448	0.2958	0.3284	0.2066	0.2454
Female	0.1137	0.163	0.2254	0.2742	0.3019	0.1869	0.2202
Unknown	0.1085	0.1583	0.2211	0.2696	0.2999	0.1808	0.2163
Micro Average	0.1187	0.1733	0.2374	0.2877	0.319	0.1988	0.2361
Macro Average	0.1149	0.1671	0.2304	0.2799	0.3101	0.1914	0.2273
Gender-Specific Model							
Male	0.14	0.1986	0.2653	0.3157	0.348	0.2272	0.2652
Female	0.1137	0.163	0.2254	0.2742	0.3019	0.1869	0.2202
Unknown	0.1132	0.1652	0.2297	0.2792	0.312	0.1889	0.2243
Micro Average	0.1309	0.1868	0.2523	0.3022	0.3339	0.2137	0.2504
Macro Average	0.1223	0.1756	0.2401	0.2897	0.3206	0.201	0.2366

Table 5: 1km results for combined, male, female and gender-specific models.

results show that photos created and tagged by male users are always located more accurately than those created and tagged by female users. Even when using a model created exclusively with female data, male photos are located 8% more accurately than female photos. The combined model locates male photos 16.8% more accurately, while the gender specific index locates male photos 23% more accurately.

The relative difference between the best and worst indexes for the 1km model is 7.2%: although we do not present full results here, the equivalent results here, the equivalent different for the 100km models is 1.6%. This suggests that that gender difference in location models seem to be more important at the hyper-local level. In Table 1 we saw that the models built from purely male data had more coverage in terms of the number of unique locations represented, and that this difference is much larger for the 1km models than the 100km. Also, the data used to create the male models contain many more unique terms than the female data: many of these extra terms may be useful for locating photos at a hyper-local level. If we add to this the fact that male test photos are consistently located more accurately than female test photos, regardless of the data used to create the model, we believe that suggests that male users tend to tag their photos with more geographically descriptive terms. We also note that the results for the unknown gender category are closer to the female results than the male results.

5.3.1 Gender-specific models

If the results show clearly that male photos tagged by male users contain more useful location information that those of female users, what they do not show is a clear improvement from using gender-specific indexes. In general, using only the male model gives optimal or close to optimal results.

As noted in Table 1, many of the locations of the photos in the test set are not represented in the data used to create the location models, which means that these photos could not possibly be located correctly by the models. Since the level of unrepresented locations varies from model to model, this factor in itself could be causing some of the gender-based differences in performance. Table 6 presents the results when test photos that are not represented in the model are ignored, showing that just over 22% of male photos and 19% of female photos can be placed in the correct 1km cell if the ground truth location is represented in the location model.

In these results, the gender specific results for male photos (i.e. male photos, male model) are over 6% better than the male results against the female model (0.2208 compared with 0.2076). The gender-specific female results are 2.6% better than the female photos evaluated against the male model. This suggests there is a genuine difference in how men and women describe locations at the hyper-local level represented by 1km cells. The difference in the macro-average scores from Table 6 for the gender-specific and other models, however, is quite minor. We would interpret these results as confirming that there are, in fact, geniune differences in how different men and women describe locations at the hyper-local level. Nevertheless, a combined model including both male and female data is able to locate photos based on their tags almost as well as a gender-specific model.

6. CONCLUSIONS

In this paper we have used geo-tagged Flickr photos to create gender-specific language models of location, with the objective of exploring to what extent gender differences influence such models. Analysis of a very large dataset showed that men are much more likely to geo-tag photos and, particularly at the hyper-local level, take photos in much more unique locations than women: a random sample of photos

Method	Accuracy
Combined Model	
Male	0.2174
Female	0.1913
Micro Average	0.212
Macro Average	0.2043
Male Model	
Male	0.2208
Female	0.1858
Micro Average	0.2137
Macro Average	0.2033
Female Model	
Male	0.2076
Female	0.1906
Micro Average	0.204
Macro Average	0.1991
Gender-Specific Model	
Male	0.2208
Female	0.1906
Micro Average	0.2148
Macro Average	0.2057

Table 6: 1km results when test photo locations not represented in the models are ignored.

from male users contained 36% more unique locations (in terms of 1km cells) than a sample of female users. Our results show that it is always easier to predict the location of male users photos based on their tags compared with female users' photos, a difference that can be as high as 25%. The results also show that, at the hyper-local level, male models perform relatively poorly at locating female photos and vice versa, although the improvement of gender-specific models over gender-agnostic models is relatively minor.

It is also possible that other factors, which we do not control for, are responsible for some if these differences. For example, maybe female users have less geographically diverse images due to lower income, which is turn correlated with gender[3]. In turn, if women are tagging their home location more often, for example, that could cause a difference in the manner in which these locations are being tagged (i.e. home users and visitors might tag the same location differently). It may to interesting, in future work, to also focus separately on how locations are tagged by visitors and by the people who live there, and to see how this correlates with gender.

This work represents a first step in exploring the relationship between demographics and location models based on user generated content, and has shown that there are quantitative differences in how the genders experience and describe locations through their Flickr photos and their tags. In future work we would like to explore other demographic factors with regard to location, in particular age, and other platforms, such as Twitter. Modeling location in a way that takes demographic factors into account can have a variety of applications: for example, the mining of user generated content to recommend tourists itineraries based on geo-tagged media could benefit significantly from demographic information, as it would allow different types of recommendations to be designed for different type of people.

[3]http://en.wikipedia.org/wiki/Male-female_income_disparity_in_the_United_States

7. REFERENCES

[1] S. Argamon, M. Koppel, J. Fine, and A. R. Shimoni. Gender, genre, and writing style in formal written texts. *TEXT*, 23:321–346, 2003.

[2] L. Backstrom, J. Kleinberg, R. Kumar, and J. Novak. Spatial variation in search engine queries. In *WWW '08: Proceeding of the 17th international conference on World Wide Web*, pages 357–366, New York, NY, USA, 2008. ACM.

[3] L. Backstrom, E. Sun, and C. Marlow. Find me if you can: improving geographical prediction with social and spatial proximity. In *WWW '10: Proceedings of the 19th International Conference on World Wide Web*, pages 61–70, New York, NY, USA, 2010. ACM.

[4] L. Chen, B.-G. Hu, L. Zhang, M. Li, and H. Zhang. Face Annotation for Family Photo Album Management. *International Journal of Image and Graphics*, 3(1):81–94, 2003.

[5] Z. Cheng, J. Caverlee, and K. Lee. You Are Where You Tweet: A Content-Based Approach to Geo-locating Twitter Users. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*. ACM, 2010.

[6] D. Crandall, L. Backstrom, D. Huttenlocher, and J. Kleinberg. Mapping the world's photos. In *Proceedings of the 18th International Conference on World Wide Web*, pages 761–770. ACM, 2009.

[7] J. Hays and A. Efros. IM2GPS: estimating geographic information from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition, 2008. CVPR 2008*, pages 1–8, 2008.

[8] N. O'Hare and V. Murdock. Modeling locations with social media. *Information Retrieval*, 2012.

[9] I. Ounis, C. Lioma, C. Macdonald, and V. Plachouras. Research Directions in Terrier: a Search Engine for Advanced Retrieval on the Web. *Novatica/UPGRADE Special Issue on Web Information Access*, 2007.

[10] J. M. Ponte and W. B. Croft. A Language Modeling Approach to Information Retrieval. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '98)*, pages 275–281, Melbourne, Australia, 1998.

[11] A. Popescu and G. Grefenstette. Image Tagging and Search - A Gender Oriented Study. In *WSM2010: 2nd ACM Workshop on Social Media*. ACM, 2010.

[12] A. Popescu and G. Grefenstette. Mining User Home Location and Gender from Flickr Tags. In *Proceedings of the Fourth International Conference on Weblogs and Social Media, ICWSM*. The AAAI Press, 2010.

[13] P. Serdyukov, V. Murdock, and R. van Zwol. Placing Flickr Photos on a Map. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 484–491. ACM, 2009.

[14] I. Weber and C. Castillo. The demographics of web search. In *SIGIR '10: Proceeding of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 523–530, New York, NY, USA, 2010. ACM.