

Callisto: Tag Recommendations by Image Content

Mathias Lux¹, Arthur Pitman¹, Oge Marques²

¹ Klagenfurt University, Universitätsstr. 65-67
9020 Klagenfurt, Austria
mlux@itec.uni-klu.ac.at, arthur.pitman@uni-klu.ac.at

² Florida Atlantic University, 777 Glades Rd.
Boca Raton, FL 33431 USA
omarques@fau.edu

Abstract. While assigning uncontrolled keywords to photos – a process called tagging – solves a lot of problems with retrieval of visual information still many photos published on the web go untagged or are tagged with non-descriptive or too few tags. In this extended abstract we demonstrate our tag recommendation prototype *Callisto* incorporating statistical tag co-occurrence as well as image contents. We also outline qualitative experiments that have shown the difference in results between solely statistical and content based approaches.

Keywords: Tagging, recommendation, content based image retrieval

1 Introduction

User generated content is on the rise millions of people upload and publish digital photos on a daily basis. Flickr alone claims thousands of uploaded photos per minute with reports stating numbers from 3,000-6,000 uploads depending on the time of day. Many of these photos are uploaded as-is, with only a minimum amount of metadata attached. This typically includes the EXIF metadata and the name of the image automatically assigned by the camera. A considerable share of photos is tagged, which means the photo is annotated by a set of keywords, but by far not every single photo was tagged by the uploader. As retrieval of the photos heavily depends on the annotations the amount and quality of tags is critical to every retrieval scenario. Therefore we identify a need for applications supporting the users in the tagging process. Classic approaches are limited to co-occurrence analysis of tags and therefore typically suggest tags that are most frequent and not very distinctive such as beautiful, topphoto or flickrdiamond. However, our assumption is that in many scenarios tags describing the actual content of the images, like flower, sunset or wood, are needed. In this paper we present a software prototype, which is able to suggest tags based on (i) one or more initial tags and (ii) image content. This novel combination allows for a more content-related suggestion of tags and might help users to find more and more descriptive tags to annotate their uploaded images.

2 Callisto

Our tag recommendation prototype, called *Callisto*, allows for input of a photo as well as one or more start tags. As depicted in Fig. 1 *Callisto* downloads images based on the start tag and ranks the set of retrieved images based on image content. Photos with highest ranks are taken into account for tag analysis, which then leads to the recommendation of a number of tags (typically 5-10).




Fig. 1. Main four steps in the recommendation process. Images are retrieved based on the provides start tag, then re-ranked based on image content. From this list the top N images are considered for tag analysis, which leads in step 4 to the recommendation of X tags.

The screenshots in Fig. 2 show *Callisto* in use. Left screenshot shows a photo of a fire juggling act, while the right one shows a scene of a person juggling clubs. With the start tag juggling the classical statistical approach (named “Sugg. Stat” in Fig. 2) are the same for both photos: juggler, fire, juggle, balls, etc. However, taking into account content based low level features leads to different recommendations per image (named “Sugg. NCP” in Fig. 2). For the fire juggling photo the tag *fire* is ranked first, while for the club juggling photo *juggler* is the first.

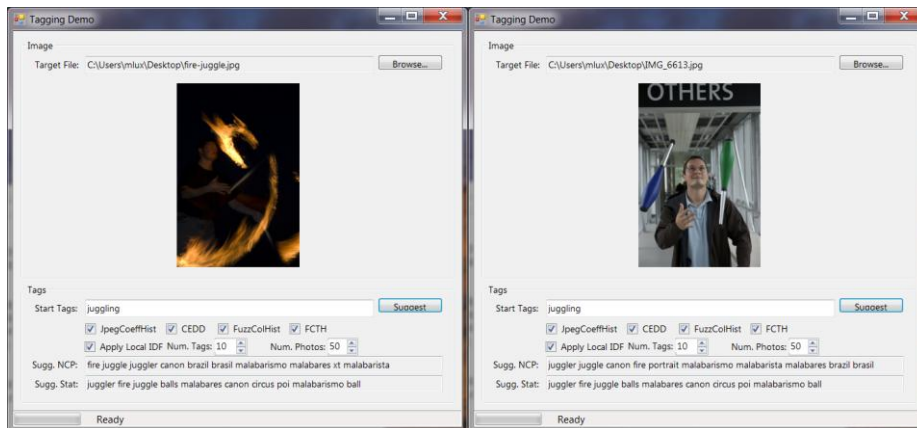


Fig. 2. Screenshot of the *Callisto* tag recommendation system for two different photos with the same start tag.

Callisto utilizes the data of Flickr (<http://www.flickr.com>) and heavily depends on the quality of annotations of the photos retrieved for tag recommendation. However in our experience the tag recommendations got better the more photos were considered in the process (step 1 in Fig. 1). With a critical mass of start tags recommendation also yields good results without a start tag and can be employed for auto-tagging.