A Meteoroid on Steroids: Ranking Media Items Stemming from Multiple Social Networks

Thomas Steiner
Google Germany GmbH
ABC-Str. 19
20354 Hamburg, Germany
tomac@google.com

ABSTRACT

We have developed an application called Social Media Illustrator that allows for finding media items on multiple social networks, clustering them by visual similarity, ranking them by different criteria, and finally arranging them in media galleries that were evaluated to be perceived as aesthetically pleasing. In this paper, we focus on the ranking aspect and show how, for a given set of media items, the most adequate ranking criterion combination can be found by interactively applying different criteria and seeing their effect on-the-fly. This leads us to an empirically optimized media item ranking formula, which takes social network interactions into account. While the ranking formula is not universally applicable, it can serve as a good starting point for an individually adapted formula, all within the context of Social Media Illustrator. A demo of the application is available publicly online at the URL http://social-media-illustrator.herokuapp.com/.

Categories and Subject Descriptors

H.3.3 [Information Search and Retrieval]: Clustering

General Terms

Algorithms

Keywords

Ranking, Event Summarization, Social Networks

1. INTRODUCTION

When people witness events like concerts, sports matches, or meteoroid impacts, they more and more share media items like photos and videos that depict these events publicly on social networks. In the past, we have worked on methods [3, 5, 9] for the automatic extraction, deduplication, and clustering of media items stemming from multiple social networks. Up to now, we have ordered the retrieved media items chronologically, by social network, or by cluster size, and thereby completely neglected social network interactions as ranking signals. Though truly added value lies in exploiting these social network interactions in order to obtain a more representative ranking of the potentially overwhelmingly many media items retrieved for a given event.

WWW 2013 Companion, May 13–17, 2013, Rio de Janeiro, Brazil.
ACM 978–1–4503–2038–2/13/05.

2. RELATED WORK

In [6], San Pedro and Siersdorfer propse a methodology for automatically ranking and classifying photos from the photo sharing platform Flickr according to their attractiveness for Flickr members. They work with extracted user feedback and annotations available on Flickr to train machine learning models based on image features like sharpness and colorfulness. While their method is tailored to Flickr, our approach is based on a social network interaction abstraction layer on top of the social networks Facebook, Twitter, Google+, Instagram, YouTube, Flickr, MobyPicture, Twitpic, and Lockerz. Jaffe et al. describe [2] a ranking and summary algorithm for geo-tagged photo sets based on spatial patterns as well as textual-topical patterns and photographer identity cues. Their algorithm can be expanded to support social, temporal, and other factors. The shown maps-based application necessarily requires geo-tagged media items, which is rarely the case with media items retrieved from social networks due to privacy concerns. In [1], Davidson et al. describe the different criteria video quality, user specificity, and diversification that determine the video ranking in the YouTube recommendation system. These criteria include view count, the ratings of the video, commenting, favoriting, and sharing activity around the video. Finally, Wiyartanti et al. introduce in [11] a ranking algorithm for user-generated videos based on social activities.

3. SOCIAL NETWORK INTERACTIONS

3.1 Abstraction Layer

Social networks have different paradigms of social interactions. In [5], we have introduced an abstraction layer on top of the native data formats of all considered social networks in order to gain an agnostic view on them. Regardless of the native data representation format of the social network of origin, the abstraction layer unifies and streamlines the available data for each media item to a greatest common divisor of all social networks. These interaction paradigms must be exposed by the social networks via specific API calls in order to be considered. In Table 1, we detail how we abstract the social interactions in question on each social network. We differentiate between unknown values that are returned as unknown, i.e., where the information is not exposed, and 0 values, where the value is known to be zero.

3.2 Merging Social Interactions

In the context of our previous research, we have developed a tile-wise histogram-based media item deduplication algorithm with additional high-level semantic matching cri-

Likes	Shares	Comments	Views
Facebook Like	Facebook Share	Facebook Comments	YouTube Views
Google+ +1	Google+ Share	Google+ Comments	Flickr Views
Instagram Like	Twitter native ReTweet	Instagram Comments	Twitpic Views
Flickr Favorite		Twitter manual RT, @Replies	MobyPicture Views
YouTube Like		Twitpic Comments	
YouTube Favorite		MobyPicture Comments	
Twitter Favorite		Flickr Comments	

Table 1: Abstract social network interaction paradigms and their underlying native counterparts

teria that is tailored to photos and videos stemming from multiple social networks. If a set of media items is visually similar enough to be clustered under the criteria detailed in [5], we treat the whole of the cluster as if it were just one media item. These criteria are pair-wise tile histogram similarity that does not exceed a given threshold and the same number of detected faces per media item. In consequence, we specify a merging strategy for the associated social interactions of the individual media items in the particular cluster. We treat unknown values as 0. The alternative to this solution is to exclude unknown values from the merging step. However, as in practice a considerable amount of social interaction values are unknwon, we are forced to proceed with the abovementioned simplification. The algorithm accumulates individual social interactions and assigns the accumulated social interactions to the cluster.

4. RANKING MEDIA ITEM CLUSTERS

In this section, we describe a ranking formula to rank a set of media clusters that match a given query. In the ranking formula, we consider several well-defined ranking criteria that were detailed in [10], namely these are visual, audial, textual, temporal, social, and aesthetic. For a given set of media item clusters, a ranking is calculated as follows.

$$\alpha \times \mathit{likes} + \beta \times \mathit{shares} + \gamma \times \mathit{comments} + \delta \times \mathit{views} + \epsilon \times \mathit{clusterSize} + \zeta \times \mathit{recency} + \eta \times \mathit{quality} \quad (1)$$

The factors *likes*, *shares*, *comments*, and *views* stem from the individual media items as described in Subsection 3.1 and Subsection 3.2. The factor *clusterSize* corresponds to the size of the current cluster. The factor *recency* is calculated as follows. If the youngest media item in the cluster is less than or exactly one day old, the value is 8, for two days it is 4, for three days it is 2, and for each day more, the value is 1. The factor *quality* is a representation of the presence of faces and a media item's photo or video quality. Empirically optimized default values that can be fine-tuned for a concrete media item set were determined as follows: $\alpha = 2$, $\beta = 4$, $\gamma = 8$, $\delta = 1$, $\epsilon = 32$, $\zeta = 2$, and $\eta = 8$.

Once a final ranking for all media items has been found, the top-n media items are compiled to different kinds of media galleries that in two user studies were shown to be perceived as aesthetically pleasing [8]. We differentiate between the Loose Order, Varying Size style, where certain media items can be featured more prominently by making them bigger at the cost of loosely disrespecting the ranking-implied order and the Strict Order, Equal Size style, which strictly respects the ranking-implied order [8].

5. IMPLEMENTATION DETAILS

The application has been implemented in Node.js, a server side JavaScript software system designed for writing scalable Internet applications. Programs are created using event-driven, asynchronous input/output operations to minimize overhead and maximize scalability. The clustering and ranking logic is kept on the client side, while the media item retrieval logic is kept on the server side. As the clustering logic needs read access to the pixel data of media items via the canvas element's getImageData function, all media items need to be proxied locally. Face detection works fully on the client side based on a library made available by Liu [4]. The interface is fully interactive and event-driven. Figure 1 and Figure 2 show screenshots of the deployed application.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented an application called Social Media Illustrator with a special focus on its social interactions abstraction layer and ranking capacities. The application has been successfully evaluated to produce both meaningful and beautiful visual and audial summaries for recent events. The majority of these summaries were made available online.[1] One example of such can be seen in Figure 2, which shows popular social media reactions for the meteoroid impact event on 15 February 2013, when a small asteroid entered the atmosphere of Earth, became visible as a bright fireball and exploded in an air burst over Chelyabinsk.

Future work will focus on adding more visualization formats that will support text-to-speech once the text synthesis part of the Web Speech API [7] has landed in Web browsers. This will allow for true story-telling, where the associated microposts for a media item can be read as it is shown, potentially in an interactive slideshow format. Further, we plan to add more clustering options that will allow for also clustering by extracted named entities [9] besides the currently visual clustering.

Concluding, with our Social Media Illustrator application, we have contributed an effective and efficient tool to deal with social media overload and to identify the few needles in the social network haystack.

7. REFERENCES

[1] J. Davidson, B. Liebald, J. Liu, P. Nandy, T. Van Vleet, U. Gargi, S. Gupta, Y. He, M. Lambert, B. Livingston, and D. Sampath. The YouTube Video Recommendation System. In Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10, pages 293–296, 2010.

[1] http://twitpic.com/photos/tomayac, accessed 02/21/2013

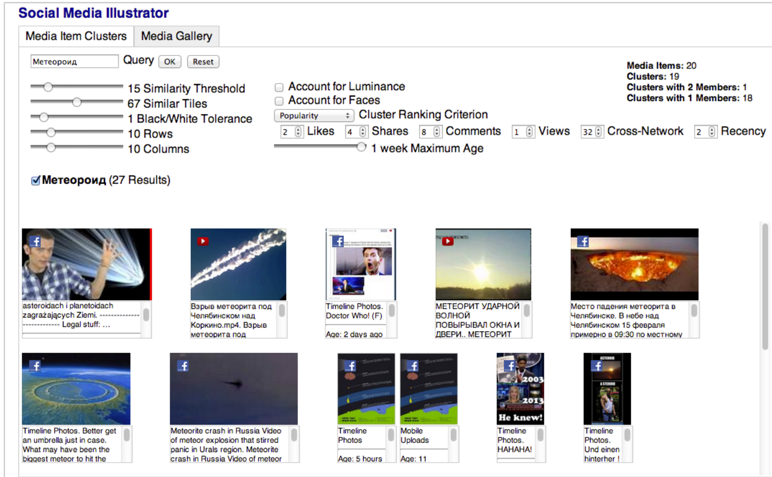

Figure 1: Media Item Clusters tab of the Social Media Illustrator application with individual and clustered (bottom middle) media items from Facebook and YouTube, ranked by popularity for the Russian query Метеороид

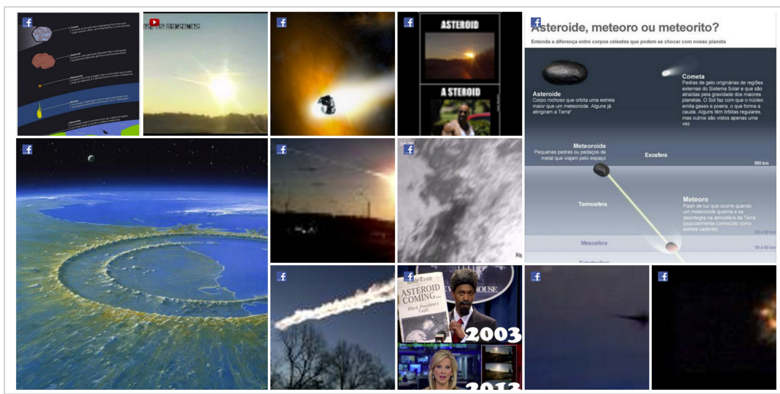

Figure 2: Zoomed view of the Media Gallery tab of the application showing an automatically generated media gallery in loose order, varying size style featuring ranked media items stemming from Facebook and YouTube for the query Метеороид

[2] A. Jaffe, M. Naaman, T. Tassa, and M. Davis. Generating Summaries and Visualization for Large Collections of Geo-referenced Photographs. In Proceedings of the 8th ACM International Workshop on Multimedia Information Retrieval, pages 89–98, 2006.

[3] H. Khrouf, G. Atemezing, G. Rizzo, R. Troncy, and T. Steiner. Aggregating Social Media for Enhancing Conference Experience. In Real-Time Analysis and Mining of Social Streams, Papers from the 2012 ICWSM Workshop, June 2012.

[4] L. Liu. JavaScript Face Detection Explained. http://liuliu.me/eyes/javascript-face-detection-explained/, 2012. Accessed March 26, 2013.

[5] G. Rizzo, T. Steiner, R. Troncy, R. Verborgh, J. L. Redondo García, and R. Van de Walle. What Fresh Media Are You Looking For?: Retrieving Media Items From Multiple Social Networks. In Proceedings of the 2012 International Workshop on Socially-Aware Multimedia, SAM '12, pages 15–20. ACM, 2012.

[6] J. San Pedro and S. Siersdorfer. Ranking and Classifying Attractiveness of Photos in Folksonomies. In Proceedings of the 18th International Conference on World Wide Web, WWW '09, pages 771–780, 2009.

[7] G. Shires and H. Wennborg. Web Speech API Specification. Final Report, W3C, Oct. 2012.

[8] T. Steiner and C. Chedeau. To Crop, Or Not to Crop: Compiling Online Media Galleries. In Proceedings of the 22nd International Conference Companion on World Wide Web, WWW '13, May 2013.

[9] T. Steiner, R. Verborgh, J. Gabarro, and R. Van de Walle. Adding Meaning to Facebook Microposts via a Mash-up API and Tracking its Data Provenance. In Next Generation Web Services Practices (NWeSP), 2011 7th International Conference on, pages 342–345, Oct. 2011.

[10] T. Steiner, R. Verborgh, J. Gabarro, and R. Van de Walle. Defining Aesthetic Principles for Automatic Media Gallery Layout for Visual and Audial Event Summarization based on Social Networks. In Quality of Multimedia Experience (QoMEX), 2012 Fourth International Workshop on, pages 27–28, July 2012.

[11] L. Wiyartanti, Y.-S. Han, and L. Kim. A Ranking Algorithm for User-generated Video Contents Based on Social Activities. In Digital Information Management, 2008. ICDIM 2008. Third International Conference on, pages 260–265, Nov. 2008.