ORIGINAL SCIENTIFIC PAPER
RECEIVED: 27-03-2014
ACCEPTED: 28-04-2014

ACTA GRAPHICA 225

Exploratory Analysis of User-Generated Photos and Indicators that Influence Their Appeal

Authors

Urban Sedlar*

University of Ljubljana,
Faculty of Electrical Engineering,
Slovenia
**E-mail: urban.sedlar@fe.uni-lj.si*

Abstract:

In this paper we analyze whether simple indicators related to photo quality (brightness, sharpness, colour palette) and established content detection techniques (face detection) can predict the success of photos in obtaining more "likes" from other users of photo-sharing social networks. This provides a unique look into the habits of users of such networks. The analysis was performed on 394.000 images downloaded from the social photo-sharing site Instagram, paired with a de-identified dataset of user liking activity, provided by a seller of a social-media mobile app. Two user groups were analyzed: all users in a two month period (N = 122.260) and a highly selective group (N = 3.982) of users that only like <10% of what they view. No correlation was found with any of the indicators using the whole (non-selective) population, likely due to their bias towards earning virtual currency in exchange for 'liking'. However, in the selective group, small positive correlation was found between 'like' ratio and image sharpness (r=0.09, p<0.0001) and small negative correlation between 'like' ratio and the number of faces (r=-0.10, p<0.0001).

Keywords:

User-Generated Content, Photography, Social Networks, Image Processing, Exploratory Analysis, Face Detection

1. Introduction

In recent years, photo-sharing applications have taken off with unforeseen popularity. They have become ubiquitous and due to their convenience and social features (sharing, 'liking', re-posting images, commenting) they act as social networks of their own (*Hochman and Schwartz 2012*). Some evidence of this trend can be summed up by the acquisition of Instagram by Facebook for $1 billion in 2012.

On photo sharing social networks the amount of followers and 'likes' are equivalents of social currency, as is the case with social networks in general (*Kietzmann, J.H. et al, 2011*). Since the social graph on such networks is typically asymmetric (following an account does not mean the account automatically reciprocates and follows the user back), the relationship dynamics and statistics of photo 'liking' are in most cases governed by user popularity and, presumably, the quality of the photos.

Recently a new market has emerged for trading with the virtual currency of *photo likes* and *followers*. Nowadays, many providers exist that specialize in bulk delivery of either. In that light, several mobile application developers have tapped the potential and provided a free marketplace for exchanging *likes* and *followers*.

We have obtained a dataset generously provided by one such developer. This gives us a unique opportunity to examine if simple photo quality metrics and content detection techniques (such as face detection) can predict the success of the photo in obtaining more *likes* from other users of photo sharing social networks.

The research presented here is based on low-level indicators, with the addition of face detection. Similar research using low-level metrics has been conducted to classify professional photo quality based on user ratings (*Ke, Y. et al, 2006*), while also more elaborate approaches exist that take into account the composition of the image itself (*Bhattacharya, S. et al, 2010*).

2. Material and methods

The dataset analyzed in this paper comes from an iOS app specializing in providing a marketplace for users to exchange photo *likes* on the photo sharing social network Instagram. The users go through the following steps to participate in the market.

- User is presented with an unending stream of photos, submitted by others.
- For each of these photos, they can either *like* it (i.e., click the "Like" button to indicate they indeed like it; this action can be seen by other users of the photo sharing site), or skip it.
- If they *like* it (i.e., they click the "Like" button), they earn a unit of virtual currency
- For every two units of currency earned, they can "buy" a 'like' for a photo of their own.

However, in this process, buying does not necessarily imply *obtaining*, since photos that are unappealing or obviously inappropriate might get skipped without obtaining a single 'like'. Each photo is shown to a user only once; whether 'liked' or skipped, it is not shown again.

In this paper we attempt to determine if any of a number of possible metrics influence the general appeal or *likability* of the photo, i.e., a ratio of how many times the photo received a *like* in relation to how many times it had been viewed by users.

For the analysis, de-identified logs of user activity were provided by the seller of the mobile app, spanning 2 months from 20.12.2014 to 20.02.2014. Logs included URLs of 498.950 photos hosted by the social photo-sharing site Instagram.

We attempted to download all referenced photos. Of those, roughly 21% (104.607) were not available any more during the crawling period (01.03 to 10.03.2014), either due to the removal by the users themselves, or due to flagging by other users for reasons of inappropriateness (nudity, etc.) or copyright infringement. This shows significant link rot, which is a widely recognized problem of web content. Thus, only the remaining 79% or 394.341 photos were successfully fetched in their maximum available resolution (either 640x640, 612x612 or 480x480 pixels). Pictures contained no EXIF metadata, so further analysis of EXIF parameters was not possible. The total size of the dataset was 32GB.

So as to establish if photo quality and various other content-related metrics determine the *likability*, the following indicators were calculated:

- Blurriness
- Number of faces present in the image
- Total area occupied by faces
- Number of colours
- Colour or greyscale flag
- Minimum, maximum and average brightness (value)
- File size
- Image entropy

In addition, for each photo the following numbers were tallied from the logs:

- Total number of times the image was served to users (N_S)
- Number of times the image was actually liked by a user (N_L)
- Number of times the image was skipped by a user, calculated as $N_S - N_L$
- From N_L and N_S the "Like ratio" was calculated as $LR = N_L/N_S$

Furthermore, since the population sample was biased towards earning virtual currency regardless of the photo quality, an additional count was performed using only *highly selective* users. For that purpose, only actions involving users that 'liked' no more than 10% of all viewed images were counted. This has yielded two new dimensions, N_{S10}, N_{L10} and a metric "Like ratio of selective users (10%)", $LR_{10} = N_{L10}/N_{S10}$.

Below, the methods used for calculating and estimating the content-related indicators are described in detail.

2.1 FACE DETECTION

Face detection was performed on all successfully downloaded images using OpenCV (*OpenCV, 2014*); the included Haar cascade classifier was used (i.e., "haarcascade_frontalface_alt.xml"), with minimum face size set to 20x20 pixels. Since photo-sharing sites are a creative outlet for many, not all faces were expected to be aligned straight. Due to the fact that Haar transforms are not rotationally invariant (except at very small rotation angles), face detection was attempted multiple times: each photo was rotated from -90° to +90° in 5° increments, with face detection performed and detection results saved at each of 37 (18+1+18) possible angles where there was a match. Additionally, maximal total area occupied by faces was stored, allowing us to distinguish close-up portraits, such as "selfies", from other types of shots. In the results section, the maximum number of faces found at any single angle between -90° and +90° is used. The reason for this is the fact that overlapping detections are possible due to the small rotation increments chosen; thus, summing up all detections would in many cases count the same face more than once. Rotationally invariant face detection methods (Kamruzzaman et al, 2010; Hladnik, 2013) were not chosen due to their higher computational expensiveness and the size of the dataset.

2.2 BLURRINESS ESTIMATION

One of the obvious photo acceptance factors is the question whether the subject of the photo is in focus. This can be hard to determine, however, it is possible to quantify the amount of sharp edges in a photo, which represents an inverse metric. For this a Laplacian filter was used to calculate a spatial second derivative of the image by convolving it with a 3x3 kernel (eq. 1), thus highlighting areas of rapid intensity change. Finally, the obtained intensity values were summed up over the entire intensity map to arrive at a single figure quantifying image sharpness (the lower the value, the blurrier the picture).

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix} \qquad (1)$$

Further, the sharpness value was divided by the number of pixels to account for different image sizes and normalized to [0, 1].

2.3 COLOUR COUNT AND COLOUR MODE

For every picture, a number of unique colours was obtained and stored without using quantization. The absence of quantization means that the colour count was higher than perceptible; however, the number should implicitly contain some information about possible use of colour range-narrowing image filters. To determine if a black and white image was submitted, the image was converted to greyscale and the colours counted again. If the number of colours was equal in both cases, the image was marked as greyscale.

2.4 FILE SIZE AND IMAGE ENTROPY

File size was measured in bytes on disk, after download, for the purpose of determining a possible correlation of likeability with download time. In addition, file size is directly proportional to image entropy (eq. 2); both measures indicate the visual complexity of the image, which is a relevant metric to consider. Image entropy was calculated by (i) converting the image to greyscale, (ii) calculating the histogram of the frequency of each of 256 greyscale components, (iii) normalizing the histogram components and (iv) summing up all non-zero components p_i, multiplied by $-\log_2(p_i)$.

$$E = \sum_{i} p_i \log_2(p_i) \qquad (2)$$

2.5 MINIMUM, MAXIMUM AND AVERAGE BRIGHTNESS, HISTOGRAM WIDTH

Brightness (i.e., value) was determined in the greyscale mode; minimum, maximum and arithmetic mean pixel values were stored for each image. Histogram width was calculated as *HW = max brightness - min brightness.*

2.6 A NOTE ON THE TOOLS USED

The de-identified logs were obtained as plaintext files; they were parsed using the Python programming language and the data was inserted into a MySQL database. A simple Python script was used to loop through all records and attempt to download each image file using command-line *wget* utility. Face detection and Laplacian filtering were performed using the OpenCV-Python binding library; all other operations were performed using Python Image Library (PIL, 2014).

3. Results

The following table gives a summary of the dataset and some of the metrics.

Table 1. Dataset summary

Dataset information	
Dataset size (unique picture URLs)	498.950
Downloaded pictures	394.341
of those: 640x640 px	338.136 (85.7%)
of those: 612x612 px	56.052 (14.2%)
other sizes:	153
Unavailable pictures	104.609 (21%)
Total size of downloaded data	32 GB
Face detection	
Pictures without faces detected	223.340 (56.6%)
1 face detected	123.489 (31.3%)
2 or more faces	47.512 (12.1 %%)
Colour statistics	
Colour images	384.607
Greyscale images	9.734

3.1 POPULATION SAMPLE

The population sample during the analyzed period comprised a total of 122.260 unique users; due to the nature of the app, the full sample was biased, since it consisted of people eager to amass as much of the virtual currency as possible. To compensate for the incentive of earning the virtual currency regardless of the photo content, we performed the analysis using two sets of users: (A) full set, encompassing all 122.260 users, and (B) only users that exhibited high selectivity, with ratio of viewed pictures vs. liked pictures of 10:1 (a total of 3.982 users). The histogram in Figure 1 shows the distribution of user selectivity; the first bin of [0,0.1] was used as the group B.

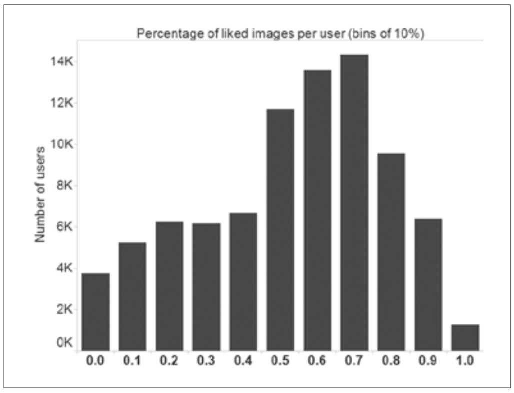

Figure 1. Histogram of user selectivity; the last bin shows users that like >90% of what they view, while the first bin, which represents the group B, shows users that like <10% of what they view.

3.2 FILE SIZE

Larger images tend to take more time to download, especially when using mobile data connectivity; in addition, they typically contain more information, since file size is directly correlated with entropy (in our case, $r=0.52$, $p<0.0001$). However, no correlation was found between file size and like ratio using the entire population A ($r=0.004$, $p<0.01$), or with the selective group B ($r=0.05$, $p<0.0001$).

3.3 NUMBER OF COLOURS; BLACK AND WHITE VS. COLOUR IMAGES

Number of images divided into 10k colour count bins can be seen in Figure 2. Of 394.341 photos, 9734 (or 2.5%) were greyscale (256 or less grey levels), as detected by using the method described above; these fall into the first bin between 0 and 10k colours. Non-greyscale monochromatic (toned) images were not detected by the method used, but could explain the high count in first bin.

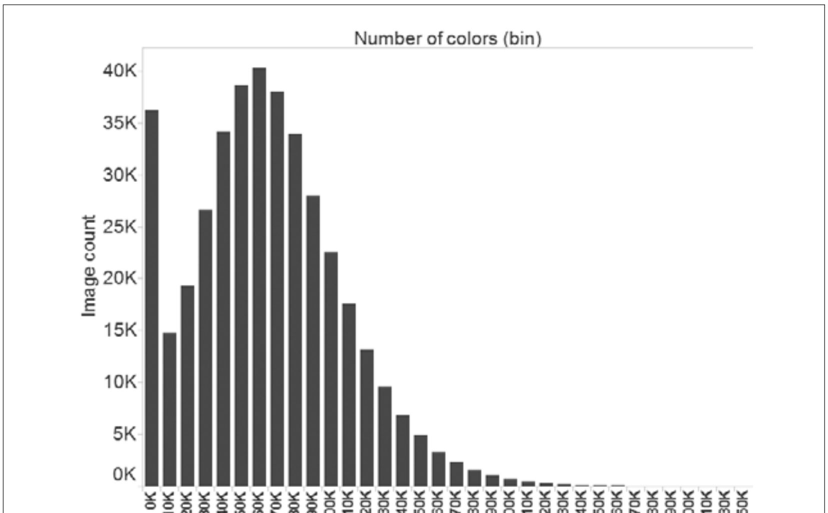

Figure 2. Distribution of colors; the large bar in bin [0-10k] is in part due to large amount of grayscale images.

3.4 SHARPNESS

No correlation was found between sharpness and 'like' ratio when using the entire population (group A) ($r=0.04$, $p<0.0001$). Small positive correlation was found in the highly selective group B ($r=0.09$, $p<0.0001$), which is also evident from the histogram in Figure 3.

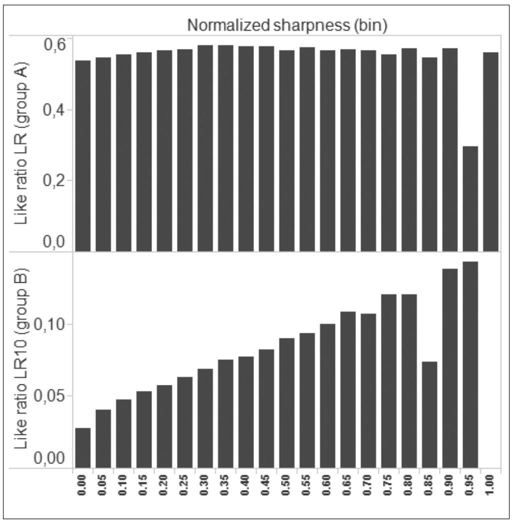

Figure 3. Average like ratio at different sharpness levels

3.5 BRIGHTNESS AND HISTOGRAM WIDTH

Small correlation was found between brightness (value) and 'like' ratio in group A ($r=0.08$, $p<0.0001$) and no correlation in group B ($r=0.05$, $p<0.0001$). No correlation was found with histogram width in either group (group A: $r=0.01$, $p<0.0001$; group B: $r=0.03$, $p<0.0001$).

3.6 PRESENCE OF FACES

171.001 of 394.341 photos (43.4%) contained faces. The majority of those only showed 1 face (31.3% of the entire set). The distribution of number of faces between 1 and 20 can be seen in Figure 4a; the maximum total number of faces found in a single image was 42. The histogram in Figure 4b shows the distribution of total area occupied by faces. Photos without faces were excluded from the histogram.

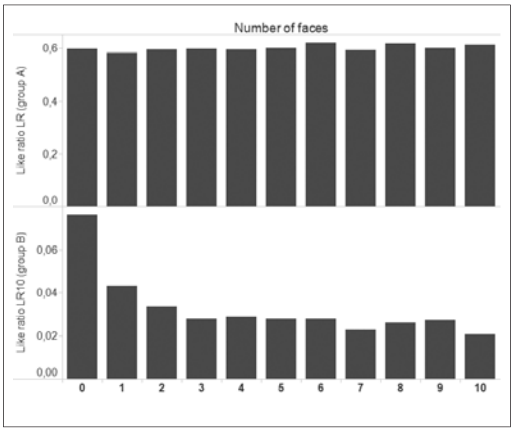

Figure 5. Average like ratio of images with different numbers of faces

As with the other metrics, group A exhibits no correlation with the number of faces detected; group B, however, shows small negative correlation ($r=-0.1$, $p<0.0001$), preferring the photos without faces more; this can be also seen in the histogram in Figure 5.

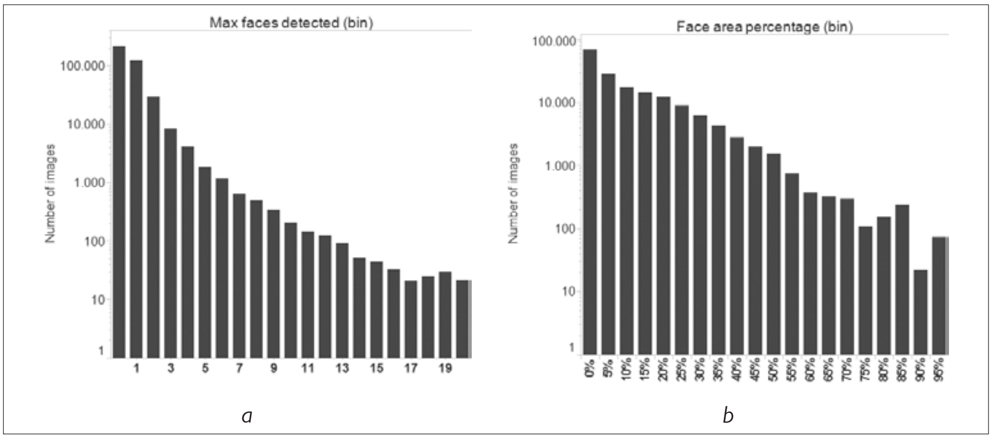

Figure 4. (a) distribution of number of faces (1-20), (b) distribution of cumulative face area

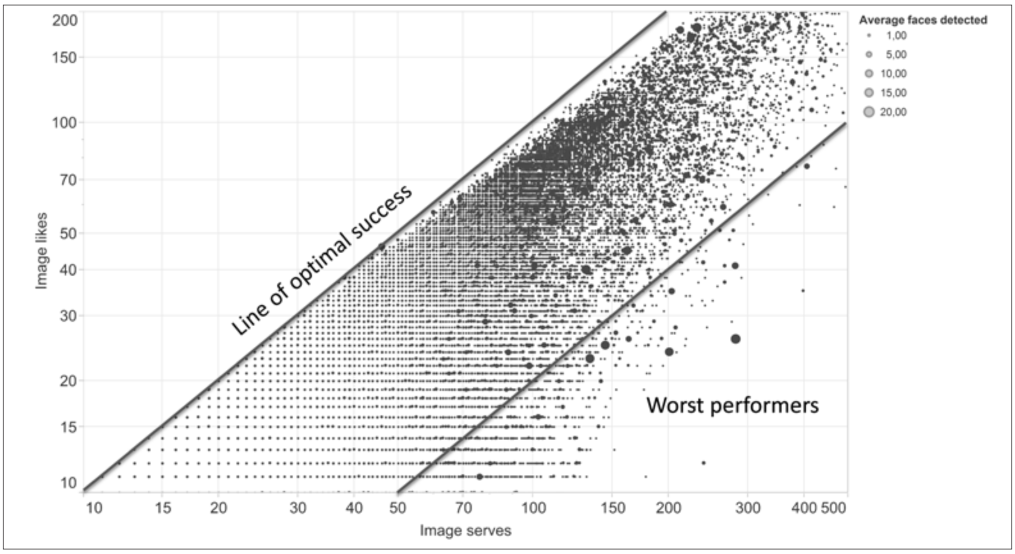

Figure 6: images displayed on a scatter plot (log-log) based on the number of serves (i.e., views), number of likes and average number of faces detected. Line of optimal success is shown, as is the line of 20% success ratio.

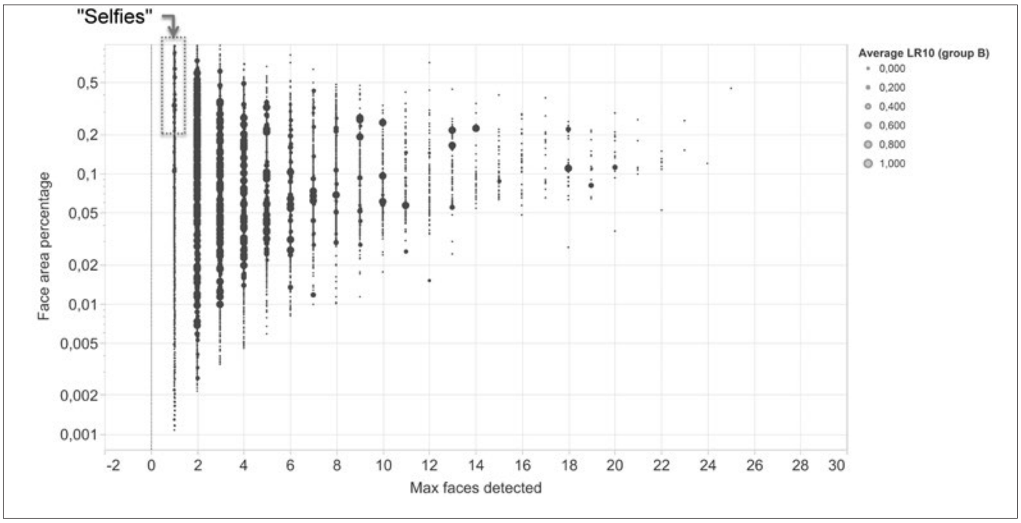

Figure 7: Average 'like' ratio of selective users (LR10), shown as dot size. It can be seen that in scatter plot of face area percentage (log) and number of faces, "selfies" (here defined as pictures with 1 face between 20% and 100% of image size) get worse performance than images with 2 or more faces.

4. Discussion

A surprisingly large amount of pictures show a single face; of those, 31.300 occupy >20% of the image, which would match the size obtained when users take the photos themselves. Thus, we classified all such photos as "selfies" (defined as "a type of self-portrait that is typically taken with a hand-held digital camera or camera phone"). This would indicate high incidence of "digital narcissism" on Instagram. The hypothesis that the presence of faces in the picture determines the likeliness of other users 'liking' a photo was proven correct, however not in a way that was expected. After filtering out the biased part of the population, the remaining selective user group exhibited negative correlation, which means they were *less* likely to like photos with faces present. This could be in part explained by the fact that users are reluctant to show they like another person based solely on their looks, as it might reveal too much of a personal taste, or that they simply *don't* like the looks of that person.

Similarly, there was small positive correlation between the sharpness of an image and the 'like' ratio, which confirms that an image with the subject in focus (presence of sharp edges) should generally have higher appeal than an image with no sharp edges. Surprisingly, no correlation was found with histogram width in either group. These findings reflect, more than anything, a bleak trend that mobile photography is headed for a state where quantity is more important than quality, and where liking something has become a mechanical activity without much regard for the contents.

5. Conclusions

In this paper we presented a unique insight into modern photo-sharing social networks through an analysis of a large dataset from the perspective of Instagram users. To the best of our knowledge, this is the first analysis exploring the relation between the behaviour of social networks users and image contents on a dataset of this scale. One of the key drivers for this study was the expectation that certain easily determinable factors influence the outcome of photo-liking activities, and that the size of the dataset would accentuate such features. However, the sample population as a whole might have been too specific to allow us to find such factors with the methods chosen here. To remedy that, a subset of the population was chosen, but it is unknown how biased they are compared to the general population. We showed that the population of users as a whole exhibited no noticeable preferences for any of the calculated indicators, while the selective user group showed small preference for sharper images and small negative preference for images with presence of faces.

The study represents an important step towards a better understanding of modern social networks and the habits of their users. In the scope of the future work, it would be interesting to see a similar analysis performed on a general, non-biased population, a population biased in a different way, or taking into account other meta-data available from social networks (e.g. Instagram API). The latter, however, would be much easier if the dataset were not de-identified, and this would likely present a problem from the standpoint of privacy.

References

BHATTACHARYA, S., SUKTHANKAR, R., & SHAH, M. (2010). A framework for photo-quality assessment and enhancement based on visual aesthetics. In *Proceedings of the international conference on Multimedia* (pp. 271-280). ACM.

HLADNIK, A. (2013) Image Compression and Face Recognition: Two Image Processing Applications of Principal Component Analysis.

HOCHMAN, N., & SCHWARTZ, R. (2012). Visualizing Instagram: Tracing cultural visual rhythms. In *Proceedings of the Workshop on Social Media Visualization (SocMedVis) in conjunction with the Sixth International AAAI Conference on Weblogs and Social Media (ICWSM–12)* (pp. 6-9).

KAMRUZZAMAN, S. M., SIDDIQI, F. A., ISLAM, M., HAQUE, M., & ALAM, M. S. (2010). Rotation Invariant Face Detection Using Wavelet, PCA and Radial Basis Function Networks. *arXiv preprint arXiv:1009.4974*.

KE, Y., TANG, X., & JING, F. (2006). The design of high-level features for photo quality assessment. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on* (Vol. 1, pp. 419-426). IEEE.

KIETZMANN, J.H., HERMKENS, K. MCCARTHY, P. AND SILVESTRE, (2011) B. S. Social media? Get serious! Understanding the functional building blocks of social media, Business Horizons, Volume 54, Issue 3, May–June 2011, Pages 241-251, ISSN 0007-6813.

OPENCV OPEN SOURCE COMPUTER VISION FRAMEWORK (2014), Available from: http://opencv.org/; [10 March 2014].

PIL PYTHON IMAGING LIBRARY (2014), Available from: http://www.pythonware.com/products/pil/; [10 March 2014].