

Unsupervised Sentiment Analysis for Social Media Images

Yilin Wang, Suhang Wang, Jiliang Tang, Huan Liu, Baoxin Li

Arizona State University

Tempe, Arizona

{Yilin.Wang.1, Suhang.Wang, Jiliang.Tang, huan.liu, baoxin.li}@asu.edu

Abstract

Recently text-based sentiment prediction has been extensively studied, while image-centric sentiment analysis receives much less attention. In this paper, we study the problem of understanding human sentiments from large-scale social media images, considering both visual content and contextual information, such as comments on the images, captions, etc. The challenge of this problem lies in the “semantic gap” between low-level visual features and higher-level image sentiments. Moreover, the lack of proper annotations/labels in the majority of social media images presents another challenge. To address these two challenges, we propose a novel Unsupervised SEntiment Analysis (USEA) framework for social media images. Our approach exploits relations among visual content and relevant contextual information to bridge the “semantic gap” in prediction of image sentiments. With experiments on two large-scale datasets, we show that the proposed method is effective in addressing the two challenges.

1 Introduction

Recent years have witnessed the explosive popularity of image-sharing services such as Flickr¹ and Instagram². For example, as of 2013, 87 millions of users have registered with Flickr³. Also, it was estimated that about 20 billion Instagram photos are shared to 2014⁴. Since by sharing photos, users could also express opinions or sentiments, social media images provide a potentially rich source for understanding public opinions/sentiments. Such an understanding may in turn benefit or even enable many real-world applications such as advertisement, recommendation, marketing and health-care. The importance of sentiment analysis for social media images has thus attracted increasing attention recently [Yang *et al.*, 2014; You *et al.*, 2015].

¹www.flickr.com

²www.instagram.com

³<http://en.wikipedia.org/wiki/Flickr>

⁴<http://blog.instagram.com/post/80721172292/200m>

Current methods of sentiment analysis for social media images include low-level visual feature based approaches [Jia *et al.*, 2012; Yang *et al.*, 2014], mid-level visual feature based approaches [Borth *et al.*, 2013; Yuan *et al.*, 2013] and deep learning based approaches [You *et al.*, 2015]. The vast majority of existing methods are supervised, relying on labeled images to train sentiment classifiers. Unfortunately, sentiment labels are in general unavailable for social media images, and it is too labor- and time-intensive to obtain labeled sets large enough for robust training. In order to utilize the vast amount of unlabeled social media images, an unsupervised approach would be much more desirable. This paper studies *unsupervised sentiment analysis*.

Typically, visual features such as color histogram, brightness, the presence of objects and visual attributes lack the level of semantic meanings required by sentiment prediction. In supervised case, label information could be directly utilized to build the connection between the visual features and the sentiment labels. Thus, unsupervised sentiment analysis for social media images is inherently more challenging than its supervised counterpart. As images from social media sources are often accompanied by textual information, intuitively such information may be employed. However, textual information accompanying images is often incomplete (e.g., scarce tags) and noisy (e.g., irrelevant comments), and thus often inadequate to support independent sentiment analysis [Hu and Liu, 2004; Hu *et al.*, 2013]. On the other hand, such information can provide much-needed additional semantic information about the underlying images, which may be exploited to enable unsupervised sentiment analysis. How to achieve this is the objective of our approach.

In this paper, we study unsupervised sentiment analysis for social media images with textual information by investigating two related challenges: (1) how to model the interaction between images and textual information systematically so as to support sentiment prediction using both sources of information, and (2) how to use textual information to enable unsupervised sentiment analysis for social media images. In addressing these two challenges, we propose a novel Unsupervised SEntiment Analysis (USEA) framework, which performs sentiment analysis for social media images in an unsupervised fashion. Figure 1 schematically illustrates the difference between the proposed unsupervised method and existing supervised methods. Supervised methods use label informa-

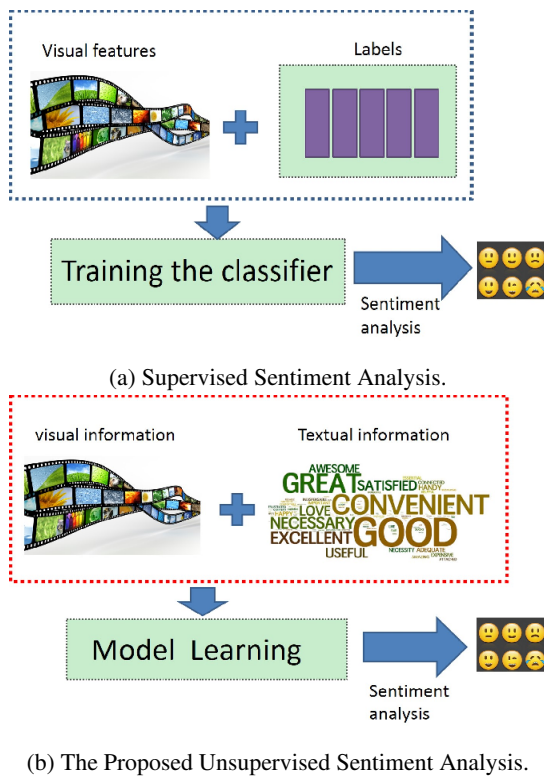


Figure 1: Sentiment Analysis for Social Media Images.

tion to learn a sentiment classifier; while the proposed method does not assume the availability of label information but employ auxiliary textual information. Our main contribution can be summarized as below:

- A principled approach to enable unsupervised sentiment analysis for social media images.
- A novel unsupervised sentiment analysis framework USEA for social media images, which captures visual and textual information into a unifying model. To our best knowledge, USEA is the first unsupervised sentiment analysis framework for social media images; and
- Comparative studies and evaluations using datasets from real-world social media image-sharing sites, documenting the performance of USEA and leading existing methods, serving as benchmark for further exploration.

References

- [Borth *et al.*, 2013] Damian Borth, Rongrong Ji, Tao Chen, Thomas Breuel, and Shih-Fu Chang. Large-scale visual sentiment ontology and detectors using adjective noun pairs. In *Proceedings of the 21st ACM international conference on Multimedia*, pages 223–232. ACM, 2013.
- [Hu and Liu, 2004] Minqing Hu and Bing Liu. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177. ACM, 2004.

- [Hu *et al.*, 2013] Yuheng Hu, Fei Wang, and Subbarao Kambhampati. Listening to the crowd: automated analysis of events via aggregated twitter sentiment. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 2640–2646. AAAI Press, 2013.
- [Jia *et al.*, 2012] Jia Jia, Sen Wu, Xiaohui Wang, Peiyun Hu, Lianhong Cai, and Jie Tang. Can we understand van gogh’s mood?: learning to infer affects from images in social networks. In *Proceedings of the 20th ACM international conference on Multimedia*, pages 857–860. ACM, 2012.
- [Yang *et al.*, 2014] Yang Yang, Jia Jia, Shumei Zhang, Boya Wu, Juanzi Li, and Jie Tang. How do your friends on social media disclose your emotions? 2014.
- [You *et al.*, 2015] Quanzeng You, Jiebo Luo, Hailin Jin, and Jianchao Yang. Robust image sentiment analysis using progressively trained and domain transferred deep networks. 2015.
- [Yuan *et al.*, 2013] Jianbo Yuan, Sean Mcdonough, Quanzeng You, and Jiebo Luo. Sentribute: image sentiment analysis from a mid-level perspective. In *Proceedings of the Second International Workshop on Issues of Sentiment Discovery and Opinion Mining*, page 10. ACM, 2013.