

NicePic!: A System for Extracting Attractive Photos from Flickr Streams

Sergej Zerr*, Stefan Siersdorfer*, Jose San Pedro**, Jonathon Hare***, Xiaofei Zhu*

*L3S Research Center, Hannover, Germany
{zerr,siersdorfer,zhu}@L3S.de

**Telefonica Research, Barcelona, Spain
jsanpedro@mac.com

***Electronics and Computer Science, University of Southampton, Southampton, UK
jsh2@ecs.soton.ac.uk

ABSTRACT

A large number of images are continuously uploaded to popular photo sharing websites and online social communities. In this demonstration we show a novel application which automatically classifies images in a live photo stream according to their attractiveness for the community, based on a number of visual and textual features. The system effectively introduces an additional facet to browse and explore photo collections by highlighting the most attractive photographs and demoting the least attractive.

Categories and Subject Descriptors

H.3.5 [Online Information Services]: Web-based services

1. INTRODUCTION

The rapid increase in size of online communities and the availability of large amounts of shared visual data make discovering relevant content a difficult task. For instance, thousands of new photos are uploaded to Flickr every minute making effective automatic content filtering techniques a necessity.

Flickr photos are accompanied by a variety of metadata such as tags, number of views, user comments, upload date, etc. The Flickr search interface exploits the explicit and implicit ratings in the metadata to infer rankings. For instance, the number of views is an indicator for the popularity of a photo. Adding a photo to one's favorite list is probably the most direct positive indicator of relevance assignment in Flickr, and is an explicit expression of interest in the photo.

However, for recently uploaded photos community feedback in any form might not yet be available. Furthermore, many photos are just sparsely annotated which might prevent text-based search and mining methods from retrieving this potentially attractive content.

Permission to make digital or hard copies of part or all 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. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

SIGIR '14, July 6–11, 2014, Gold Coast, Queensland, Australia.

ACM 978-1-4503-2257-7/14/07.

<http://dx.doi.org/10.1145/2600428.2611183>.

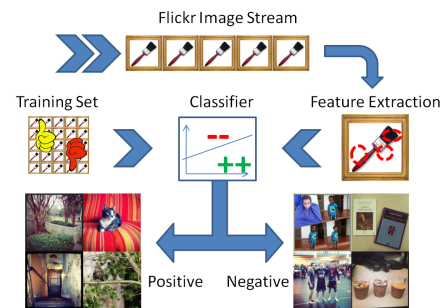


Figure 1: System architecture overview.

In this work we demonstrate the NicePic!¹ application, based on a web service for automatically classifying and ranking photos according to their attractiveness. We exploit the vast amount of social feedback available in Flickr, to obtain a training set of photos considered as more or less attractive by the community. This allows us to build classification and regression models based on multi-modal visual and textual features, and to apply them to identify new attractive content. In a wider system context, such techniques can be useful to enhance ranking functions for photo search [4], and, more generally, to complement mining and retrieval methods based on text, other metadata and social dimensions.

NicePic! allows users to explore the live flickr photo stream in an attractiveness-centered way. NicePic! shows the top photos submitted to Flickr during the last hour, day and week according to our aesthetic inference model, enabling users to rapidly see the best content uploaded to Flickr in the recent past. In addition, users can introduce query terms in the search box to get results relevant to their information needs. These results are organized into two columns, separating the most attractive from the least attractive photographs.

2. SYSTEM ARCHITECTURE

This section describes the main components of the system. The architecture is illustrated in Figure 2. Firstly, we build a training set of images with and without favorite assignments. In the second step we extract visual, and if available, textual *features*, from the images. We then train a SVM *classifier* which is used by our system for identifying potentially attractive visual content. Finally, the user can

¹Demo and data are available at <http://l3s.de/nicepic/>.

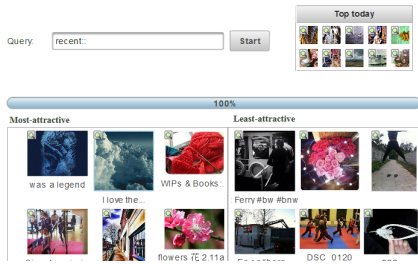


Figure 2: Search Graphical User Interface.

access the application from arbitrary *clients*, including desktops and mobile devices featuring web browsing capabilities. In the following we provide a brief overview of the system components and show how results are presented to the user. A fully detailed description of the underlying scientific approach can be found in our work [3].

Data. We randomly selected time periods of 20 minutes from a time span of 5 years 2005-2010. From each of the periods we selected at most 5 pictures with the highest number of favorite assignments as positive example as well as the same number of photos without favorite assignments as negative examples. We stopped after obtaining a set of 200,000 photos from each class.

Features. Even though aesthetic and artistic quality cannot be quantitatively computed, it has been shown that certain visual features [1] of images have significant correlation with them. For instance, appealing images tend to have higher *colorfulness*, increased *contrast* and *sharpness*. Textual features like title and tags can also provide precise information about the image quality. These correlations were described in [3] which forms the base for this demonstration.

Classification. In the next step we built a classifier using the SVMlight [2] classification software and the dataset of consisting of pictures with and without favorite assignments as described in Section 2. Our quality measures were the precision-recall curves as well as the precision-recall break-even points for these curves. The break-even point (BEP) is the precision/recall value at the point where precision equals recall, which is equal to the F_1 measure and is the harmonic mean of precision and recall. The curves for the classification experiments for selected visual features described in Section 2 are presented in Figure 3. The combination of textual and visual features have shown the best applicability resulting in BEP 0.84. Classification with only visual features alone also produces promising results (BEP 0.67), and can be useful if no or insufficient textual annotations are available as is usually the case for freshly uploaded photos. Extraction of visual features is a computationally intensive process, using only textual features (if available) provides a cheaper alternative with a BEP of 0.79.

Search. In order to create a list of images ranked by quality, we estimated the likelihood of image attractiveness using the output of the SVM classifier trained on a set of images labeled as “attractive” or “non attractive”. One of the three classifiers (textual, visual, or textual+visual) is selected for each image based on the availability of the image features. We use the Flickr API as the underlying search provider for our NicePic! service. It is possible to process the most recent uploaded pictures, a user image stream, or a selection for an arbitrary keyword.




Figure 3: P/R curves for the features and their combination.

3. GUI & DEMONSTRATION OVERVIEW

In the demonstration we will primarily show how the NicePic! search system works. The user interface of the application simply consists of a text box and a keyword search can be performed pressing the “Search” button. The difference to other engines is mainly in the search result representation. NicePic! divides the results into two sets: left (“Most Attractive”) and right (“Least Attractive”). The server in the background continuously obtains the most recent photos from the Flickr stream and classifies them. The GUI also shows the top attractive and non attractive photos in the last week, 24h and 1h as selected by the classifier. Additionally the user can use the keyword search function and obtain the most/least attractive photos for a particular query, or particular user. For example query “recent:” would analyze the most recent photos (the maximum number can be selected by user) and query “user::username” would pick only photos uploaded by a single user.

The user starts with the execution of the search. As soon as the first results are ready, they are immediately displayed. The view with the top 20 most-attractive and least-attractive images is continuously updated until the set is classified completely. The user can move the mouse pointer over a result image and see additional information including the attractiveness value computed by the classifier. Clicking on the image opens a larger view.

Future Work. In the future we plan to experiment with other datasets such as photos taken with mobile devices and exploit other visual features such as SURF, as well as the EXIF data.

4. ACKNOWLEDGMENTS

This work was partially funded by the European Commission FP7 under grant agreement No. 287704 (CUBRIK) and COST Action KEYSTONE IC1302

5. REFERENCES

- [1] J. S. Hare, S. Samangooei, and D. Dupplaw. Openimaj and imagerrier: Java libraries and tools for scalable multimedia analysis and indexing of images. In *ACM MM’11*.
- [2] T. Joachims. Making large-scale support vector machine learning practical. *Advances in kernel methods: support vector learning*, pages 169–184, 1999.
- [3] J. San Pedro and S. Siersdorfer. Ranking and classifying attractiveness of photos in folksonomies. In *WWW ’09*.
- [4] J. San Pedro, T. Yeh, and N. Oliver. Leveraging user comments for aesthetic aware image search reranking. In *WWW ’12*.