Tag Suggestions from Social Media Profiles

Petri RANTANEN [a], Pekka SILLBERG [a], Jari SOINI [a,1] and Hannu JAAKKOLA [a]
[a] *Tampere University of Technology, Pori Department*

Abstract. Attaching any kind of clue – event, location, person, tag or keyword – to a photo eases the process of searching. Often the problem is that the user feels that it is difficult to think of good tags or that the tagging process is too tedious or cumbersome. At the same time, users use social media daily, and write about topics they feel are important and that they are actively interested in. This paper presents a method for extracting metadata (tag suggestions) from social media profiles and illustrates the use of the tags for photo tagging by means of a web-based photo application.

Keywords. Metadata, tags, content-based analysis, social media

1. Introduction

Searching for a particular photo from one's photo collection can be a time-consuming task. In a small set of photos one could recall from memory where a photo might be, but as the size of the collection grows larger, the more impractical it becomes to simply trust one's memory. It might also be possible that someone other than the photographer is browsing the photo collection for a certain photo. Attaching any kind of clue – event, location, person, tag, or keyword – to the photo eases the process of searching. Often the problem is that users feel that it is difficult to think of good tags. At the same time, the same users use social media daily, and write about topics they feel are important and that they are actively interested in. In many cases these activities – posts on social media and taking pictures – contain similar elements and describe the same events. Thus, it should be possible to extract usable tags from social media profiles to be used with pictures.

Automatic tag generation for images and tagging in general has been extensively researched in recent years. Nevertheless, no comprehensive solution has been found, and often one approach might work for one use case, but not for another. Also, many of the pre-existing systems use the image content or related images to generate new tags essentially requiring a pre-defined set of images used for training of the analysis engines or for reference. [1, 2, 3, 4]

Additional concern is the selection of the tags – or classes used for image analysis. The web client shown in this paper has been previously used in a user trial, which studied the tagging behavior of the users as well as the functionality and ideas of the user interface [5]. One of the results of the study was the feedback on the usability of the tags. In the user trial we used a pre-defined set of more than 300 tags, which

[1] Corresponding Author.

represented the classes used to teach visual features to the analysis engines. Even though the tags (e.g. airplane, person, beach…) worked acceptable in describing the features or objects in the photos and performed well for photo similarity search, the users reported that they would not normally use the tags. Unfortunately, it is fairly difficult to figure out which tags the users would actually prefer to use.

In our scenario, we attempt to generate tags based on the content the user has provided on other services – in this case, Facebook and Twitter. There has been some research on the subject of utilizing content from external services [6]. In our case, text summarization is used to produce tags from the user's social media profiles – tweets, status messages, groups, and events. The assumption is that the user's activities in the social media relate to pictures submitted to other sites, even though no direct connection exists between the sites in question. Text summarization has been used in the generation of tag for photo content with encouraging results [7, 8], though in our case the source material does not necessarily contain any photos and is simply used refine the suggestion list offered to users when associating tags to photos.

The research presented in this paper is based on the work done in a large national research project called Data to Intelligence (D2I) [9]. D2I is a four-year Finnish Strategic Research and Innovation Agenda (SRIA) program that aims to study and develop methods and tools for the management, processing, and utilization of large amounts of data captured from the environment, Internet, and many other sources. The basis of the task-based analysis system [10], the photo tagging user interface [5, 11], and the summarizer design [12] has been presented in earlier publications related to the project.

The contents of this paper are as follows. Chapters 2 and 3 briefly describe the technical approach and the overall architecture of the designed system. Chapter 4 illustrates how the generated tags can be visualized to the user. And finally, chapters 5 and 6 provide insight on current research issues and possible future research directions, and provide the reasoning behind the chosen approach used for automatic tag generation.

2. Social Media Summarization

Figure 1 illustrates the basic flow of the summarization process. The numbers in the figure describe the order of events or steps. The initial step of the social profile summarization is to request the user's consent to access the profile. In principle, it would also be possible to perform limited summarization on accounts publicly available on the Internet. In our use case we are specifically interested in the users' personal accounts. Commonly, asking for consent is a one-time operation, which is only required when the user connects his/her Facebook or Twitter account to the service.

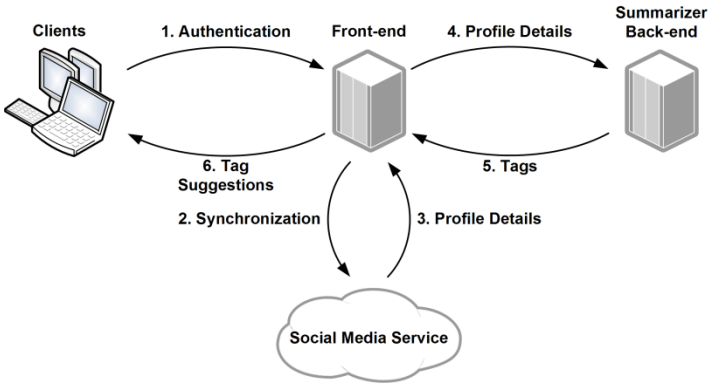

Figure 1. Basic flow of summarization process.

For Facebook, the consent request and the following API requests are authorized using the OAuth2 authentication scheme [13]. The OAuth2 scheme for requesting access tokens [14] is relatively simple to implement, and there are also various libraries available to handle the process for the commonly used social media services. It is also the recommended way to access Facebook APIs. The generated access tokens are valid for 60 days and are refreshed on each use, which means that executing the profile summarization on a regular basis should keep the token valid without the user having to authorize it separately on each occasion.

With Twitter OAuth2 is not supported, and the older OAuth must be used [15]. The authentication flow differs from Facebook's OAuth2 implementation, but the actual API usage is similar after the user consent has been given and the required access tokens have been generated. In Twitter's case, the access tokens will never expire, and thus there is no need for refresh tokens. With OAuth, similarly to OAuth2, it is possible for the access tokens to become invalid for other reasons than the predefined token expiration time (for example, the user might revoke the access), and it might be required to ask for consent from the user at some later date after the initial request.

After the user has provided consent, the front-end will retrieve the social media profiles for the summarization. The synchronization process can be manually executed by the user, or scheduled to occur automatically at pre-defined time intervals (e.g. once per week). The frequency of the summarization can have a large impact on the usability of the system. Summarizing often will provide more up-to-date tags, but on the other hand, will require more processing power on the analysis back-ends. A more rapid summarization cycle can also produce other practical issues caused by the API query limitations enforced by the social media services.

Retrieving an entire Facebook account or a Twitter account is a relatively fast operation, taking from a few seconds to several minutes, depending on the size of the account, and it is easy to control the duration of the retrieval by paging the data in smaller chunks, or by simply retrieving only the newest content.

Depending on the scopes provided for the initial consent request, it is possible to either retrieve the entire social media profile or a limited set of data. In general, consent should be requested only for the required data, and one should not ask users to give access to all their content simply because it is possible. In our case, for Facebook we retrieve user's events, groups, likes, and status messages in addition to the basic account details (user's name, location, biography, etc...). It is also possible to retrieve

the content that the user's friends have allowed the user to access, but in our case the use of this material is very limited, because we are more interested in the user him/herself than his/her friends. For Twitter, there is not that much variation on what can be retrieved, and the content is mostly limited to the basic account details and the user's tweets.

In practice, most of the generated tags originate from the user's tweets and status messages. This is caused by configuration settings on the summarizer, which give more weight to the aforementioned content. The configuration can be modified and it can even be modified on a per user basis to provide different weighting values for different users, though setting the weights manually is perhaps something an average user would rather not do. The reasoning behind the default weighting bias is that some content (e.g. groups or biography) does not always provide very up-to-date information on the user's activities or interests. Additionally, some content (e.g. groups or likes) can be thought to be tags themselves and as such, do not need to be summarized, and are for research purposes given lower priority in order to emphasize the extracted tags.

In any case, the retrieved content is converted by the front-end service to a uniform task format used to deliver the data to the back-ends (in this case, the summarizer). The format is similar, though the content is not identical for both use cases (Facebook and Twitter). Tweets are almost identical to Facebook status messages and can be represented with a common structure in the task format, but some data does not exist for both services (e.g. Twitter does not have the concept of groups or events). Basically, in the produced task workload, the two cases are differentiated only by the actual workload content and the front-end specified service identifier. Without the unified task format all back-ends would need to support all formats provided by every service connected to the front-end service, which would require a lot of extra work for the people responsible for the back-end implementation. The downside is that in some cases it might be very difficult to represent – or restructure – the retrieved data in the uniform format.

On the basis of the submitted task the summarizer will produce a list of tags. The task scheduling is described in more depth in [10], and the summarization process is explained in [12]. The generated tags are stored on the back-end and other content retrieved from the social media profiles – and the profiles themselves – can either be stored or discarded. In our use case we are simply interested in the generated tags, and as such, the actual profile content can be discarded. In theory, it could be possible to synchronize the generated tags back to the originating service. For example, if the Facebook status messages contained images, these images could be tagged with the generated keywords using Facebook APIs. In our use case these image-tag relations are kept on the front-end and are not synchronized outside the service.

3. Photo Tagging with Suggestions

The tags extracted from the user’s social media account can be used for photo tagging. This is enabled by implementing a text completion – or tag suggestion – into the edit mode, which was presented in [11]. The input fields for adding of new tags are now constantly updated as the user types characters into them. The functionality is similar to that found in, for example, input fields in web browsers. The summarized example account containing content about table ice hockey games is seen in Figure 2 and the extracted tag suggestions are illustrated in Figure 3. The user has typed in *hock*, and the

system returns ice hockey related suggestions (*hockey*, *hockey player*, etc…) for a new tag.

Figure 2. A Facebook account used in summarization.

Tags presented in the drop down list are in descending order based on the matching score. The lists used to produce the suggestions are prioritized primarily on user ranking of the tags, secondly, by summarizer "rating" of the words. The suggestions can also contain the user's custom tags (if the user has previously added new tags) and tags previously extracted from other sources. The "origin" (how and where the tags were extracted from) of the tags is stored in the system and can be retrieved, though in our use case, it is assumed that the user does not care about the origin, but is simply interested in the tags themselves.

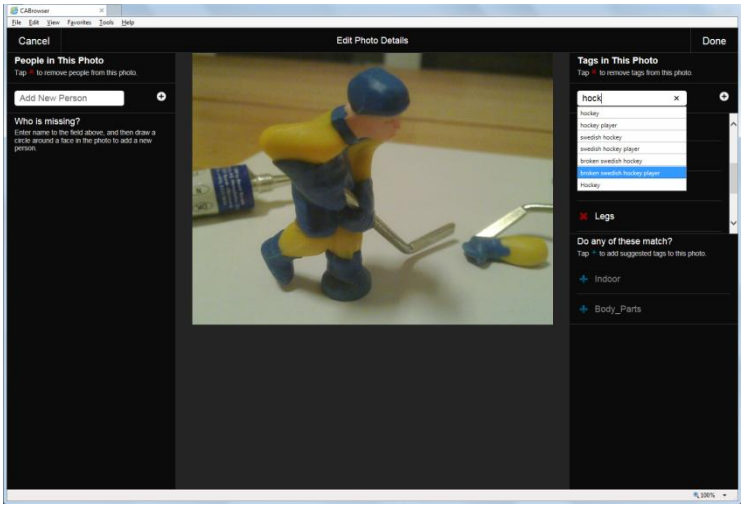

Figure 3. Screenshot of the edit mode.

The edit UI uses JavaScript to access the REST API that provides suggestions, which means that it is also possible to obtain suggestions in other applications. The edit mode also shows tags (Figure 3, below *Do any of these match?* text) which are detected

by analyzing the photo in question. The tags are not by definition considered "tag suggestions" in the context of this paper, though in practice, the list of tags given when starts typing a new tag can contain tags previously extracted by using visual analysis of the photo. Moreover, the user is free to choose from these "recommended" tags or start typing a new tag, in which case the suggestion API is used to provide missing tags. The list of tags inserted (or accepted from the recommendations or suggestions) by the user is listed below the input field. In Figure 3 only the tag *Legs* is visible in the tag list.

4. Tag Evaluation

The tags discovered by the social media profile summarization can also be retrieved separately from the other tags created by the user. To support the evaluation of the summarization algorithm, we created a web application that retrieves the twenty best-ranked tags, and lays them out as different sized and colored circles. An example of tags from a summarized profile can be seen in Figure 4. A higher score, or relevancy, translates into a larger circle size.

Figure 4. Screenshot of the tag evaluation application.

In the scope of evaluating the summarization accuracy [12], the test user may select one of the three states for a tag: neutral (shades of blue), accepted (green background with underlined text), or rejected (red background with strikethrough text). The different background colors of the neutral tag circles simply illustrate the number of words in the tag (i.e. unigrams, n-grams), and do not bear any functional purpose. Although the tag evaluation application is not aimed at end users, the tag state can be taken into account when calculating the matching score. The accepted tag can be given

a higher score for it to appear higher in the tag suggestion list shown in Figure 3. Similarly, a rejected tag gets a lower score, and may not appear in the list at all.

5. Discussion

The assumption is that users post similar or related content on all the services they use. Thus, extracting content from one service should provide supporting content for other services. In our use case, a user might tweet about his holiday trip even though the tweets do not contain any pictures, and later upload the photos of the trip to another service, which is perhaps better suited for storing or sharing pictures.

Instinctively this seems like a valid hypothesis, although it is difficult to find research to support or refute it. It is easy to find publications concerning individual social media websites, but much more difficult to find research concerning users' usage patterns over multiple sites. Especially, comparisons of sites of different "types" (for example, blogging site content compared to image sharing site content) are hard to find. One reason for this might be that the statistics provided by the services themselves are often anonymous in nature, and limited to what the service in question is willing to publish.

However, this is not a deciding factor in our scenario. Using the summarization technique, it is also possible, for example, to summarize a Facebook account to provide tag suggestions for pictures stored on Facebook. It is also possible to target the summarization process for status messages (or tweets) with pictures, increasing the correlation of tags with the pictures, although this does not necessarily provide viable tags for photo collections not referenced by status messages.

Based on earlier studies related to the design and implementation of the text summarizer utilized in our case [12], about half of the tags extracted from social media profiles were felt by the users to match the content of their profiles. It should be noted that the users were not especially asked to associate the tags with previously uploaded pictures, but simply to give an opinion on whether they might use the tags in the future. Thus, a larger user study is required before any conclusions on the usability of the extracted tags can be made.

To original problem of using tags, which the users did not feel they would use in a practical scenario, for the basis of training material for visual analysis engines, the tags extracted from the social media profiles unfortunately do not provide much help. In practice most of the extracted tags are too abstract (or simply too noisy) for the creation of valid classifiers for visual analysis – that is, it is difficult to find large enough set of example photos, which correlate unambiguously to the tags to use as training material. Even though, in our experience, the tags did not work adequately as new classifiers, the extracted metadata can potentially have other uses. In the context of this paper only simple tag suggestions were extracted, but using text summarization – and further analysis of the social media profiles – it is possible to gather other useful information about the user, such as the user's hobbies, groups, participated events and other details helpful for detecting the user's personal preferences across multiple online services.

Additionally, one direction for future research would be to take better advantage of the time aspect of the generated tags, which is not extensively utilized. The summarizer implementation does consider word frequency, and as such, commonly used tags will be ranked higher. In practice, repeatedly used words in a short time period in the past and no use of the words in the future can still bias the results. The

front-end also stores the timestamps for tag updates, for example, when the user updates a tag or when the summarizer updates the tag with a more accurate confidence value (weight). It is possible to perform a naive ranking of the suggestion results based on the timestamp or filter out tags, which have not been updated for a while, although the current user interface provides no easy way to access this functionality.

6. Summary

This paper described a method and a general principle for extracting tags from social media profiles by means of social media summarization. The use of extracted metadata as tag suggestions in a web-based photo application was demonstrated and a web user interface for evaluating the extracted tags was presented. The solution presented in this paper is a concrete example of how content-based analysis of social media profiles can be utilized to help users to manage their multimedia content.

References

[1] A.K. Milicevic, A. Nanopoulos and M. Ivanovic, Social tagging in recommender systems: a survey of the state-of-the-art and possible extensions, *Artificial Intelligence Review* **33**, Issue 3 (2010), 187–209.

[2] D. Zhang, M.M. Islam and G. Lu, A review on automatic image annotation techniques, *Pattern Recognition* **45**, Issue 1 (2012), 346–362.

[3] M. Wang, B. Ni, X. Hua and T. Chua, Assistive Tagging: A Survey of Multimedia Tagging with Human-Computer Joint Exploration, *ACM Computing Surveys (CSUR)* **44**, Issue 4 (2012).

[4] T. Pavlidis, Why meaningful automatic tagging of images is very hard, Proceedings of the 2009 IEEE International Conference on Multimedia and Expo (2009), 1432–1435.

[5] M. Aqdas and M. Nieminen, Understanding the Usage and Requirements of the Photo Tagging System, *Human IT* (2014), 117–161.

[6] P.J. McParlane and J.M. Jose, Exploiting Twitter and Wikipedia for the Annotation of Event Images, Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval (2014), 1175–1178.

[7] Digile, Data to Intelligence (D2I), http://www.datatointelligence.fi

[8] G. Wang, D. Hoiem and D. Forsyth, Building text features for object image classification, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2009), 1367–1374.

[9] Z-J. Zha, M. Wang, J. Shen and T-S. Chua, Text Mining in Multimedia, *Mining Text Data*, Springer US (2012), 361–384.

[10] P. Rantanen, P. Sillberg and J. Soini, Content analysis system for Images, Proceedings of the 16th International Multiconference Information Society (2013), 241–244.

[11] P. Sillberg, P. Rantanen and J. Soini A content based tool for searching, connecting and combining digital information – case: smart photo service, Proceedings of the 16th International Multiconference Information Society (2013), 249–252.

[12] T. Forss, S. Liu and K. Björk, Extracting People's Hobby and Interest Information from Social Media Content, *Terminology and Knowledge Engineering* (2014).

[13] Facebook, Facebook Login for the Web with the JavaScript SDK, https://developers.facebook.com/docs/facebook-login/web

[14] Internet Engineering Task Force (IETF), The OAuth 2.0 Authorization Framework, https://tools.ietf.org/html/rfc6749 https://tools.ietf.org/html/rfc6750

[15] Twitter, Inc., OAuth Send secure authorized requests to the Twitter API, https://dev.twitter.com/oauth