

A Survey of Geo-tagged Multimedia Content Analysis within Flickr

Evangelos Spyrou, Phivos Mylonas

► **To cite this version:**

Evangelos Spyrou, Phivos Mylonas. A Survey of Geo-tagged Multimedia Content Analysis within Flickr. Lazaros Iliadis; Ilias Maglogiannis; Harris Papadopoulos; Spyros Sioutas; Christos Makris. 10th IFIP International Conference on Artificial Intelligence Applications and Innovations (AIAI), Sep 2014, Rhodes, Greece. Springer, IFIP Advances in Information and Communication Technology, AICT-437, pp.126-135, 2014, Artificial Intelligence Applications and Innovations. <10.1007/978-3-662-44722-2_14>. <hal-01391037>

HAL Id: hal-01391037

<https://hal.inria.fr/hal-01391037>

Submitted on 2 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



A survey of geo-tagged multimedia content analysis within Flickr

Evangelos Spyrou¹ and Phivos Mylonas²

¹ Computational Intelligence Laboratory (CIL)
Institute of Informatics and Telecommunications
National Center for Scientific Research Demokritos, Athens, Greece

`espyrou@iit.demokritos.gr`

² Department of Informatics
Ionian University, Corfu, Greece
`fmylonas@ionio.gr`

Abstract. Our survey paper attempts to investigate how recent and undoubted emerge in enriched, geo-tagged social networks' multimedia content sharing works to the benefit of their users and whether it could be handled in a formal way, in order to capture the meaningful semantics rising from this newly introduced user experience. It further specializes its focus by providing an overview of current state-of-the-art techniques with respect to geo-tagged content access, processing and manipulation within the popular Flickr social network. In this manner it explores the role of information retrieval, integration and extraction from the technical point of view, coupled together with human social network activities, like, for instance, localization and recommendations based on pre-processed collaborative geo-tagged photos, resulting into more efficient, optimized search results.

Keywords: Flickr, tags, photos, multimedia, social networks

1 Introduction

Current digital era is characterized by a single, yet very important observation: an extremely large amount of digital multimedia content is shared online every moment by people interacting within the so-called “social networks”. This online social networking explosion shows no signs of abating, with almost twice as many Internet users having an online social profile than two years ago, helping to make Facebook³ the most viewed website in the European Union, according to recent research efforts. On top of that, more people are using the Internet to create their own multimedia content than ever before, with 73% of online users having a social networking profile [1], compared with 37% in 2008 [2].

Flickr⁴ is an image and video hosting website created by a Vancouver-based company named “Ludicorp” back in 2004 and currently owned by Yahoo!⁵. What

³ <http://www.facebook.com>

⁴ <http://www.flickr.com>

⁵ <http://www.yahoo.com>

makes it special among other social networks is its aspect as an online community, within which users are able to interact by sharing comments about photography and create groups of particular interests. The Verge reported in March 2013 that “Flickr had a total of 87 million registered members and more than 3.5 million new photos uploaded daily”⁶. Each photo may contain metadata added by its photographer, such as tags that describe either its visual content or location, or a free text description. It also contains metadata added by the camera that has been used, such as date taken, camera settings, camera model, etc. Few Global Positioning System (GPS) enhanced cameras automatically geo-tag the photos they take, but in principal this is done manually, by the photographer. The vast majority of images uploaded to Flickr are taken by common users or amateur photographers. The textual metadata associated with the image often serves as a reminder of the context of the image for the photographer and his social circle [3], [4].

In principle, every part of a photo may be tied to a geographic location, but in most typical applications, only the position of the photographer is associated with the entire digital photo. As the reader may imagine, this small detail implicates and significantly burdens most multimedia content search and retrieval tasks. In the most typical example, photos of a landmark may have been taken from very different positions apart and in order to identify all photos of this particular landmark within an image database, all photos taken within a reasonable circular distance from it must be considered. Now, when such geo-tagged photos are uploaded to online multimedia content sharing communities, such as Flickr, Panoramio⁷ or Instagram⁸, that enable the construction of infinite connections among their users [5], a photo can be placed onto a map to view the location the photo was taken. In this way, social network users can browse photos from a map, search for photos from a given area, and find related photos of the same place from other users; these tasks are considered elementary in order to build additional, ad-hoc value-added digital services on top, like automated route/trip planning or like, to our most recent knowledge, the popular “NOW” app; the latter uses geo-tagged Instagram photos to find nearby events happening now⁹.

The act of automatically providing or calculating meaningful photo’s geo-tags (the so-called “geo-tagging” process) opens a huge research topic for the researchers’ community, mainly to the direction of being able to analyze them, to identify and determine social patterns amongst them. However, issues of credibility on the volunteered user-generated geo-tagging should become of broader research interest in various areas [6], [7], motivating us to further investigate this topic in the following, focusing on the popular Flickr social network. At this point, it should be noted that our work differentiates from previous similar surveys, since they emphasized either on geotagged content without focusing

⁶ http://en.wikipedia.org/wiki/Flickr#cite_note-4

⁷ <http://www.panoramio.com>

⁸ <http://www.instagram.com>

⁹ <http://techcrunch.com/2013/01/11/now-app/>

specifically on user generated and manually geotagged photos of Flickr [8], or in social media in general [9].

2 Multimedia Content Retrieval

The very first research community dealing with up-to-date multimedia community research challenges is the one depicted by Information Retrieval in general and multimedia content retrieval in particular. Since Flickr is mainly a photo sharing website, the fact that it attracted the interest of the image retrieval community is considered to be rather natural. The main approach followed by researchers is to use either textual metadata or visual properties of photos and often combine them in an effort to improve the accuracy of their respective algorithms. As depicted in the following, such research efforts vary from textual ones, to ones based on low level visual or even hybrid characteristics.

2.1 Aren't tags text, after all?

As discussed, the first approach to tackle the problem at hand is based solely on text retrieval, a branch of information retrieval where information is manipulated primarily in the form of text, having each photo represented solely by its textual features, i.e. the manually generated tags. Abbasi et al. [10] identified landmarks using tags and Flickr groups, without exploiting geospatial information. They used SVM classifiers trained on thematic Flickr groups, in order to find relevant landmark-related tags. Ahern et al. [11] analyzed tags associated with geo-referenced Flickr images so as to generate knowledge. This knowledge was a set of the most “representative” tags for an area. They used a TF-IDF approach and presented a visualization tool, namely the *World Explorer*, which allowed users explore their results. Serdyukov et al. [12] adopted a language model which lies on the user collected Flickr metadata and aimed to annotate an image based on these metadata. Their goal was to place photos on a map, i.e. provide an automatic alternative to manual geo-tagging. Venetis et al. [13] examined techniques to create a “tag-cloud”, i.e. a set of terms/tags able to provide a brief yet rich description of a large set of terms/tags. They presented and defined certain user models, metrics and algorithms aiming at this goal. Lerman et al. [14] aimed to personalize text-based search results by adding information about users' relations. Finally, Larson et al. [15] tried to detect whether tags correspond to physical objects, and also the scale of these objects, using a natural language approach.

2.2 Shall we consider visual characteristics?

The second approach focuses on the visual aspects of multimedia content analysis. Research efforts in this area discard textual annotations and focus on low-level visual features. Wang et al. [16] proposed a training algorithm and applied it on the problem of image similarity. They worked on a Flickr data set, under the

assumption that two images are considered similar if they belong to the same group. Chatzilari et al. [17] used region level annotations and visual features, in an effort to recognize objects with a semi-supervised approach. They started from a set of Flickr photos that contain the same object. Philbin and Zisserman [18] created a graph based on visual features and tried to group similar Flickr photos, from a corpus of 1M photos. Avrithis et al. [19] retrieved similar Flickr photos by using a 2-level clustering processing, both by means of geo-tags and visual features. Yanai et al. [20] focused on the relationship between words and locations. They used visual features and tried to associate them with certain locations, using an entropy based approach. Li et al. [21] used SVMs trained on visual features to classify a 30M data set. They observed that by incorporating temporal information, the accuracy of the results was significantly improved. Joshi and Luo [22] used visual detectors and incorporate bags-of-geotags within a probabilistic framework, in order to detect activities and events in photos. Yu and Luo [23] combined visual context with location information in order to detect concepts in photos. Luo et al. [24] fused information extracted from both a Flickr data set and a set of satellite images, in order to detect events. Batko et al. [25] used MPEG-7 visual features and search into a set of over 50M photos from Flickr. Seah et al. [26] created visual summaries on the results of visual queries on a data set of Flickr images that in contrast to previous works, e.g., the one of [27], they attempted to generate concept-preserving summaries. Finally, Liu et al. [28] incorporated the social aspect of photos, in order to re-rank search results more according to both social and visual relevances.

2.3 A little bit of both - the hybrid approach!

Since the visual content of images may provide a powerful description, many research efforts try to combine visual descriptions with textual metadata. Barrios et al. [29] presented an image retrieval system that combines textual and visual content. They downloaded and stored locally images from Flickr and used simple color and texture visual descriptors, along with the title, description and tags, for each image. Crandall et al [30] used visual, temporal and geospatial information to automatically identify places and/or events in city and landmark level. They also added temporal metadata information to improve classification performance. With the same motivation, Quack et al. [31] divided the area of interest into non-overlapping, square tiles, then extracted and used visual, textual and geospatial features. They handled tags by a modified TF-IDF ranking and linked their results to Wikipedia¹⁰. Gammeter et al. [32] overlaid a geospatial grid over earth and matched pairwise retrieved photos of each tile using visual features. Then they clustered photos into groups of images depicting the same scene. The metadata were used to label these clusters automatically, using a TF-IDF scheme. Moëllic et al [27] aimed to extract meaningful and representative clusters from large-scale image collections. They proposed a method based on a shared nearest neighbors approach that treats both visual features and tags. Li

¹⁰ <http://www.wikipedia.org>

et al [33] proposed an algorithm that learns tag relevance by voting from visually similar neighbors. They did not use geospatial data, nor limited their approach on landmarks/places of interest and aimed to retrieve semantically similar images. Moxley et al. [34] classified mined geo-referenced tags as places, by extending [35], landmarks by clustering image datasets considering mutual information and prior knowledge from Wikipedia and visual terms using the mutual information between visual descriptors and tags. Ulges et al. [36] adopted a context- based approach, assuming that users place semantically similar photos in Flickr groups. Fan et al. [37] proposed a system, namely *JustClick* which exploits both visual and textual information and after a search and retrieval process, it recommends photos using an interactive interface. Simon et al. [38] created visual summaries of large image data set based mainly on visual features, but also exploiting tags. Kennedy and Naaman [39] used visual features and tags, in order to extract the most representative tags and views for landmarks, working on a corpus of 110K Flickr photos from San Fransisco. Finally, Liu et al. [40] were the first to consider user uploading patterns, geotagging behaviors, and the relationship between the temporal and the spatial gap of two photos from the same user.

3 Automatic Tag/Geo-tag Generation

In a slightly different approach, special attention has been given to methods exploiting the automatic generation of tags, a process often called “(tag-)recommendation”, as well as the prediction of geo-tags, i.e., of the geographic coordinates where a photo has been taken, a process often referred to as “localization”. In the following, we briefly present both approaches, summarizing most important research works in the fields.

3.1 Tag Recommendation

Initially and as expected, tag recommendation approaches often adopt traditional tag processing techniques. In this manner, Chen et al. [41] proposed a system that automatically recommends tags for photos and also for adding photos into appropriate popular groups. For the latter case, they used SVM predictors in order to identify concepts and used these results so as to search for groups. Then, they used these groups to harvest more tags and attach them to their photos. Anderson et al. [42] presented a system, namely *TagEz* which combined both textual and visual features, so as to recommend tags. Their results indicated that the use of textual metadata outperformed both visual and combined features. Chaundry et al. [43] presented an approach for tag assignment to geographic areas, using a TF-IDF scheme and logistic regression, for various levels of detail. Hsieh and Hsu [44] exploited visual similarity and after a tag expansion process, aim to automatically annotate photos. Kennedy et al. [45] selected representative tags from urban areas using a multimodal approach. Their results indicate that the use of visual features can drastically improve precision. Sigurbjörnsson and Van Zwol [46] extracted tag co-occurrence statistics

and tag aggregation algorithms, in order to recommend tags by investigating and evaluating four different strategies. Furthermore, they introduced a “promotion” function, whose role was to promote the most descriptive tags. Garg and Weber [47], [48] presented a system that while users tagged their photos, it dynamically suggested related tags by considering similar groups to user’s preferences. Moxley et al [49] presented *SpiritTagger* tool, in order to recommend tags for Flickr photos of urban regions, which is unaware of the user’s tags and lies on visual properties and geographic distance, in order to select similar photos. Popescu and Moëllic [50] presented *Monuanno*, a system that uses visual features to automatically annotate georeferenced landmark images. Kleban et al. [51] presented a world scale system for tag recommendation, based on geotags and visual features. Finally, Chen and Shin [52] used both textual and social features of tags and a machine learning approach, in order to extract representative tags that can be related to the users favorite topics.

3.2 Content Localization

On the other hand, automatic geo-tag generation has gained huge research interest, mainly due to the vast available Flickr database of geo-tagged photos. Kelm et al. [53] adopted a hierarchical approach and tried to automatically predict geo-tags for Flickr videos. Their technique lies on both textual and visual features and also uses external resources, such as Geonames¹¹ and Wikipedia. Van Laere et al. [54] trained naive Bayes classifiers at different spatial resolutions. They used only textual features and worked at various spatial resolutions, for a set of 55 european cities. De Rouck et al [55] used language probabilistic models that have been trained on Flickr photos, in order to geo-tag Wikipedia pages. Their approach outperformed Yahoo! Placemaker¹² and their results indicated that the increasing growth of tagged content in Flickr would continuously improve their accuracy. Friedland et al. [56] combined textual and visual features and worked on a the MediaEval 2010 data set. They concluded that solely visual information proves inadequate for accurate geo-localization, but when combined with textual it can assist on the improvement of the accuracy. Hauff and Houben [57] added information considering user’s activities in Twitter¹³. However, even if their results were promising, the median location error was still far from usable. Van Laere et al. [58] divided areas into disjoint regions and then used statistics and a Naive Bayes classifier. Van Laere et al. [59] proposed a tag-based approach that uses language models and similarity search, in order to estimate geo-tags based on a training set. Friedland et al. [60] worked on Flickr videos and used both textual and visual metadata. Their results may seem poor, however they were superior to all other contributions of MediaEval 2010¹⁴. Kalantidis et al. [61] presented *Viral*, a system that aims to localize photos uploads by users,

¹¹ <http://www.geonames.org/>

¹² <http://www.programmableweb.com/api/yahoo-placemaker>

¹³ <http://www.twitter.com>

¹⁴ <http://www.multimediaeval.org/mediaeval2010/>

based on the visual similarity to geo-tagged images. They used a database of more than 2M photos taken from 40 cities. In previous work [62] we suggested a probabilistic framework which aimed to place Flickr data on a map based on their tags.

Joshi et al. [63] proposed a probabilistic framework for tag-based localization. Their work was extended by Gallagher et al. [64] who used a large geotagged corpus from Flickr, extracted several visual features and used location probability maps for tags. They integrated them and tried to localize photos. Hays and Efros [65] present *IMGPS*, a system for image localization using visual features. It should be noted that they provided a probability distribution over the Earth. Kalogerakis et al. [66] extended this work by adding temporal information, in an effort to extract information about image sequences. O’ Hare and Murdock [67] presented a statistical language modeling approach, in order to identifying locations in arbitrary text. They investigated several ways to estimate models, based on the term and the user frequencies. To this goal. they used a set of public, geo-tagged photos in Flickr as ground truth. Hare et al. [68] estimated a continuous probability density function (PDF) over the Earth and combined textual with a number of weighted visual features. Their approach on tags differs from the others as they do not filter any of the tags, but they rather use them for evidence, i.e. certain words may be associated with certain countries. Finally, Li et al [69] removed “noisy” photos, i.e. photos that cannot contribute sufficiently to location estimation. They extracted both local and global features and instead of using the whole dataset, they performed clustering and use the resulting centroids, instead.

4 Conclusions

In this position paper we attempted to conduct a detailed survey and provide a brief summarization of current state-of-the-art techniques regarding geo-tagged Flickr content access, processing and manipulation issues. In this manner, we explored related research efforts mostly focused on information retrieval tasks. Our intention was to identify the trends in the surveyed area and organize them in a novel way that would integrate and add understanding to the work in the field with respect to the Flickr social network, so as for fellow researchers to be able to seek and reference related information efficiently. Among our future work is the extension of this survey to other Flickr application domains, other popular social networks and even other content types, such as text snippets according to the social network under investigation.

References

1. Dugan, M., Smith, A.: Social media update 2013. <http://pewinternet.org/Reports/2013/Social-Media-Update.aspx> (2013)
2. Lenhart, A., Purcell, K., Smith, A., Zickuhr, K.: Social media & mobile internet use among teens and young adults. <http://pewinternet.org/Reports/2010/Social-Media-and-Young-Adults.aspx> (2010)

3. Nov, O., Naaman, M., Ye, C.: Analysis of participation in an online photo-sharing community: A multidimensional perspective. *Journal of the American Society for Information Science and Technology* **61**(3) (2010) 555–566
4. Van House, N.A.: Flickr and public image-sharing: distant closeness and photo exhibition. In: *Proc. of ACM CHI*. (2007)
5. Van Dijck, J.: Flickr and the culture of connectivity: Sharing views, experiences, memories. *Memory Studies* **4**(4) (2011) 401–415
6. Flanagin, A.J., Metzger, M.J.: The credibility of volunteered geographic information. *GeoJournal* **72**(3-4) (2008) 137–148
7. Spielman, S.E.: Spatial collective intelligence? credibility, accuracy, and volunteered geographic information. *Cartography and Geographic Information Science* (ahead-of-print) (2014) 1–10
8. Luo, J., Joshi, D., Yu, J., Gallagher, A.: Geotagging in multimedia and computer vision: a survey. *Multimedia Tools and Applications* **51**(1) (2011) 187–211
9. Wu, J., Sun, H., Tan, Y.: Social media research: A review. *Journal of Systems Science and Systems Engineering* **22**(3) (2013) 257–282
10. Abbasi, R., Chernov, S., Nejd, W., Paiu, R., Staab, S.: Exploiting flickr tags and groups for finding landmark photos. In: *Advances in Information Retrieval*. Springer (2009) 654–661
11. Ahern, S., Naaman, M., Nair, R., Yang, J.H.I.: World explorer: visualizing aggregate data from unstructured text in geo-referenced collections. In: *Proc. of the ACM/IEEE-CS JCDL*. (2007)
12. Serdyukov, P., Murdock, V., Van Zwol, R.: Placing flickr photos on a map. In: *Proc. of ACM SIGIR*. (2009)
13. Venetis, P., Koutrika, G., Garcia-Molina, H.: On the selection of tags for tag clouds. In: *Proc. of ACM WSDM*. (2011)
14. Lerman, K., Plangprasopchok, A., Wong, C.: Personalizing image search results on flickr. (2007)
15. Larson, M., Kofler, C., Hanjalic, A.: Reading between the tags to predict real-world size-class for visually depicted objects in images. In: *Proc. of ACM MM*, ACM (2011) 273–282
16. Wang, G., Hoiem, D., Forsyth, D.: Learning image similarity from flickr groups using stochastic intersection kernel machines. In: *Proc. of IEEE ICCV*. (2009)
17. Chatzilari, E., Nikolopoulos, S., Papadopoulos, S., Zigkolis, C., Kompatsiaris, Y.: Semi-supervised object recognition using flickr images. In: *Proc. of IEEE CBMI*. (2011)
18. Philbin, J., Zisserman, A.: Object mining using a matching graph on very large image collections. In: *Proc. of IEEE ICVGIP*. (2008)
19. Avrithis, Y., Kalantidis, Y., Toliás, G., Spyrou, E.: Retrieving landmark and non-landmark images from community photo collections. In: *Proc. of ACM MM*. (2010)
20. Yanai, K., Kawakubo, H., Qiu, B.: A visual analysis of the relationship between word concepts and geographical locations. In: *Proc. of ACM ICIVP*. (2009)
21. Li, Y., Crandall, D.J., Huttenlocher, D.P.: Landmark classification in large-scale image collections. In: *Proc. of IEEE ICCV*. (2009)
22. Joshi, D., Luo, J.: Inferring generic activities and events from image content and bags of geo-tags. In: *Proc. of ACM CBIVR*. (2008)
23. Yu, J., Luo, J.: Leveraging probabilistic season and location context models for scene understanding. In: *Proc. of ACM CBIVR*. (2008)
24. Luo, J., Yu, J., Joshi, D., Hao, W.: Event recognition: viewing the world with a third eye. In: *Proc. of ACM MM*. (2008)

25. Batko, M., Falchi, F., Lucchese, C., Novak, D., Perego, R., Rabitti, F., Sedmidubsky, J., Zezula, P.: Building a web-scale image similarity search system. *Multimedia Tools and Applications* **47**(3) (2010) 599–629
26. Seah, B.S., Bhowmick, S.S., Sun, A.: Summarizing social image search results. In: *Proc. of ACM WWW Companion*. (2014)
27. Moëllic, P.A., Haugeard, J.E., Pitel, G.: Image clustering based on a shared nearest neighbors approach for tagged collections. In: *Proc. of ACM CIVR*. (2008)
28. Liu, S., Cui, P., Luan, H., Zhu, W., Yang, S., Tian, Q.: Social-oriented visual image search. *Computer Vision and Image Understanding* **118** (2014) 30–39
29. Barrios, J.M., Diaz-Espinoza, D., Bustos, B.: Text-based and content-based image retrieval on flickr: Demo
30. Crandall, D.J., Backstrom, L., Huttenlocher, D., Kleinberg, J.: Mapping the world’s photos. In: *Proc. of ACM WWW*. (2009)
31. Quack, T., Leibe, B., Van Gool, L.: World-scale mining of objects and events from community photo collections. In: *Proc. of ACM CIVR*. (2008)
32. Gammeter, S., Bossard, L., Quack, T., Gool, L.V.: I know what you did last summer: object-level auto-annotation of holiday snaps. In: *Proc. of IEEE ICCV*. (2009)
33. Li, X., Snoek, C.G., Worring, M.: Learning tag relevance by neighbor voting for social image retrieval. In: *Proc of ACM MIR*. (2008)
34. Moxley, E., Kleban, J., Xu, J., Manjunath, B.: Not all tags are created equal: Learning flickr tag semantics for global annotation. In: *Proc. of IEEE ICME*. (2009)
35. Rattenbury, T., Good, N., Naaman, M.: Towards automatic extraction of event and place semantics from flickr tags. In: *Proc. of ACM SIGIR*. (2007)
36. Ulges, A., Worring, M., Breuel, T.: Learning visual contexts for image annotation from flickr groups. *Multimedia, IEEE Trans. on* **13**(2) (2011) 330–341
37. Fan, J., Keim, D.A., Gao, Y., Luo, H., Li, Z.: Justclick: Personalized image recommendation via exploratory search from large-scale flickr images. *IEEE Trans. on Circuits and Systems for Video Technology* **19**(2) (2009) 273–288
38. Simon, I., Snavely, N., Seitz, S.M.: Scene summarization for online image collections. In: *Proc. of ICCV*. (2007)
39. Kennedy, L.S., Naaman, M.: Generating diverse and representative image search results for landmarks. In: *Proc. of ACM WWW*. (2008)
40. Liu, B., Yuan, Q., Cong, G., Xu, D.: Where your photo is taken: Geolocation prediction for social images. *Journal of the Association for Information Science and Technology* (2014)
41. Chen, H.M., Chang, M.H., Chang, P.C., Tien, M.C., Hsu, W.H., Wu, J.L.: Sheepdog: group and tag recommendation for flickr photos by automatic search-based learning. In: *Proc. of ACM MM*. (2008)
42. Anderson, A., Ranghunathan, K., Vogel, A.: Tagez: Flickr tag recommendation. *Association for the Advancement of Artificial Intelligence* (2008)
43. Chaudhry, O., Mackaness, W.: Automated extraction and geographical structuring of flickr tags. In: *Proc. of GEOProcessing*. (2012)
44. Hsieh, L.C., Hsu, W.H.: Search-based automatic image annotation via flickr photos using tag expansion. In: *Proc. of IEEE ICASSP*. (2010)
45. Kennedy, L., Naaman, M., Ahern, S., Nair, R., Rattenbury, T.: How flickr helps us make sense of the world: context and content in community-contributed media collections. In: *Proc. of ACM MM*. (2007)
46. Sigurbjörnsson, B., Van Zwol, R.: Flickr tag recommendation based on collective knowledge. In: *Proc. of ACM WWW*. (2008)

47. Garg, N., Weber, I.: Personalized tag suggestion for flickr. In: Proc. of ACM WWW. (2008)
48. Garg, N., Weber, I.: Personalized, interactive tag recommendation for flickr. In: Proc. of ACM RecSys. (2008)
49. Moxley, E., Kleban, J., Manjunath, B.: Spirittagger: a geo-aware tag suggestion tool mined from flickr. In: Proc. of ACM MIR. (2008)
50. Popescu, A., Moëllic, P.A.: Monuanno: automatic annotation of georeferenced landmarks images. In: Proc. of ACM CIVR. (2009)
51. Kleban, J., Moxley, E., Xu, J., Manjunath, B.: Global annotation on georeferenced photographs. In: Proc. of ACM CIVR. (2009)
52. Chen, X., Shin, H.: Tag recommendation by machine learning with textual and social features. *Journal of Intelligent Information Systems* **40**(2) (2013) 261–282
53. Kelm, P., Schmiedeke, S., Sikora, T.: A hierarchical, multi-modal approach for placing videos on the map using millions of flickr photographs. In: Proc. of ACM SBNMA. (2011)
54. Van Laere, O., Schockaert, S., Dhoedt, B.: Combining multi-resolution evidence for georeferencing flickr images. In: Scalable Uncertainty Management. Springer (2010) 347–360
55. De Rouck, C., Van Laere, O., Schockaert, S., Dhoedt, B.: Georeferencing wikipedia pages using language models from flickr. (2011)
56. Friedland, G., Choi, J., Lei, H., Janin, A.: Multimodal location estimation on flickr videos. In: Proc. of ACM WSM. (2011)
57. Hauff, C., Houben, G.J.: Geo-location estimation of flickr images: social web based enrichment. In: Adv. in Information Retrieval. Springer (2012) 85–96
58. Van Laere, O., Schockaert, S., Dhoedt, B.: Towards automated georeferencing of flickr photos. In: Proc. of ACM GIR. (2010)
59. Van Laere, O., Schockaert, S., Dhoedt, B.: Finding locations of flickr resources using language models and similarity search. In: Proc. of ACM ICMR. (2011)
60. Friedland, G., Choi, J., Janin, A.: Video2gps: a demo of multimodal location estimation on flickr videos. In: Proc. of ACM MM. (2011)
61. Kalantidis, Y., Tolias, G., Avrithis, Y., Phinikettos, M., Spyrou, E., Mylonas, P., Kollias, S.: Viral: Visual image retrieval and localization. *Multimedia Tools and Applications* **51**(2) (2011) 555–592
62. Spyrou, E., Mylonas, P.: Placing user-generated photo metadata on a map. In: Proc. of IEEE SMAP. (2011)
63. Joshi, D., Gallagher, A., Yu, J., Luo, J.: Exploring user image tags for geo-location inference. In: Proc. of IEEE ICASSP. (2010)
64. Gallagher, A., Joshi, D., Yu, J., Luo, J.: Geo-location inference from image content and user tags. In: Proc. of IEEE CVPR. (2009)
65. Hays, J., Efros, A.A.: Im2gps: estimating geographic information from a single image. In: Proc. of IEEE CVPR. (2008)
66. Kalogerakis, E., Vesselova, O., Hays, J., Efros, A.A., Hertzmann, A.: Image sequence geolocation with human travel priors. In: Proc. of IEEE ICCV. (2009)
67. OHare, N., Murdock, V.: Modeling locations with social media. *Information retrieval* **16**(1) (2013) 30–62
68. Hare, J., Davies, J., Samangooei, S., Lewis, P.H.: Placing photos with a multimodal probability density function
69. Li, J., Qian, X., Tang, Y.Y., Yang, L., Mei, T.: Gps estimation for places of interest from social users' uploaded photos. *IEEE Transactions on Multimedia* **15**(8) (2013) 2058–2071