

ACM MM 2015 Grand Challenge

Yahoo-flickr Event Summarization Challenge

Challenge overview

With thousands of photos and videos being uploaded every minute to social photography sharing websites – sometimes automatically after capture without a lot of useful metadata or annotations available – it is not easy to find structure amongst all multimedia items.

In this grand challenge the objective is to automatically uncover structure within a collection of 100 million photos/videos in the form of detecting and identifying events, and summarizing them succinctly for consumer consumption.

Achieving the objective requires successfully bringing together multiple modalities (textual, visual, aural) and multiple levels of understanding (concepts, topics, semantics, quality), as well as multiple ways of showcasing (presentation, visualization, interaction).

Task description

Given a collection of photos and videos, the goal is to automatically separate them into a set of possibly overlapping events. The events to detect can range from private one-off parties (e.g. a birthday) to public multi-day occasions (e.g. a music festival), where the events can be recurring (e.g. every year), their locations dispersed (e.g. across various venues), and their structure can be multi-level (e.g. sub-events within a larger event).

The photos/videos in the dataset are to be considered in order of upload, such that proposed systems ideally are able to incrementally summarize events as they occur (although this is not a requirement and systems may recompute at each moment in time), as well as retrospectively update previous summaries when new, better, media is uploaded. The ultimate objective is to automatically produce a summary of each detected event using a small set of representative photos and/or videos.

Dataset

The dataset to be used is the publicly available YFCC100M collection, which contains approximately 99 million photos and 1 million videos, each of which has been assigned a Creative Commons license. The dataset is available at <http://bit.ly/yfcc100md>.

Events

To evaluate the proposed automatic event detection and summarization systems, the organizers have provided several events of which and several timestamps at which the participants have to present the current state of their proposed system in terms of the summaries of those events that occurred before, or were still ongoing during, the specified moments in time. The timestamps listed below refer to the number of seconds that have passed since 1 January 1970.

For each given event and timestamp, your aim is to automatically detect and summarize different aspects of the event, by only considering photos uploaded on

or before the timestamp. The events and timestamps are, in no particular order of importance:

Occupy Movement

The Occupy Movement¹ is the international branch of the Occupy Wall Street movement that protests against social and economic inequality around the world, its primary goal being to make the economic and political relations in all societies less vertically hierarchical and more flatly distributed. This is a global multi-day event.

Timestamp 1: 1316390400 (19 September 2011 @ 12:00am UTC)
Timestamp 2: 1320105600 (1 November 2011 @ 12:00am UTC)
Timestamp 3: 1348012800 (19 September 2012 @ 12:00am UTC)

Batkid

Batkid² is the superhero name of Miles Scott, an American child and cancer survivor who is in remission. His wish was to be “Batkid”, a sidekick of the eponymous comic book superhero Batman, the subject of books, radio, television, and films. Once the request went out, thousands of volunteers, city officials, businesses and supporters rallied to turn the city of San Francisco into “Gotham City” for one of the largest and most elaborate Make-A-Wish projects ever staged. This is a local single-day event.

Timestamp 1: 1384556400 (15 November 2013 @ 11:00pm UTC)
Timestamp 2: 1384592400 (16 November 2013 @ 9:00am UTC)
Timestamp 3: 1388534400 (1 January 2014 @ 12:00am UTC)

Olympic Games

The modern Olympic Games³ are the leading international sporting event featuring summer and winter sports competitions in which thousands of athletes from around the world participate in a variety of competitions. This is a global and bi-yearly recurring multi-day event that encompasses many sub-events, each of which can again be broken down into mini-events (e.g. 2014 Sochi Winter Olympics → Snowboarding → Half pipe → Women).

Timestamp 1: 1340582400 (25 June 2012 @ 12:00am UTC)
Timestamp 2: 1349049600 (1 October 2012 @ 12:00am UTC)
Timestamp 3: 1420070400 (1 January 2015 @ 12:00am UTC)

Eyjafjallajökull Eruption

Over the past 500 years, Iceland’s volcanoes have erupted a third of the total global lava output. The eruption under Eyjafjallajökull (“glacier of Eyjafjöll”) in 2010⁴ was notable because the volcanic ash plume disrupted air travel in northern Europe for several weeks. This is a seemingly local event with a short duration, but it affected many around the globe for several weeks.

Timestamp 1: 1270080000 (1 April 2010 @ 12:00am UTC)
Timestamp 2: 1271548800 (18 April 2010 @ 12:00am UTC)
Timestamp 3: 1275350400 (1 June 2010 @ 12:00am UTC)

¹ https://en.wikipedia.org/wiki/Occupy_movement

² <https://en.wikipedia.org/wiki/Batkid>

³ https://en.wikipedia.org/wiki/Olympic_Games

⁴ https://en.wikipedia.org/wiki/2010_eruptions_of_Eyjafjallajökull

Holi

Holi⁵ is a spring festival, also known as the festival of colors or the festival of love. This is a predominantly regional and yearly recurring event with a short duration.

Timestamp 1: 1236470400 (8 March 2009 @ 12:00am UTC)

Timestamp 2: 1238544000 (1 April 2009 @ 12:00am UTC)

Timestamp 3: 1301616000 (1 April 2011 @ 12:00am UTC)

Byron Bay Bluesfest

The Byron Bay Bluesfest⁶, formerly the East Coast International Blues & Roots Music Festival, is an annual music festival held for five days over the Easter long weekend at Tyagarah, near Byron Bay, New South Wales, Australia. It features a large selection of blues and roots performers, both international and Australian. This is a local and yearly recurring multi-day event that encompasses several sub-events.

Timestamp 1: 1306886400 (1 June 2011 @ 12:00am UTC)

Timestamp 2: 1420070400 (1 January 2015 @ 12:00am UTC)

Hanami

Hanami⁷ is the Japanese traditional custom of enjoying the transient beauty of flowers, flowers in this case almost always referring to those of the cherry (“sakura”) or, less frequently, plum (“ume”) trees. From the end of March to early May, sakura bloom all over Japan, and around the first of February on the island of Okinawa. The blossom forecast is announced each year by the weather bureau, and is watched carefully by those planning hanami as the blossoms only last a week or two. This is a regional and yearly recurring event with a medium duration.

Timestamp 1: 1359676800 (1 February 2013 @ 12:00am UTC)

Timestamp 2: 1370044800 (1 June 2013 @ 12:00am UTC)

Submissions

The submission should include the URL that points to the proposed system, as well as a 4-page paper that describes the system in detail. Please follow the guidelines of ACM Multimedia 2015 Grand Challenge for the paper submission.

Evaluation metric

Your paper and system submission will be peer-reviewed and acceptance will principally depend on the proposed technical contributions to the state of the art. At this point, your system has the role of supporting your paper but does not yet need to be final: it should function and summarize the pre-selected events at the given timestamps, but you will have the opportunity to improve it upon acceptance to the conference.

At the conference, the final ranking of submissions will be determined by a panel of judges from the multimedia community, which will evaluate the submissions based on the novelty and originality of their approach, perceived accuracy of the summaries for the preselected events at the preselected timestamps, and the perceived quality of how the summaries are presented.

⁵ <http://en.wikipedia.org/wiki/Holi>

⁶ http://en.wikipedia.org/wiki/Bluesfest_Byron_Bay

⁷ <http://en.wikipedia.org/wiki/Hanami>

Reviewers and judges alike will focus on how well your system would be suited to automatically and incrementally uncover structure within a large collection of media, rather than it being primed for only detecting the events that are part of this challenge.

Notes:

- Not all photos/videos can be, or need to be, assigned to an event, while others may be assigned to multiple simultaneously occurring events.
- We intentionally do not provide specific guidelines on how to present the event summaries in order to encourage creative solutions. The summaries may be static or dynamic, interactive or cued, generic or personalized, etc. – anything is possible.
- We are looking for promising, high-quality solutions that theoretically and practically will be able to solve the event detection and summarization problem at web-scale, either now or sometime in the near future. It is not necessary that all photos/videos from the dataset are processed, but it is considered a plus.

Awards and prizes

The top three teams having the best solutions to the grand challenge will receive award certificates and cash prizes: 1st wins \$250, 2nd wins \$150, and 3rd wins \$100 (currency is USD).

Participation

We encourage, but do not require, team-based participation. Each team can have one or more members, and an individual can be a member of multiple teams. No two teams, however, can have more than half of their members in common.

Timeline

- 1 May 2015: Preselected events and timestamps announced.
- 6 July 2015: Paper submissions due.
- 27 July 2015: Notification of acceptance.
- 15 August 2015: Camera-ready submissions due.
- 26-30 October 2015: Challenge solutions to be presented at ACM MM.

Note that the dataset is already available, so teams may already start addressing the challenge right away.

Organizers

Bart Thomee (Yahoo Labs) – bthomee@yahoo-inc.com

David A. Shamma (Yahoo Labs) – aymans@acm.org