2014 9th International Workshop on Semantic and Social Media Adaptation and Personalization

An overview of Flickr challenges and research opportunities

Evaggelos Spyrou
Computational Intelligence Laboratory (CIL)
Institute of Informatics and Telecommunications
National Center for Scientific Research Demokritos,
Athens, Greece
E-mail: espyrou@iit.demokritos.gr

Phivos Mylonas
Department of Informatics
Ionian University,
Corfu, Greece
E-mail: fmylonas@ionio.gr

***Abstract*—It is rather true that nowadays diversity and a potential for growth characterize the field of popular social networks like Flickr. Based on this observation, we provide herein a brief overview of the basic characteristics, challenges and research opportunities of the latter. Some historical facts of its origin are presented, followed by a review of its digital photo sharing and tagging capabilities, focusing on the most popular related research fields, like travel applications, knowledge extraction and human activity tracking. Throughout the paper the status of current research and directions for future research are highlighted in an effort to provide useful information to fellow researchers.**

I. Motivation and Introduction

A social network is in principle considered to be a digital, on-line place where people could create profiles and build a personal network that connects them to other people. During recent years, such networks have rocketed from a niche activity into a phenomenon that engages tens of millions of Internet users every day. The growth in the popularity of these virtual networks has been coupled together with extremely large amounts of digital multimedia content shared on-line every moment, by people interacting within these social networks. The biggest majority of such multimedia content uploaded and shared to one of the most popular social networks, Flickr[1], are taken by so-called everyday users or even amateur "event capturers". Moreover, the fact that personal multimedia content has become easy to grasp and capture by rather cheap hardware and, as Van House [1] empirically identified, is being used according to four basic uses: (a) memory, narrative and identity; (b)relationships; (c) self-representation; and (d) self-expression, aided to its mass market access, however, resulting into other types of problems. For instance, the fact that this type of user-generated multimedia content is difficult to efficiently get access to, or to be processed and effectively manipulated by people in a meaningful amount of time and effort spent; tasks that become day-by-day very difficult and challenging to tackle.

So, it is common research knowledge that the overwhelming distribution of such dynamically generated content over its end users, online communities and devices requires ways for efficient representation and organization, in order to be exploited in applications and services. In their empirical study, Zeng and Wei [2] investigated potential relationships between social ties and similarities of the type of digital content that people create online. Among their observations, they found that around the time that a social tie between two individuals is formatted, they begin to create more similar content compared to what they had created before. Interestingly, this similarity tends to evolve in different ways when observing different subgroups of user pairs.

In this framework, traditional analysis approaches of analyzing data in terms of objects, and/or concepts and other isolated entities are often quite insufficient, since they do not take into account important properties of online shared multimedia content, such as spatial or other related types of information – the so-called ''content metadata''. Although in general, the application of qualitative and quantitative multimedia content analysis techniques to assess metadata records goes back in time and does not advance current research state-of-the-art, it is the nature of today's framework conditions that point out the special interesting role of *geo-tags* in the process, that makes the difference. Geo-tags are considered important for online multimedia analysis and annotation – in the dynamic environment of a social network, human behavior and activities are better described and exploited in terms of enriched content metadata. We should note herein that the roots of this annotation process lie within the analog photo era, where users wrote some "metadata" information such as place and date, behind photos.

In principle, geotagging is the process of adding geographical identification metadata to various media such as a geotagged digital photos and consists typically of latitude and longitude coordinates. When tagging a photo, e.g., containing a landmark with the location it has been taken (i.e. "geo-tagging"), there are two ways that may be considered as "correct": a) to geo-tag according to the actual location of the landmark depicted; and b) to geo-tag based on the location of the photographer. Semantically, the first approach appears more meaningful, however it is not trivial how to tag when a photo contains more landmarks. To understand this, one should consider e.g. the examples of panoramic photos depicting a whole area from above. The latter approach is the one that is actually followed by GPS-enabled cameras and mobile phones. It should be obvious that this small detail implicates and significantly burdens most multimedia content search and retrieval tasks. So, we could imagine that in the context of

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

978-1-4799-6814-5/14 $31.00 © 2014 IEEE
DOI 10.1109/SMAP.2014.19

a large photo sharing website, such as Flickr, photos of a landmark may have been taken from very different positions apart and in order to identify the majority of the photos of this particular landmark, one should consider all photos taken within a reasonable circular distance.

When such geo-tagged photos are uploaded to online multimedia content sharing communities, such as Flickr, Panoramio[2] or Instagram[3], that enable the construction of infinite connections among their users [3], typically a photo may be placed onto a map to mark this way 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 geotagged Instagram photos to find nearby events happening "now"[4].

It is rather obvious that geotagging may help users find a wide variety of location-specific information, whereas geotagging-enabled information services can also potentially be used to find location-based multimedia resources. Still, credibility issues on user-generated geotagging should become of broader research interest in various areas [4], [5], motivating us to further focus on this topic in the following, within the framework of the popular Flickr social network. In addition, quite on the contrary to most recent previous analogue surveys on the topic, such as [6] or [7], that tackle geotagged content in general without focusing on its user-generated aspect, our work focuses solely on user-generated digital photos shared within Flickr. In the following sections we shall briefly present some fundamental social network characteristics, as well as a selection of the most interesting research opportunities that arise in the field.

II. The Flickr Website and its Social Aspects

A. A brief presentation of Flickr

Flickr was initially launched by a small company named 'Ludicorp', in 2004 and was immediately acquired by Yahoo![5] only a year later. At first sight, one would describe it as an image and video hosting website. However, by a deeper look, one should notice its aspect as an online virtual community. Within this community, users are able to interact among others in the following ways: a) they can share comments with other users; b) they can follow other users; and c) they are able to create groups, whose members share the same interests. This facilitation of intra-user communication has been suggested as the main originality of Flickr [8]. To be more accurate, Flickr focused on the communication among amateur photographers, which although represent the minority in terms of their population within the community, they play the most important role, as they are the main content producers, they tend to socialize and encourage new activities.

More than 14 billion of images have been currently uploaded to Flickr. Most of them contain user generated textual and location metadata. The first, typically aim to provide a free, non-structured description of its visual content and/or location, while the latter the location the photo has been taken, in terms of its geographic coordinates. Also they may contain automatically generated (i.e. by the camera) metadata, 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, i.e. by the photographer and in this case the accuracy is not guaranteed. This way, Flickr has become the largest collection of community collected geo-tagged photos.

B. Flickr Sharing and Tagging Photos

The act of adding descriptive keywords to photos is generally denoted as "tagging". It is notable that the vast majority of uploaded photos on Flickr has been tagged with a few descriptive keywords, although these may differ significantly among "consumers" and "producers" of digital content [9]. It also has been generally accepted that certain tasks associated with multimedia data management may benefit by a certain amount of metadata [10]. But why do people upload in common view and moreover tag their photos? Angus and Thelwall [11] investigated the main motivations for using Flickr and for tagging photos. Their sample users originated from the US and Denmark and their research indicated that users employ Flickr both as a personal archive and also as a means of sharing photos with friends and family. However, reasons for tagging lie in the social organization and communication possibilities that Flickr offers. Flickr is thus, treated more as a virtual community and less as a website for commercial gains.

Moreover, Flickr users often interact on purpose, by commenting and asking questions about uploaded photos, as it has been shown by Canningham and Mahui [12]. Additionally, the research of Nov et al. [13] indicated that self and public motivations for tagging are highly correlated with the quantity of tags generated by a user. These motivations include an organizational and a communicational aspect, with the first dealing with the facilitation of photo retrieval and the latter with providing the best detailed description to other users. Finally, Ames and Naaman [14] observed that users may often be encouraged to tag by external factors such as point-of-capture annotation (e.g., on the mobile device) or by tag suggestions. They concluded that it is more likely that users would annotate their photos when they are given certain motivations and affordances.

The procedure of tagging photos within Flickr is towards the aforementioned way. Users are able to tag their photos as part of the uploading process or when they view them. An important question herein deals with the accuracy of user provided textual metadata on Flickr. To this goal, Winget [15] made a study towards the "correct" way of organizing information on the web, with case studies derived from Flickr. Her main conclusions were that users intend on providing accurate textual descriptions, however since these appear rather arbitrary, there is the need of organizing them with structured vocabularies. Moreover, in their survey, Wang et al. [16] concluded that the main open challenges in tagging are the involvement of a large number of humans in the process and the automatic (i.e. generated by a computer) tagging in large scale collections.

[2]http://www.panoramio.com
[3]http://www.instagram.com
[4]http://techcrunch.com/2013/01/11/now-app/
[5]http://www.yahoo.com

III. Research Opportunities on Flickr

Flickr has been very popular around the research community during the last few years both for being the largest collection of community collected geo-tagged photos and for offering a public API[6] for accessing these photos along with their textual metadata. Thus, the majority of research on community collected photo metadata and geo-data uses part of its database as a continuously growing test-bench, whose size is larger than the one that the majority of state-of-the-art algorithms is able to handle.

Encouraging cooperation between research fellows and users from Flickr, by promoting the opportunities offered by the latter towards the efficient access, processing and effective manipulation of digital multimedia content offered within Flickr includes notably the setting up of related scientific challenges, but it can also involve the organization of awareness raising and information dissemination activities. Still, the tasks of semantic characterization, organization and efficient exploitation of user-generated multimedia content towards their meaningful exploitation remain of great importance within recent research community efforts.

Since Flickr is mainly a photo sharing website, the fact that it initially attracted the interest of the image retrieval community is considered to be rather expected. Research efforts focused on:

1) the use of either *textual* metadata, i.e. approaches based on text retrieval; information is manipulated primarily in the form of text, having each photo represented solely by its manually generated tags.
2) the *visual* aspects of multimedia content analysis; textual annotations are discarded and information is manipulated having each photo represented solely by its low-level visual features.
3) the *combination* q of visual descriptions with textual metadata; both descriptions are fused, in an effort to improve the accuracy.

With different motivations, special attention has been given to methods focusing solely on tags. Such approaches may be divided into two main categories:

1) the automatic generation of tags, a process often called "(tag-)recommendation", i.e. extending the descriptions of photos by automatically adding more tags.
2) 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 previous work [17], we extensively surveyed Flickr-related research efforts on the aforementioned categories. In this work we continue our survey in the remaining research areas. As a result, in the following we shall summarize and present four meaningful research group categories.

A. Travel Applications

Since a large percentage of Flickr content is of touristic travel nature, e.g., landmarks, places of interest, events, etc., many research efforts focus on the exploitation of such content for applications that target potential travelers and/or tourists. The latter form typically applications one can download on his mobile device to use during his trip. Most such travel apps focus on things like cutting down journey times, pointing their users in the direction of authentic local events or simply showing them how to ask after e.g. a New York landmark.

1) Visual Reconstruction: The most important step towards the implementation of such applications is considered to have been taken by Agarwal et al. [18] that used a large dataset of Flickr photos that have been taken in Rome, Italy and to other cities and tried to create a 3D reconstructed model in 24 hours using a cloud-based computer architecture. Snavely et al. [19] extracted viewpoints of photos and used them to create paths which were then used fro image-based rendering. This work has been extended in [20], where 3D models were created within a system called *Photo Tourism*, which facilitated users' navigation. Li et al. [21] combined 2D appearance with 3D geometric constraints and using iconic scene graphs they created summaries and 3D reconstructions. Kaminsky et al. [22] created visual reconstructions and with the aid of geo-tags they aligned them with aerial photos. Tuite et al. [23] presented *PhotoCity*, an online game that aimed to train its players so that they became "experts" at taking photos at targeted locations and in great density, with a goal to create 3D building models. They evaluated their approach by reconstructing large portions of two university campuses. The main difference compared to other works is that this way they intended to include certain areas that otherwise did not have much photographic coverage on sites such as Flickr. The presented experimental results and the comparisons confirmed the aforementioned goal.

2) Recommendation Systems: The second category of tourist applications focuses on the recommendation of places of interest. Typical applications in this area focus on automatically discovering main attractions, letting users decide which to visit. Kisilevich et al. [24] used geotagged photos and data clustering techniques in order to determine urban areas of interest, analyzed spatial and temporal distributions so as to identify events and rank places of interest based on their popularity. Van Canneyt et al. [25] proposed a recommendation system for trends in tourist attractions in cities. They used a data set of approx. 665K photos from Flickr. They dealt with recommendations both as a ranking and as an assignment problem and adopted a probabilistic approach. In the first case they ranked places of interest according to their popularity, by also taking into account temporal information about the user, while in the second, the user selected a few places of interest and time-slots he/she was available and their system proposed the best coverage of these. Chen and Roy [26] detected events by exploiting tags, date information and geo-tags and used wavelets in order to handle noisy data efficiently. Their approach performed well in periodic events, while being less accurate in aperiodic. Hollenstein and Purves [27] used a set of 8M Flickr images and a vernacular geography approach, in order to study how accurate are the user generated tags, how can urban areas be described using tags and how can tag images allow for the understanding of the location and extent of vernacular regions. Cao et al. [28] proposed a tourism recommendation system. They used mean-shift clustering and built a set of representative images and tags for each cluster which they use to match users' input. Sang et al. [29] proposed

[6] http://www.flickr.com/api

a framework for personalized image search. Therein, results are ranked based on the tags that the user has already used. Nitta et al. [30] proposed an approach to detect events using tags, geotags and temporal metadata of Flickr photos, by first defining event classes, i.e. semantically related groups of events. Finally, Yao et al. [31] presented a graph based framework which aimed in a unified manner to be used for friend recommendation, image tagging and personalized image search. They used Flickr as a testbed, while claiming that their approach can be easily adapted to other social networks.

3) Trip/Route Suggestion Systems: The third category of tourist applications extends the former, in the sense that it not only recommends main attractions, but also tries to organize the users' schedule and help them visit as many as they wish in a time efficient way. Jain et al. [32] extracted photos around the location of a trip and created a graph. They found tours that start from this location and visit popular places using certain distance constraints. Their system, namely *Antourage* tried to cover all popular landmarks, however they did not consider factors such as time spent at locations and reachability. Popescu and Grefenstette [33] aimed to estimate expected visiting times for tourist attractions, by exploiting Flickr metadata. This work was extended by Popescu et al. [34] where the authors mined spatial and temporal tourist information from Flickr and tried to discover trips in popular cities. Hao et al. [35] presented *Travelscope*, a system that created virtual tours by mining Flickr data. Sun et al. [36] clustered images spatially, identified landmarks within them and ranked them based on their popularity. They aimed to recommend to the user minimum distances with maximum tourism popularity. Majid et al. [37] presented an approach for personalization and recommendation of tourist locations. They obtained users' preferences from their travel history in a city and used this information to recommend locations in an another, unknown city.

B. Extracting Knowledge from Flickr Metadata

In an effort to utilize collective intelligence in the process, many research efforts have leaned towards automatic knowledge representation and organization, using intelligence that has been gathered from Flick textual metadata. As a result, researchers often make use of groups, that may be regarded as a simplistic form of a semantic hierarchy. Schmitz [38] applied a subsumption-based model and used a vocabulary that had been created from Flickr tags and from 9M images, in order to create an ontology. Firan et al. [39] used an ontology, Naive Bayes and SVM classifiers, in order to detect events. Using those events, they ended up classifying Flickr photos. Lu and Li [40] constructed semantic topic-hierarchies and then mapped Flickr groups onto them, so as to construct group-hierarchies. Their initial experiments indicated that these hierarchies may facilitate the browsing experience of users. Negoescu et al. [41] grouped Flickr groups using a cluster approach, in order to create "hypergroups". This way, they allowed smaller groups to be easily discovered by potential users. Prangprasopchok and Lerman [42] constructed quite detailed folksonomies by aggregating relations from different users, following a generic approach which can be applied to several other systems. Lee et al. [43] first extracted points of interest and then by using a clustering approach and by applying association rules mining, they tried to detect associative point of interest patterns. Finally, Xie et al. [44] proposed an algorithm of an augmented folksonomy graph (AFG), which was used in order to incorporate multi-faceted relations in social media. They also used a novel density-based clustering method so as to discover latent user community from AFG by combining contents and tags of multimedia resources.

C. Human Activity Tracking

Considering Flickr as a social network, many researchers studied the relationships among users, user activities and behavior within it. Stvilia and Jorgensen [45] investigated Flickr member activities and how these may assist on the automatic metadata creation. They worked on a set of historical photographs and their respective comments and discussions. Negoescu et al. [46] tried to model users and groups with a common tag-based representation, using a probabilistic topic-based analysis. Cha et al. [47] showed empirically that photos spread due to social links and this propagation process is limited within close connections of users and also slower than their initial expectations. Valafar et al. [48] studied user interactions by analyzing their temporal properties. Marlow et al. [49] presented a taxonomy of tagging systems, with Flickr among them. Their goal was to facilitate analysis and research of such systems. Cox et al. [50] interviewed Flickr users so as to understand the use of the website in conjunction with the users' practices in photography, i.e., they considered Flickr from the scope of hobbyist photographers. They concluded that Flickr should be analyzed as a social network of amateur photographers, rather than a simple photo storage facility. Mislove et al. [51] used empirical data and investigated the link formation process. They showed that links tend to be created by users that have many links, and also users tend to link to those users that are close to them in the social network structure. Finally, Lerman and Jones [52] investigated "social browsing", i.e., the strong correlations users have with their contacts in Flickr. They showed that the primary way that users follow when they aim to discover photos is not by tag-based search or subscription to specialized groups, but browsing through photo streams of their contacts.

D. Other Application Domains

Last but not least, the variety of Flickr content allows for several other research efforts on application domains that do not fall under previously discussed categories. For instance, Jin et al. [53] applied regression- and diffusion-based prediction models on certain Flickr textual and visual features and used them for social studies such as politics, economics and marketing. They experimented on the prediction of product sales and on the American presidential election of 2008. In the first case, they showed that Flickr may monitor the worldwide adoption of products, while for the latter it provided hints that may assist for the prediction of the election results. Similarly, Singh et al. [54] combined Tweets and Flickr posts, in order to study spatio-temporal events. Lei et al [55] adopted a multimodal methodology, based on both acoustic and textual features and aimed to identify cities using machine learning approaches. They observed that in some cases acoustic features are enough for correct classification. Clements et al. [56] used only geo-tags and proposed a weighting scheme, in order to identify similar places at world and city level. Wu et al. [57]

proposed Flickr distance, as a means of measuring relations between semantic concepts, construct a concept network and applied it to concept clustering and image annotation. A similar approach by Zhang et al. [58] aimed to discover and visualize tag relationships from spatial and temporal similarities. Cox et al. [59] proposed metrics in order to characterize and compare Flickr groups. Jaffe et al. [60] generated summaries by selecting the most representative photos, using geo-tags and a clustering approach. Wang et al. [61] proposed a generative probabilistic model and used it for group recommendation. A similar work by Chen et al. [62], focused on the creation of visual summaries to be used as tourist maps, by capturing the most important points of interest. Leung and Newsam [63] used datasets from Flickr and Geograph[7], in an effort to derive maps of "what-is-where" on the surface of earth. Yanai et al. [64], [65] tried to detect cultural differences regarding certain world widespread concepts. Hao et al. [66] presented a methodology for the automatic creation of travelogues, i.e., a kind of a travel-related experience logging. They retrieved geo-tagged photos from Flickr and embed them in these travelogues. Finally, Baber et. al. [67] asked tourists to capture photos of a monument, with the goal to be able to support subsequent question-asking. The results of their study indicated that "much tourist photography represents a special form of image capture" in which tourists tend to gravitate towards the best vantage points to take their own versions of photos seen in brochures.

IV. Conclusions

In this paper we made an effort to survey research efforts that use Flick as the source of their data and focus mainly on the broader areas of semantic and social, user-generated, multimedia content adaptation and personalization. More specifically, we emphasized on methodologies dealing with visual reconstruction approaches, recommendation systems and route suggestion systems, as well as on techniques that aim to extract knowledge from metadata and on approaches that track human activities. Our intention was to complement our previous related work [17], which focused on information retrieval aspects and at the same time summarize in an organized way all trends in the aforementioned areas, so as for fellow researchers to be able to identify an efficient point of reference. Among our future work is the extension of this survey to other popular social networks, like Twitter, Facebook and Youtube[8], and different content types, such as text snippets and video sequences, as well.

References

[1] N. A. Van House, "Flickr and public image-sharing: distant closeness and photo exhibition," in *Proc. of ACM CHI*, 2007.

[2] X. Zeng and L. Wei, "Social ties and user content generation: Evidence from flickr," *Information Systems Research*, vol. 24, no. 1, pp. 71–87, 2013.

[3] J. Van Dijck, "Flickr and the culture of connectivity: Sharing views, experiences, memories," *Memory Studies*, vol. 4, no. 4, pp. 401–415, 2011.

[4] A. J. Flanagin and M. J. Metzger, "The credibility of volunteered geographic information," *GeoJournal*, vol. 72, no. 3-4, pp. 137–148, 2008.

[5] S. E. Spielman, "Spatial collective intelligence? credibility, accuracy, and volunteered geographic information," *Cartography and Geographic Information Science*, no. ahead-of-print, pp. 1–10, 2014.

[6] J. Luo, D. Joshi, J. Yu, and A. Gallagher, "Geotagging in multimedia and computer visiona survey," *Multimedia Tools and Applications*, vol. 51, no. 1, pp. 187–211, 2011.

[7] J. Wu, H. Sun, and Y. Tan, "Social media research: A review," *Journal of Systems Science and Systems Engineering*, vol. 22, no. 3, pp. 257–282, 2013.

[8] C. Prieur, D. Cardon, J.-S. Beuscart, N. Pissard, and P. Pons, "The stength of weak cooperation: A case study on flickr," *arXiv preprint arXiv:0802.2317*, 2008.

[9] J. Cheng and D. Cosley, "How annotation styles influence content and preferences," in *Proc. of ACM HT*, 2013.

[10] J. Shen, M. Wang, S. Yan, and X.-S. Hua, "Multimedia tagging: past, present and future," in *Proc. of ACM MM*, 2011.

[11] E. Angus and M. Thelwall, "Motivations for image publishing and tagging on flickr," in *Proc. of 14th Elpub*, 2010.

[12] S. J. Cunningham and M. Mahoui, "Interacting with and through a digital library collection: commenting behavior in flickr's the commons," in *Proc. of the ACM/IEEE-CS JCDL*, 2013.

[13] O. Nov, M. Naaman, and C. Ye, "What drives content tagging: the case of photos on flickr," in *Proc. of ACM CHI*, 2008.

[14] M. Ames and M. Naaman, "Why we tag: motivations for annotation in mobile and online media," in *Proceedings of the ACM CHI*, 2007.

[15] M. Winget, "User-defined classification on the online photo sharing site flickr or, how i learned to stop worrying and love the million typing monkeys," *Advances in Classification Research Online*, vol. 17, no. 1, pp. 1–16, 2006.

[16] M. Wang, B. Ni, X.-S. Hua, and T.-S. Chua, "Assistive tagging: A survey of multimedia tagging with human-computer joint exploration," *ACM Computing Surveys (CSUR)*, vol. 44, no. 4, p. 25, 2012.

[17] E. Spyrou and P. Mylonas, "A survey of geo-tagged multimedia content analysis within flickr," in *Proc. of MHDW*, 2014.

[18] S. Agarwal, Y. Furukawa, N. Snavely, I. Simon, B. Curless, S. M. Seitz, and R. Szeliski, "Building rome in a day," *Communications of the ACM*, vol. 54, no. 10, pp. 105–112, 2011.

[19] N. Snavely, R. Garg, S. M. Seitz, and R. Szeliski, "Finding paths through the world's photos," in *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3. ACM, 2008, p. 15.

[20] N. Snavely, S. M. Seitz, and R. Szeliski, "Photo tourism: exploring photo collections in 3d," *ACM transactions on graphics (TOG)*, vol. 25, no. 3, pp. 835–846, 2006.

[21] X. Li, C. Wu, C. Zach, S. Lazebnik, and J.-M. Frahm, "Modeling and recognition of landmark image collections using iconic scene graphs," in *Proc. of ECCV*, 2008.

[22] R. S. Kaminsky, N. Snavely, S. M. Seitz, and R. Szeliski, "Alignment of 3d point clouds to overhead images," in *Proc. of IEEE CVPR*, 2009.

[23] K. Tuite, N. Snavely, D.-y. Hsiao, N. Tabing, and Z. Popovic, "Photocity: training experts at large-scale image acquisition through a competitive game," in *Proc. of ACM SIGCHI Conf. on Human Factors in Computing Systems*, 2011.

[24] S. Kisilevich, D. Keim, N. Andrienko, and G. Andrienko, "Towards acquisition of semantics of places and events by multi-perspective analysis of geotagged photo collections," *Geospatial Visualisation*, pp. 211–233, 2013.

[25] S. Van Canneyt, S. Schockaert, O. Van Laere, and B. Dhoedt, "Time-dependent recommendation of tourist attractions using flickr," 2011.

[26] L. Chen and A. Roy, "Event detection from flickr data through wavelet-based spatial analysis," in *Proc. of ACM CIKM*, 2009.

[27] L. Hollenstein and R. Purves, "Exploring place through user-generated content: Using flickr tags to describe city cores," *Journal of Spatial Information Science*, no. 1, pp. 21–48, 2013.

[28] L. Cao, J. Luo, A. Gallagher, X. Jin, J. Han, and T. S. Huang, "Aworldwide tourism recommendation system based on geotaggedweb photos," in *Proc. of IEEE ICASSP*, 2010.

[29] J. Sang, C. Xu, and D. Lu, "Learn to personalized image search from

[7] http://www.geograph.com

[8] http://www.youtube.com

the photo sharing websites," *Multimedia, IEEE Transactions on*, vol. 14, no. 4, pp. 963–974, 2012.

[30] N. Nitta, Y. Kumihashi, T. Kato, and N. Babaguchi, "Real-world event detection using flickr images," in *MultiMedia Modeling*. Springer, 2014, pp. 307–314.

[31] T. Yao, Y. Liu, C.-W. Ngo, and T. Mei, "Unified entity search in social media community," in *Proc. of ACM WWW*, 2013.

[32] S. Jain, S. Seufert, and S. Bedathur, "Antourage: Mining distance-constrained trips from flickr."

[33] A. Popescu and G. Grefenstette, "Deducing trip related information from flickr," in *Proc. of ACM WWW*, 2009.

[34] A. Popescu, G. Grefenstette, and P.-A. Moëllic, "Mining tourist information from user-supplied collections," in *Proc. of ACM CIKM*, 2009.

[35] Q. Hao, R. Cai, J.-M. Yang, R. Xiao, L. Liu, S. Wang, and L. Zhang, "Travelscope: standing on the shoulders of dedicated travelers," in *Proc. of ACM MM*, 2009.

[36] Y. Sun, H. Fan, M. Bakillah, and A. Zipf, "Road-based travel recommendation using geo-tagged images," *Computers, Environment and Urban Systems*, 2013.

[37] A. Majid, L. Chen, G. Chen, H. T. Mirza, I. Hussain, and J. Woodward, "A context-aware personalized travel recommendation system based on geotagged social media data mining," *International Journal of Geographical Information Science*, vol. 27, no. 4, pp. 662–684, 2013.

[38] P. Schmitz, "Inducing ontology from flickr tags," in *Proc. of WWW*, 2006.

[39] C. S. Firan, M. Georgescu, W. Nejdl, and R. Paiu, "Bringing order to your photos: event-driven classification of flickr images based on social knowledge," in *Proc. of ACM CIKM*, 2010.

[40] D. Lu and Q. Li, "Exploiting semantic hierarchies for flickr group," in *Active Media Technology*. Springer, 2010, pp. 74–85.

[41] R.-A. Negoescu, B. Adams, D. Phung, S. Venkatesh, and D. Gatica-Perez, "Flickr hypergroups," 2009.

[42] A. Plangprasopchok and K. Lerman, "Constructing folksonomies from user-specified relations on flickr."

[43] I. Lee, G. Cai, and K. Lee, "Exploration of geo-tagged photos through data mining approaches," *Expert Systems with Applications*, vol. 41, no. 2, pp. 397–405, 2014.

[44] H. Xie, Q. Li, X. Mao, X. Li, Y. Cai, and Q. Zheng, "Mining latent user community for tag-based and content-based search in social media," *The Computer Journal*, 2014.

[45] B. Stvilia and C. Jörgensen, "Member activities and quality of tags in a collection of historical photographs in flickr."

[46] R.-A. Negoescu and D. Gatica-Perez, "Modeling flickr communities through probabilistic topic-based analysis," *Multimedia, IEEE Trans. on*, vol. 12, no. 5, pp. 399–416, 2010.

[47] M. Cha, A. Mislove, and K. P. Gummadi, "A measurement-driven analysis of information propagation in the flickr social network," in *Proc. of ACM WWW*, 2009.

[48] M. Valafar, R. Rejaie, and W. Willinger, "Beyond friendship graphs: a study of user interactions in flickr," in *Proc. of ACM WOSN*, 2009.

[49] C. Marlow, M. Naaman, D. Boyd, and M. Davis, "Ht06, tagging paper, taxonomy, flickr, academic article, to read," in *Proc. of ACM HT*, 2006.

[50] A. M. Cox, P. D. Clough, and J. Marlow, "Flickr: a first look at user behaviour in the context of photography as serious leisure," *Information Research*, vol. 13, no. 1, p. 5, 2008.

[51] A. Mislove, H. S. Koppula, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Growth of the flickr social network," in *Proc. of ACM WOSM*, 2008.

[52] K. Lerman and L. Jones, "Social browsing on flickr," *arXiv preprint cs/0612047*, 2006.

[53] X. Jin, A. Gallagher, L. Cao, J. Luo, and J. Han, "The wisdom of social multimedia: using flickr for prediction and forecast," in *Proc. of ACM MM*, 2010.

[54] V. K. Singh, M. Gao, and R. Jain, "Social pixels: genesis and evaluation," in *Proc. of ACM MM*, 2010.

[55] H. Lei, J. Choi, and G. Friedland, "Multimodal city-verification on flickr videos using acoustic and textual features," in *Proc. of IEEE ICASSP*, 2012.

[56] M. Clements, P. Serdyukov, A. P. de Vries, and M. J. Reinders, "Using flickr geotags to predict user travel behaviour," in *Proc. of ACM SIGIR*, 2010.

[57] L. Wu, X.-S. Hua, N. Yu, W.-Y. Ma, and S. Li, "Flickr distance," in *Proc. of ACM MM*, 2008.

[58] H. Zhang, M. Korayem, E. You, and D. J. Crandall, "Beyond co-occurrence: discovering and visualizing tag relationships from geo-spatial and temporal similarities," in *Proc. of ACM WSDM*, 2012.

[59] A. Cox, P. Clough, and S. Siersdorfer, "Developing metrics to characterize flickr groups," *Journal of the American Society for Information Science and Technology*, vol. 62, no. 3, pp. 493–506, 2011.

[60] A. Jaffe, M. Naaman, T. Tassa, and M. Davis, "Generating summaries for large collections of geo-referenced photographs," in *Proc. of ACM WWW*, 2006.

[61] J. Wang, Z. Zhao, J. Zhou, H. Wang, B. Cui, and G. Qi, "Recommending flickr groups with social topic model," *Information retrieval*, vol. 15, no. 3-4, pp. 278–295, 2012.

[62] W.-C. Chen, A. Battestini, N. Gelfand, and V. Setlur, "Visual summaries of popular landmarks from community photo collections," in *Proc. of IEEE ASILOMAR*, 2009.

[63] D. Leung and S. Newsam, "Proximate sensing: Inferring what-is-where from georeferenced photo collections," in *Proc. of IEEE CVPR*, 2010.

[64] K. Yanai, K. Yaegashi, and B. Qiu, "Detecting cultural differences using consumer-generated geotagged photos," in *Proc. of ACM LOCWEB*, 2009.

[65] K. Yanai, H. Kawakubo, and B. Qiu, "A visual analysis of the relationship between word concepts and geographical locations," in *Proc. of ACM ICIVP*, 2009.

[66] Q. Hao, R. Cai, C. Wang, R. Xiao, J.-M. Yang, Y. Pang, and L. Zhang, "Equip tourists with knowledge mined from travelogues," in *Proc. of ACM WWW*, 2010.

[67] C. Baber, J. Cross, T. Khaleel, and R. Beale, "Location-based photography as sense-making," in *Proc. of BCS HCI*, 2008.