

What Fresh Media Are You Looking For?

Retrieving Media Items from Multiple Social Networks

Giuseppe Rizzo
EURECOM
Sophia Antipolis, France
giuseppe.rizzo@eurecom.fr

Ruben Verborgh
Ghent University – IBBT, ELIS
Ghent, Belgium
ruben.verborgh@ugent.be

Thomas Steiner
Univ. Politècnica de Catalunya
Barcelona, Spain,
tsteiner@lsi.upc.edu

José Luis Redondo Garcia
EURECOM
Sophia Antipolis, France
redondo@eurecom.fr

Raphaël Troncy
EURECOM
Sophia Antipolis, France
raphael.troncy@eurecom.fr

Rik Van de Walle
Ghent University – IBBT, ELIS
Ghent, Belgium
rik.vandewalle@ugent.be

ABSTRACT

Social networks play an increasingly important role for sharing media items related to daily life moments or for the live coverage of events. One of the problems is that media are spread over multiple social networks. In this paper, we propose a social network-agnostic approach for collecting recent images and videos which can be potentially attached to an event. These media items can be used for the automatic generation of visual summaries in the form of media galleries. Our approach includes the alignment of the varying search result formats of different social networks, while putting media items in correspondence with the status updates and stories they are related to. More precisely we leverage on: (i) visual features from media items, (ii) textual features from status updates, and (iii) social features from social networks to interpret, deduplicate, cluster, and visualize media items. We address the technical details of media item extraction and media item processing, discuss criteria for media item filtering and envision several visualization options for media presentation. Our evaluation is divided into two parts: first we assess the performances of the image process deduplication and then we propose a human evaluation of the summary creation compared with Teleportd and Twitter media galleries. A demo of our approach is publicly available at <http://eventmedia.eurecom.fr/media-finder>.

Categories and Subject Descriptors

H.3.4 [Information Systems]: Information Storage and Retrieval—World Wide Web

Keywords

Media Retrieval, Event Identification, Visual Summarization, Media Gallery, Social Networks

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

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

Copyright 2012 ACM 978-1-4503-1586-9/12/10 ...\$10.00.

1. INTRODUCTION

The widespread availability of mobile phones with higher resolution cameras has transformed citizens into media publishers and witnesses who are used to comment and share event-related media on social networks. Some examples with global impact include the shootings in Utøya, which first appeared on Twitter, the capture and arrest of Muammar Gaddafi, which first appeared on YouTube, or the emergency ditching of a plane in the Hudson river, which first appeared on Twitpic. Some news agencies¹ have even specialized in aggregating and brokering this user-generated content. In this paper, we propose a new approach for retrieving all those event-related media items that are being published by users on several social networks.

The literature in this field includes several works for event detection based on users' activities. Hence, many domain-specific methods for event detection showing good accuracy have been proposed, for example in the sports domain [?]. However, the challenge in this field is to find methods that are content-agnostic. A first category of related work includes research that aims to collect, align, and organize media for trends or events. Liu *et al.* combine semantic inference and visual analysis to automatically find media that illustrate events [?]. They interlink large datasets of event metadata and media with the Linking Open Data Cloud [?]. Data reconciliation uses visual, temporal, and spatial similarity measures for attaching photo streams to events [?]. Other ways to collect and order media from social networks are based on user-driven metadata such as geospatial information [?]. Another relevant work area is duplicate and near-duplicate media detection. Work on ordinal measures for image correspondence started in the last decade of the 20th century [?]. Recently, Chum *et al.* have proposed a near-duplicate image detection method using MinHash and TF-IDF weighting [?]. A method for both images and video has been proposed by Yang *et al.* [?]. Specialized methods for video exist as well [?], an excellent survey of which has been conducted by Lian *et al.* [?].

Beside the event detection task, numerous efforts have been spent to tackle the event summary creation. Capturing life moments and building narratives using social networks is addressed in [?], where the authors investigate the interaction between event stories and the use of social networks to

¹e.g., Citizenside (<http://www.citizenside.com>)

tell them. They proposed Storify², a Web application which supports users to perform story telling and in particular: (1) sorting and organizing the items of an experience similar to the elements of a story, (2) communicating and discussing strategies on how to guide a user towards an intended experience. The overall storytelling creation is supervised by the user, who describes the story as a crafted experience [?]. Streams of news flow through social platforms such as Twitter and YouTube. Getting the big picture from them is the objective of Storyful³. This application allows the user to navigate through the story created by other users or to create his own, aggregating content from different social networks. These two approaches require user interaction, making summary creation a supervised task. Unsupervised approaches are proposed by Teleportd⁴ and Twitter with photo and video galleries where a user can basically create a real-time gallery according to the main event he has searched. Both search APIs allow to collect the list of involved items according to the keyword or hashtag.

In this paper, we tackle the challenge of reconciling the social media items that could illustrate not only events but anything trending within a social network. We then propose visual summaries of these events, applying post-processing techniques such as image deduplication in media galleries and named-entity recognition techniques for organizing the processed media items. Finally, we propose a two-fold evaluation: first, we have assessed the ability of the system to remove duplicates and second, we have performed a user evaluation involving 7 participants that judged the relevance of the media items and the illustrativeness of the galleries. We compare our system with the Teleport and Twitter photo gallery that provide similar services. We want to emphasize that we do *not* perform event detection: the events we are dealing with are known in advance and we use specific human-chosen search terms to find illustrating media. As opposed to ImageCLEF and TRECVID, which both evaluate content-based analysis approaches for multimedia information retrieval, our approach deals with the separate additional task of extracting media from social networks in the first place.

The remainder of this paper is organized as follows. In Section 2, we provide a classification of social networks according to the media support they provide. In Section 3, we detail the process of extraction and reconciliation of media items from different sources. In Section 4, we describe the operations we perform to filter and index media items to improve the visual summaries that are then presented in a Web interface. We present the results of our two-fold evaluation in Section 5. Finally, we give our conclusions and outline future work in Section 6.

2. SOCIAL NETWORKS

A social network is an online service or media platform that focuses on building and reflecting social relationships among people who share interests and/or activities. The boundary between social networks and media platforms is rather blurry. Several media sharing platforms, such as YouTube, enable people to upload content and optionally allow other people to react to this content in the form of

comments, likes or dislikes. On other social networks (*e.g.*, Facebook), users can update their status, post links to stories, upload media content and also give readers the option to react. Finally, there are hybrid clients (*e.g.*, TweetDeck for Twitter using Twitpic) where social networks integrate with media platforms typically via third party applications. Therefore, we consider three types of support of media items with social networks:

- *First-order support*: The social network is centered on media items and posting requires the inclusion of a media item (*e.g.*, YouTube, Flickr);
- *Second-order support*: The social network lets users upload media items but it is also possible to post only textual messages (*e.g.*, Facebook);
- *Third-order support*: The social network has no direct support for media items but relies on third party application to host media items, which are linked to the status update (*e.g.*, Twitter before the introduction of native photo support).

In this paper, we consider 11 different social networks that all have powerful and stable APIs and, together, represent the majority of the market. The criteria for including media sharing platforms follow a study performed by the company Sysomos, specialized in social media monitoring and analytics [?]. Table 1 lists these platforms according to the categorization defined above.

3. MEDIA COLLECTOR

We have developed a media collector composed of media item extractors for all the media sharing networks listed in Table 1. The media collector takes as input a search term, *e.g.*, “io12” for the Google I/O event, then a parallel key-search is performed to all the social networks. Each platform has a 30 second timeout window to deliver its results. When the timeout has expired, or when all social networks have responded, a unified output is delivered. Figure 1 depicts the overall architecture of the media collector. It proposes a common alignment schema for all social networks in order to be agnostic of a particular social network. The resulting metadata for a media item are detailed below (URI examples for the search by “io12” keyword, shortened for legibility):

Media URL Deep link to the media item (*e.g.*, <http://goo.gl/zI2Tg>).

Type Type of the media item (photo or video).

Story URL URL of the micropost where the media item appeared (*e.g.*, <http://goo.gl/R41v8>).

Message Text Description of the micropost in raw format.

Clean Cleaned text description of the micropost where some characters are removed.

User URL of the micropost author (*e.g.*, <http://goo.gl/zI2Tg>).

Timestamp Reference time when the micropost was authored or the media item was uploaded.

²<http://storify.com>

³<http://storyful.com>

⁴<http://teleportd.com>

Social Network	URL	Category	Comment
Google+ MySpace Facebook Twitter	http://google.com/+ http://myspace.com http://facebook.com http://twitter.com	second-order second-order second-order second-/third-order	Links to media items are returned via the Google+ API. Links to media items are returned via the MySpace API. Links to media items are returned via the Facebook API. In second order mode, links to media items are returned via the Twitter API. In third order mode, Web scraping or media platform API usage are necessary to retrieve links to media items. Many people use Twitter in third order mode with other media platforms.
Instagram YouTube Flickr MobyPicture	http://instagram.com http://youtube.com http://flickr.com http://mobypicture.com	first-order first-order first-order first-order	Links to media items are returned via the Instagram API. Links to media items are returned via the YouTube API. Links to media items are returned via the Flickr API. Media platform for Twitter. Links to media items are returned via the MobyPicture API.
Twitpic	http://twitpic.com	first-order	Media platform for Twitter. Links to media items must be retrieved via Web scraping.
img.ly	http://img.ly	first-order	Media platform for Twitter. Links to media items must be retrieved via Web scraping.
yfrog	http://yfrog.com	first-order	Media platform for Twitter. Links to media items must be retrieved via Web scraping.

Table 1: Social networks with different support levels for media items and techniques needed to retrieve them




Figure 1: The media collector architecture: it proposes a hybrid approach for the media item extraction process using a combination of API access and Web scraping.

4. MEDIA POST-PROCESSING AND VISUALIZATION

After collecting media items from the different social networks, we perform two extra processes in order to improve the way the results are presented to the end users. The first one consists of deduplicating the media which have been retrieved. This is a crucial step for avoiding redundant information that appears for example when retrieving various versions of the same image. Our approach accomplishes this objective by using previous existing content-based image retrieval (CBIR) techniques for searching images in large-scale databases, which rely on features like color, texture, and shape [?].

The second step is focused on inferring facets that help the user to select those media items from the results that could be more interesting for him. In order to do that, we perform named-entity recognition over the available textual information using the NERD framework [?]. The extracted named entities are then used as facets that can narrow down the search space for visual factors, enabling cross-fertilization between the textual and visual analysis, which results in effective context-aware analysis possibilities [?]. For example, when searching with the term “politicians”, location can be

proposed as a suitable facet because most of the obtained results will have references to the countries where they belong to.

Once the media items are collected and post-processed, we present them under the form of media galleries. Figure 2 shows a Web user interface that enables a user to explore the set of retrieved media items and related metadata such as timestamp, original story and source platform. We propose three different media item visualizations: *i*) grid visualization, where media items are arranged in a grid as soon as the media collector retrieves them; *ii*) timeline visualization, where media items are sorted chronologically; *iii*) source visualization, where media items are grouped in a bar chart according to the source they are coming from. Aesthetic principles for automatic media gallery layout have been defined in [?].

5. EVALUATION

We performed a two-fold evaluation:

1. We analyzed the performance of the media post-processing only for the media deduplication
2. We conducted a human-based evaluation in order to




Figure 2: The Media-Finder Web application: on the left side, the media gallery resulting from a search using the term “io12”. On the right, a UI dialog zooms over a particular media item showing its related textual information, the date and hour of publication and the platform where it has been published, here, Instagram.

compare our media galleries with results generated by Teleportd and Twitter.

5.1 Media post-processing analysis

We evaluated this framework during the period of January 10 to 19, 2012 in which we have selected nine events. For those events, we collected media items and microposts using the media collector⁵. The search terms used are: *i*) Assad Speech, *ii*) CES Las Vegas, *iii*) Costa Concordia Disaster, *iv*) Cut the Rope Launch, *v*) Dixville Notch, *vi*) Free Mobile Launch, *vii*) Blackout SOPA, *viii*) Ubuntu TV Launch, *ix*) Christian Wulff Case. The dataset collected contains 448 images with an average file size of $\sim 0.7\text{MB}$ and 143 videos. Some videos are no longer available due to either an account termination or a video takedown by the user (Assad, Dixville). We observed that the process of image deduplication is by no means a solved issue. We used the Photo-Sweeper CBIR-based image duplication detection software that allows for manual algorithm and threshold selection to detect duplicates in the dataset (Table 2). For each event, we have manually selected the best settings to limit the number of duplicate misses and false positives. The main problem with the dataset is its diversity. It ranges from entirely sharp screenshots in all sorts of formats (*e.g.*, screenshots of the Google homepage for the Blackout SOPA event), to blurry cell phone images in standard photo formats (*e.g.*, photos of the stage for the Free Mobile Launch event). A common performance tweak to speed up the duplication detection process is to shrink images to quadratic bitmaps. In the context of our dataset, however, this approach is counter productive, as a screenshot of a rectangular IAB 728×90 “leaderboard” banner is treated the same as a standard 3.1 megapixels (2048×1536) cell phone photo. In practice, shrinking a wide rectangular banner to a square led to many incorrect results requiring manual deduplication with the Blackout SOPA event.

5.2 User evaluation

We conducted a second user evaluation in order to compare the relevance and illustrativeness of the media galleries

⁵The dataset as well as the visual summaries are available at <http://webmasterapp.net/social/acmmm2012/>

	Media-Finder	Teleportd	Twitter	#total
google i/o	108	20	96	224
io12	69	20	98	187

Table 3: Number of media items retrieved by each service per search term and total of unique items

generated by the Media-Finder with respect to two other similar tools, namely Teleportd and Twitter media galleries. First, we generated media galleries for these three systems on the Google I/O event⁶ using the keywords “google i/o” and “io12” as search terms. Table 3 shows the number of items collected by the three services and the total number of unique items (*i.e.* we remove the duplicates). For *google io* 218 unique items were found by the three service and for *io12*, we retrieved 187 unique items.

Second, we created two surveys composed of three different sections, one per platform used. In order do not bias the participant, we called these sections Gallery A, B and C. For each gallery, participants were asked to judge the relevance of each media item to the event by ticking a box. We also asked each participant to answer three additional questions: *Q1: how illustrative this gallery is for this event?*, *Q2: how visually diverse this gallery is for this event?*, *Q3* a free comment space. We used a Likert scale [?] from 1-7 to assess the results of *Q1* and *Q2*. The surveys are available at <http://goo.gl/QzSM6> and <http://goo.gl/7ov6Q>.

The results provided by the social platforms involved in this experiment are volatile, *i.e.* the deep link URI of the media item expires after a short period of time. As introduced in Section 3, besides the mediaurl some platforms provide the storyurl. We provided this information in order to help the participant in judging the relevance of the media item in case this one was already not viewable anymore. Each participant followed a short training, where we explained the goal of the evaluation and we recommended to judge thoroughly each media item. In total, 7 participants (6 male, 1 female) were involved in this study divided into two groups G1 (4 participants) and G2 (3 participants). The first group

⁶en.wikipedia.org/wiki/Google_io

Event	Exact Duplicate I	Loose Duplicate I	Exact Duplicate V	Loose Duplicate V
Assad Speech	0 image in 0 seq	2 images in 1 seq	0 video in 0 seq	2 videos in 1 seq
CES Las Vegas	0 image in 0 seq	9 images in 3 seq	0 video in 0 seq	2 videos in 1 seq
Costa Concordia	0 image in 0 seq	6 images in 3 seq	0 video in 0 seq	0 video in 0 seq
Cut the Rope Launch	2 images in 1 seq	15 images in 5 seq	0 video in 0 seq	14 videos in 3 seq
Dixville Notch	2 images in 1 seq	2 images in 1 seq	2 videos in 1 seq	0 video in 0 seq
Free Mobile Launch	2 images in 1 seq	16 images in 7 seq	0 video in 0 seq	0 video in 0 seq
Blackout SOPA	0 image in 0 seq	14 images in 4 seq	2 videos in 1 seq	0 video in 0 seq
Ubuntu TV Launch	0 image in 0 seq	5 images in 1 seq	4 videos in 1 seq	9 videos in 4 seq
Christian Wulff Case	4 images in 2 seq	0 image in 0 seq	0 video in 0 seq	0 video in 0 seq

Table 2: Exact and loose duplicate images (I) and videos (V) per event

		google i/o			io12		
		relevance	Q1	Q2	relevance	Q1	Q2
Total	Media-finder	0.1996	0.8997	1.2724	0.1873	1.1547	1.211
	Teleportd	0.1842	0.9512	1.2724	0.2358	1.2724	1.7995
	Twitter	0.2182	0.9759	0.8997	0.2471	0.6901	0.9512
G1	Media-finder	0.0711	0.9574	1.291	0.0942	0.5774	0
	Teleportd	0.1190	0.5	0.8165	0.1843	0.5	1.291
	Twitter	0.1619	0	0.8165	0.2317	0.5774	1
G2	Media-finder	0.1002	1	1.5275	0.2151	1.5275	1.5275
	Teleportd	0.2517	1.1547	1.5275	0.3215	2	2.6458
	Twitter	0.1286	1.1547	1.4142	0.2746	0.5774	1

Table 4: Standard deviation among the participants for each service within each group

	google i/o	io12
Media-Finder	0.4954	0.3728
Teleportd	0.0917	0.1081
Twitter	0.4404	0.5297

Table 5: Weighted score computed considering the number of items a service can retrieve according to an ideal dataset obtained by combining the three minus the duplicates.

used both the media item and the media story to judge its relevance, while the second group only used the media item. Table 4 shows the standard deviation of G1 and G2. We observe general agreement within groups.

We compute the average relevance score of each service weighting by the total number of unique items (Table 5). Table 6 shows the aggregated results of relevance, Q1 and Q2 question weighted for each group. We conclude that both the Media-Finder and the Twitter media galleries are considered relevant and illustrative while the Media-Finder galleries seem to be slightly more visually diverse (Q2).

The survey finally asked each participant to express his general feeling about the media galleries. The general summary from them is that Teleportd is strongly penalized by the reduce number of items it can retrieve. The following comment well describes the general impression: *“This gallery provides only photos an no videos, but all thumbnails are available! Only the author seems to be available, not the story behind the media, making it hard for judging the relevance to the event. I don’t have a feeling of what were the key announcements of this tech conference. However, I saw unique photos not seen in the other galleries that*

seem relevant”. A more exhaustive summary is provided by Twitter and Media-Finder. For the “io12” keyword search a participant wrote *“I appreciate the diversity of languages of the stories behind the media. However, only images were shown, no video I found almost all media relevant to this event. The main topics covered seem to be: the new Nexus device, illustration of Android, the new google glasses, some announcements of apps on Android (e.g., Pinterest), a lot of pictures of the talks and the venue ... I have the feeling that I could guess the key announcements of this tech conference.”*; while for the same search, the same participants reports about Media-Finder: *“I appreciate the diversity of media (images and videos). I found a lot of media relevant to this event ... however, I was intrigued by many media Thai stories that I couldn’t understand but that do not seem to be relevant for this event. [...] I have the feeling that I could guess the key announcements of this tech event.”*

6. CONCLUSION AND FUTURE WORK

In this paper, we presented a generic media collector for retrieving media items shared on social networks and illustrating daily life moments. We proposed a common schema in order to align the search results of these platforms. We further described a full processing chain for the media items that includes named-entity extraction and media item deduplication. We presented a user interface that displays media galleries and we assessed the performance of image deduplication. Finally, we conducted a user evaluation to compare the usefulness of the media galleries with similar services, namely Teleportd and Twitter. The analysis of this evaluation shows that the media items retrieved by Media-Finder are more diverse, although Twitter media galleries are also appreciated. We also found out that Media-Finder lacked the option to consider media that are natively hosted on

		google i/o			io12		
		relevance	Q1	Q2	relevance	Q1	Q2
Total	Media-finder	0.2031	2.4063	2.6894	0.1753	1.8649	1.9892
	Teleportd	0.043	0.3408	0.4194	0.0502	0.3861	0.5714
	Twitter	0.2254	2.3906	2.5793	0.3027	3.2546	3.027
G1	Media-finder	0.2764	2.353	2.7248	0.2122	2.0516	2.2378
	Teleportd	0.0482	0.2982	0.3669	0.0446	0.3516	0.5945
	Twitter	0.2856	2.6422	2.6422	0.3432	3.4432	2.9135
G2	Media-finder	0.10552	2.4771	2.6422	0.1261	1.6162	1.7405
	Teleportd	0.0352	0.3976	0.4892	0.0577	0.4324	0.5405
	Twitter	0.1453	2.055	2.6422	0.2486	3.0018	3.1784

Table 6: Aggregate results per service according to the relevance of retrieved results and Q1 and Q2 scales

Twitter and are clearly becoming the preferred option. An update of the Media-Finder may soon outperform Twitter in terms of media coverage. Multimedia analysis techniques and natural language processing (to parse textual information related to each media item published) should be better integrated in the processing chain as this will create a multi-modal environment where different factors are used to organize social content. Content deduplication and visual quality metrics (sharpness, contrast, etc.) can be used to further cluster media items. The identification of original content can allow users to choose a balance between popularity (favor omnipresent content) and originality (promote rare content), tasks that we plan to address in future work.

Acknowledgments

This work was partially supported by LinkedTV, OpenSEM projects and the European Commission under Grant No. 248296 FP7 I-SEARCH.

7. REFERENCES

- [1] B. Atasoy and J.-B. Martens. STORIFY: a tool to assist design teams in envisioning and discussing user experience. In *Annual Conference on Human Factors in Computing Systems*, pages 2263–2268, 2011.
- [2] D. Bhat and S. Nayar. Ordinal measures for image correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(4):415–423.
- [3] O. Chum, J. Philbin, and A. Zisserman. Near Duplicate Image Detection: min-Hash and tf-idf Weighting. In *19th British Machine Vision Conference (BMVC’08)*, pages 493–502, London, UK, 2008.
- [4] D. J. Crandall, L. Backstrom, D. Huttenlocher, and J. Kleinberg. Mapping the world’s photos. In *18th International Conference on World Wide Web (WWW’09)*, pages 761–770, Madrid, Spain, 2009.
- [5] R. Cyganiak and A. Jentzsch. Linking Open Data Cloud diagram, 2011. <http://lod-cloud.net/>.
- [6] M. Hassenzahl. Experience Design: Technology for All the Right Reasons. *Synthesis Lectures on Human-Centered Informatics*, 2010.
- [7] S. Levine. How People Currently Share Pictures On Twitter, 2011. <http://blog.sysomos.com/2011/06/02/>.
- [8] B. Li and M. Sezan. Event detection and summarization in American football broadcast video. In *Storage and Retrieval for Media Database*, pages 202–213, 2002.
- [9] S. Lian, N. Nikolaidis, and H. Sencar. Content-Based Video Copy Detection - A Survey. In *Intelligent Multimedia Analysis for Security Applications*, pages 253–273. Springer, 2010.
- [10] R. Likert. A technique for the measurement of attitudes. *Archives of Psychology*, 22(140):1–55, 1932.
- [11] X. Liu, R. Troncy, and B. Huet. Finding Media Illustrating Events. In *1st ACM International Conference on Multimedia Retrieval (ICMR’11)*, Trento, Italy, 2011.
- [12] H.-S. Min, J. Y. Choi, W. D. Neve, and Y. M. Ro. Bimodal fusion of low-level visual features and high-level semantic features for near-duplicate video clip detection. *Signal Processing: Image Communication*, 26(10):612–627, 2011.
- [13] N. P. Ramaiah and C. K. Mohan. De-duplication of Photograph Images using Histogram Refinement. In *IEEE Recent Advances in Intelligent Computational Systems (RAICS’11)*, pages 391–395, 2011.
- [14] G. Rizzo and R. Troncy. NERD: A Framework for Unifying Named Entity Recognition and Disambiguation Extraction Tools. In *Conference of the European Chapter of the Association for computational Linguistics (EACL’12)*, 2012.
- [15] 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. IEEE, 2012.
- [16] R. Verborgh, D. V. Deursen, E. Mannens, C. Poppe, and R. V. de Walle. Enabling context-aware multimedia annotation by a novel generic semantic problem-solving platform. *Multimedia Tools and Applications*, 2012.
- [17] J. Yang, J. Luo, J. Yu, and T. S. Huang. Photo stream alignment for collaborative photo collection and sharing in social media. In *3rd ACM International Workshop on Social Media (WSM’11)*, pages 41–46, Scottsdale, Arizona, USA, 2011.
- [18] X. Yang, Q. Zhu, and K.-T. Cheng. Near-duplicate detection for images and videos. In *1st ACM Workshop on Large-Scale Multimedia Retrieval and Mining (LS-MMRM’09)*, pages 73–80, 2009.