Personalized Economy of Images in Social Forums: An Analysis on Supply, Consumption, and Saliency

Sangmin Oh and Megha Pandey and Ilseo Kim and Anthony Hoogs and Jeff Baumes
Kitware Inc.
E-mail: sangmin.oh, megha.pandey, ilseo.kim, anthony.hoogs, jeff.baumes@kitware.com

***Abstract*—In this work, we focus on the novel problem of analyzing individual user's behavioral patterns regarding images shared on social forums. In particular, we view diverse user activities on social multimedia services as an economy, where the first activity mode of sharing or posting is interpreted as supply, and another mode of activity such as commenting on images is interpreted as consumption. To characterize user profiles in these two behavioral modes, we propose an approach to characterize users' supply and consumption profiles based on the image content types with which they engage. We then present various statistical analyses, which confirm that there is an unexpected significant difference between these two behavioral modes. In addition, we introduce a statistical approach to identify users with salient profiles, which can be useful for social multimedia services for blocking users with undesirable behavior or viral content promotion. We showcase the benefits of the proposed saliency detection approach and its extension to detect significant key images from a complex dataset, which exhibits the inherent multi-modal nature of user bases of social multimedia services.**

I. Introduction

Users of social multimedia services (e.g., Twitter, Flickr, and Reddit) share diverse images and videos, which are consumed by other users. Such activities are facilitated via social structure (e.g., friends, communities, and etc), as well as tools such as subscription-based data feeds. The importance of social multimedia services in our society has been increasing exponentially in recent years as a medium for culture, personal communication, news sharing, and commerce. Accordingly, with shared goals of making these services more effective, researchers have been making rapid progress in understanding user behavior, content propagation, and characteristics of viral content [1], [2], [3], [4], [5], [6], [7], [8].

In this work, we focus on the novel problem of analyzing *individual user's behavioral patterns regarding visual multimedia (i.e., images)* shared on social forums. The wide adoption of camera-equipped personal devices has facilitated the sharing of images on social media services in recent years. Even on Twitter, which is designed for microblogs with short texts, a large amount of tweets carry embedded images. Hence, the problem of understanding the types of visual content that people post, and with which they engage, is becoming more important than ever. To the best of our knowledge, although existing work [1], [2], [5], [7], [4] has studied various aspects of images and videos on social media services, analysis of diverse aspects of individual user behavior regarding engagement with visual content has been lacking, and our work provides novel insights in this direction.

In particular, we view diverse user activities on social multimedia services as an economy, where the first activity mode of sharing or posting is interpreted as '*supply*'. On the other hand, another mode of activity such as commenting on images is interpreted as '*consumption*'. Understanding these two modes is important in improving our understanding of content propagation, as well as providing improved personalized services to users. For this purpose, we have conducted a data-driven analysis of two behavioral modes using the Reddit image dataset [5], which contains images shared on *reddit.com*. Reddit is a popular social forum where people upload posts, with or without images. Then, other users can comment on (and even vote up or down) each post. The dataset contains a total of approximately $17K$ unique images, collected during a timespan of 4.5 years. Additionally, for every image, it contains a metadata log that specifies who uploaded or commented on the image, as well as timestamp information, which makes it ideal for this research.

A novel contribution of this work is that we propose an approach to characterize users' supply and consumption profiles and present a statistical analysis of the differences between each user's two activity modes. In our approach, two profiles of a user are characterized as a distribution of image categories with which the user interacts. One challenge of analyzing images in social multimedia services is that they are extremely diverse. For example, Figure 1(a) shows randomly selected images from the Reddit image dataset, which exhibits significant variation. This work presents that even for an in-the-wild dataset such as the Reddit image dataset, meaningful image groups can be obtained automatically by clustering images based on their visual features. For example, Fig. 1(b) shows sample image groups, such as people, images with texts, cartoons, and screen shots, obtained by the proposed approach. From these sample image groups, user profiles are obtained as bag-of-words (BoW) representations of the categories, effectively capturing the image styles users prefer to supply and consume.

Finally, by comparing users' supply and consumption profiles, our analysis finds that as much as ˜40% of users show fairly different supply and consumption patterns regarding images. This finding contradicts a baseline hypothesis that users are likely to post and comment on similar image categories, and suggests that personalization of social multimedia services need to be optimized from both of the angles, rather than pursuing a single monolithic model per user. Most existing research focuses solely on consumption patterns [3], [4], [5], [6], [7], and we believe our study is the first one that concretely shows that there can be significant differences between the two behavioral modes.

As another contribution, we present *a statistical approach to identify users with salient profiles*, despite the inherent multi-modal nature of user bases on social multimedia services. By definition, a salient user exhibits a unique profile. Identifying such a user can be beneficial for social media services for various purposes. For example, such a user may be a trend setter who is in the middle of creating unseen but potentially viral content that social media services may want to consider promoting heavily to draw new users into their services. In other cases, such a user may be posting illegal or prohibited content that is outside the norm of the common user base. One difficulty in identifying salient users is that the users of social media services are very diverse. There are multiple communities with different interests, which naturally creates a multi-modal distribution of user profiles. We propose a detection method to identify salient users, which is based on non-parametric kernel density estimation framework [9]. Finally, we show that certain users' profiles can change dynamically over time and that the proposed method can be extended to be used in such scenarios, providing the ability to identify key images which contribute the most towards users' saliency. We showcase the usefulness of the proposed methods through qualitative examples.

The remainder of the paper is organized as follows. In Section II, related work is reviewed. Section III describes the methodologies utilized to characterize individual user's supply and consumption profiles. It also presents our analysis on the similarities and differences between these two behavioral modes. Finally, Section IV introduces our proposed approaches for salient user detection and key images based on the images supplied by users.

II. Related Work

A group of researchers studied the problem of modeling and predicting the popularity of tweets on Twitter. Most of their work (e.g., [10], [3], [6]) focused on using only textual content within tweets and disregarded embedded images and videos. In [10], [3], the authors studied the problem of developing a general population-level estimator, i.e., predicting the general total lifetime retweet counts. [6] is more similar to ours in motivation, in that they studied the problem of predicting retweet likelihood on an individual basis, although they focused on using text inputs only. To the best of our knowledge, [4] is the only work focused on predicting the popularity of tweets using embedded visual multimedia in which the authors evaluated multiple feature types and their combinations to train a general population-level total retweet count predictor. In our work, we focus more on individual patterns and differences between the two behavioral modes.

Additionally, there has been substantial work on analyzing the propagation of content on social media services, where notable work includes studies on Flickr [1], YouTube [2], and Reddit [5]. In [1], [2], the authors focused on analyzing the patterns of popularity formation and propagation as a function of multiple factors, such as the amount of time passed since the content was posted, as well as structural properties such as the degree and density of connectivity around the originating account holder. In [5], the Reddit image dataset was collected, which is used in this work. In particular, the authors introduced a model that predicts the popularity of a post as a function of multiple detailed factors, such as the time when it was posted, its similarity to previous viral posts, and its title. Nonetheless, [5] did not study the impact of the visual content itself, which is the focus of this work.

In the computer vision community, recent work [11] studied approaches to improve image search via a personalized visual attribute classifier. Their experimental results showed that personalization can improve image search aided by adapted visual attributes. While [11] focused on constrained image domains such as shoes or scenes, our work focuses on analyzing user behavior from an in-the-wild social multimedia dataset from social forums regarding two different aspects of supply and consumption, as well as the detection of salient users. In [12], the authors proposed an approach to improve image classification aided by metadata available from images uploaded onto a social network. It focused on image classification into semantic labels, but did not attempt to analyze personalized patterns.

Recently, [8] showed that vibrant online communities exhibit constantly changing linguistic patterns and user engagement. Their work focused on text content and patterns of temporal changes. Furthermore, [7] demonstrated that each community shows evolving viral topics over time and that such topical grouping can be conducted automatically based on various metadata in addition to visual features. Our work is different in that we focus on individual user patterns with regard to visual data and study the saliency and anomaly of content, with or without dynamic conditions.

Existing work on saliency detection on social multimedia services mostly formulates the problem as detecting nodes that exhibit abnormal structural patterns such as connectivity [13], [14]. Rather than focusing on structural properties, our saliency detection approach focuses on detecting salient users based on the actual visual content they supply and consume.

III. Modeling and Analysis of Personalized Image Economy

In this section, we introduce our framework to characterize users' supply and consumption profiles. Then, we present a statistical analysis of the differences between users' two activity modes in the Reddit image dataset.

An overview of our approach is illustrated in Fig. 1. First, the set of images in the Reddit image dataset, some samples from which are shown in Fig. 1(a), are automatically grouped into different image categories based on their visual features. Fig. 1(b) shows a set of sample categories learned during the grouping process described in Sec. III-A. It can be observed that while the nature of images shared on social forums tend to be extremely diverse, they can still be grouped into fairly coherent categories based on their content and styles. For example, the different groups shown in Fig. 1(b), while noisy, show consistent visual content such as cartoons, faces with text, scenery, animals, cartoons with text, and group of people.

Each user's supply and consumption profiles are computed as BoW representations of the learned categories, effectively capturing the image content and styles the user prefers to

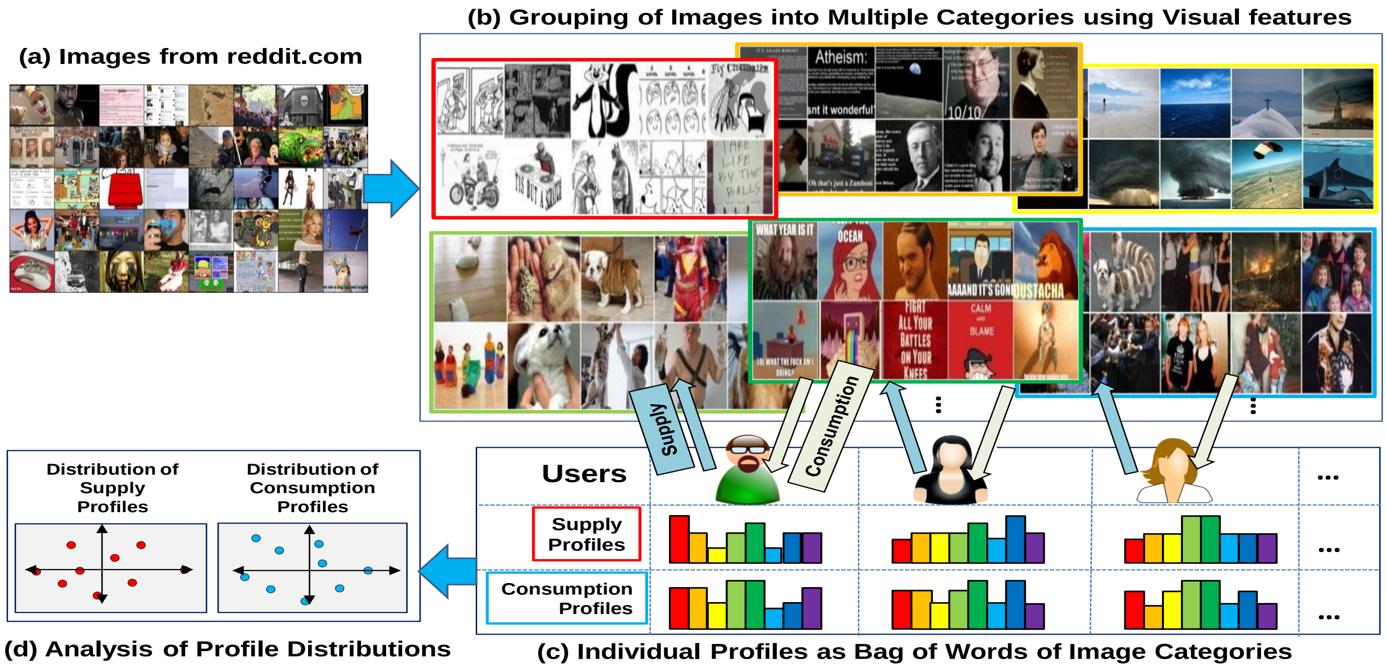

Fig. 1. Illustration of overall approach. (a) Randomly sampled images from Reddit image dataset. (b) Images are grouped into multiple categories based on their content and styles. Different groups are marked in different colors. (c) Individual user's supply and consumption profiles are computed as bag-of-words representation based on the history of posting and commenting on different image categories. (d) Profiles of multiple users are analyzed to measure similarity and difference between users' two behavioral modes. Every dot represents a user's profile. Users with salient profiles are identified in this space as well.

supply and consume. Fig. 1(c) shows nominal profiles of some users as normalized histograms, where the heights of different color bars represent the frequency of the user's supply or consumption of images from the corresponding image category. For example, the middle user in Fig. 1(c) supplies and consumes images from the cluster marked in dark blue color relatively more often than other users. Based on this process, each user's profiles can be represented as two vectors, for supply and consumption respectively, which can then be compared with each other and those of other users as well.

Finally, this section provides a novel statistical analysis on the similarity and difference between users' two behavioral modes. Fig. 1(d) shows two scatter plots, one for supply and another for consumption, where every node represents a user's profile vector. Note that while these nominal illustrations show user profiles in 2D space, the actual analysis is conducted in C-dimensional space in this work, where C is the total number of image categories.

In the following subsections, after more details about automatic image categorization and user profile computation are described in Sec III-A and Sec III-B, Sec. III-C presents our methodologies and statistical analysis results about the characteristics of user profiles. Sec. IV further presents our approaches to identify users with salient profiles in these profile spaces.

A. Learning Grouping of Images by Content and Styles

The groupings of images used in this work are automatically learned using an unsupervised clustering technique. In detail, visual features are computed from images and every image is represented as a feature vector. Then, images are clustered into C groups using standard K-means clustering. Here, C is a pre-specified parameter where we tried multiple values for C to capture diverse groups of images at different granularities, including: $50, 100, 250, 500$ While most of our analysis presented in the following sections is relatively insensitive to the value of C, different parameters do show slight variations, and we will show such variations as feasible to provide more insight. In the rest of the paper, each cluster center will be denoted as c_i and the whole set as $\{c_i | 1 \leq i \leq C\}$.

Among the diverse visual features we explored, we found that visual features which capture both color and shape provide the most satisfying grouping results in terms of coherency of content and styles within each group. For detailed implementation, we used color-SIFT from [15]. In particular, BoW representation of sparsely detected (from corner point detections) feature points, introduced in [16], provided the highest grouping quality, probably because it can capture more salient foreground from images that contain high ratio of uniform background (e.g., comics) and/or texts – see top left, top middle, and bottom middle groups in Fig. 1(b) for examples. A set of locally computed raw colorSIFT features from every image were quantized into a BoW representation based on a codebook with a size of D, resulting in a D-dimensional vector per image. In this work, $D = 1024$ was used and image feature vectors were $L1$-normalized to sum to one. In the rest of the paper, the image vector will be denoted as x_j for the j-th image.

B. Characterization of User Profiles: Supply & Consumption

In this work, we propose a novel approach to characterize each user's supply and consumption profiles as distributions of image categories which the user interacts with. Based on the C number of learned image categories, every profile is computed as an $L1$-normalized C-dimensional vector, denoted as y_k for

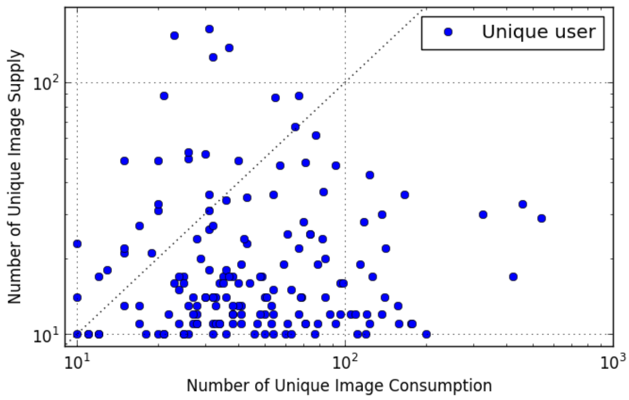

Fig. 2. Activities across users. Each point is a unique user. The horizontal and vertical coordinates correspond to the number of 'unique' images consumed or supplied by each user respectively. Log-scale is used for both axes.

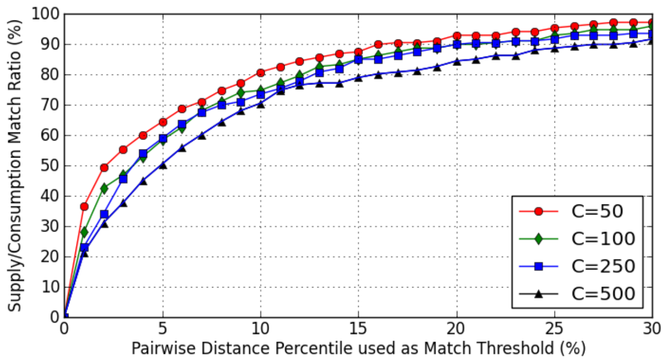

Fig. 3. Ratio of users with similar supply/consumption profiles as a function of varying threshold levels. The match threshold is varying as a function of percentile values within a distribution of all pairwise distances across users. At a relatively large threshold at the 5-th percentile level, as much as ˜40% of users show supply and consumption behavior which are fairly different.

the k-th user. We omit an indicator for supply or consumption in this notation for simplicity.

In further detail, each profile y_k is a BoW representation which captures the relative frequency with which the corresponding user accesses particular image clusters. To compute these profiles, the image metadata consisting of uploader and commenter logs is used. In the rest of the paper, the set of image indices accessed by the k-th user are denoted as a set $\{j|k\}$. When metadata indicates that a user engaged with the j-th image, either via supply or consumption, its feature vector x_j is compared with the C cluster centers. A C-dimensional soft quantization vector x'_j is then computed. Soft quantization is important to avoid unnecessarily harsh boundary effect and produce a more generalizable characterization which is less sensitive to the size C of the clusters used. Essentially, soft quantization represents each image as a weighted mixture of C image categories, where higher weights indicate more similarity and $|x'_j|_1 = 1$. For the experiments reported in this work, we used the soft quantization scheme proposed in [17]. Finally, each user's profile is computed by averaging the soft quantization vectors over all the images they engaged with, as follows: $y_k = \sum_{\{j|k\}} x'_j / |\{j|k\}|$ where $|y_k|_1 = 1$.

C. What does Personalized Image Economy Look Like?

First, we analyze how the users' activities are distributed between consumption and supply, which is visualized in Fig. 2. Here, each point corresponds to a unique user, and the horizontal and vertical coordinates correspond to the number of 'unique' images consumed or supplied by each user, respectively. To improve visibility of details in densely populated areas, log-scale is used for both the axes. By 'unique' images, we mean that when a user has commented multiple times across repeated postings of an identical image, it will be counted as only one consumption. Similarly, even if a particular user has posted the same image multiple times, it will count as only one supply. In particular, Fig. 2 includes only users who consumed and supplied at least 10 unique images each in both the modes. There is a total of 165 such users in the dataset, effectively filtering out a large number of users with very low amount of activity in either mode.

An interesting pattern observed from Fig. 2 is the wide variation in the ratio between supply and consumption across users, and the unexpectedly high ratio of users above diagonal who act more as suppliers than as consumers. Basically, the number of points under 1:1 diagonal indicates users with heavier 'consumption', and vice versa. While it is generally expected that most users will act as consumers rather than suppliers, actually a substantial number of users, as much as ˜15%, show up above the diagonal.

Furthermore, we have analyzed the similarity and difference between supply and consumption profiles for individual users, and found that there is a significant number of users whose profiles do not match – i.e., the images many people supply are fairly different from what they consume. We believe this is the first scientific analysis which reports such a finding.

For our analysis on profile matching, we have used supply and consumption profiles of the 165 users who supplied and consumed at least 10 images, which is the same set of users that appear in Fig. 2. In the rest of the paper, the k-th user's supply and consumption profiles are denoted as $y_{k,s}$ and $y_{k,c}$ respectively. Our methodology to determine whether two profiles match is to check whether the two profiles are within a certain match threshold distance t_p, i.e., $|y_{k,s} - y_{k,c}|_2 < t_p$. Then, the fraction of users whose profiles match can be computed by counting users that satisfy such condition, divided by the total number of users. Fig. 3 shows our analysis where the horizontal axis corresponds to the radius, growing from smaller to larger values, and the vertical axis is the ratio of users with matched profiles. In particular, the match threshold t_p is varying as a function of the p-th percentile within the distribution of all pairwise distances between users. In detail, the p-th percentile match threshold is defined as follows:

$$\begin{aligned} A &= \text{sorted array}\, \{\forall_{k \neq k'} d_{k,k'} | d_{k,k'} = |y_k - y_{k'}|_2\} \\ t_p &= A[|A| * (p/100)] \end{aligned}$$

where A is a sorted array (in increasing order) of all the pairwise distances between users' supply profiles collected across users, and $|A|$ indicates the size of the array.

Then, t_p is obtained as the p-th percentile value from the array A. For example, t_5 is a threshold that corresponds to the top 5 percentile within such distribution. An intuitive interpretation for t_p is that, when it is used as a threshold, for any reference user, approximately p percent of the rest of

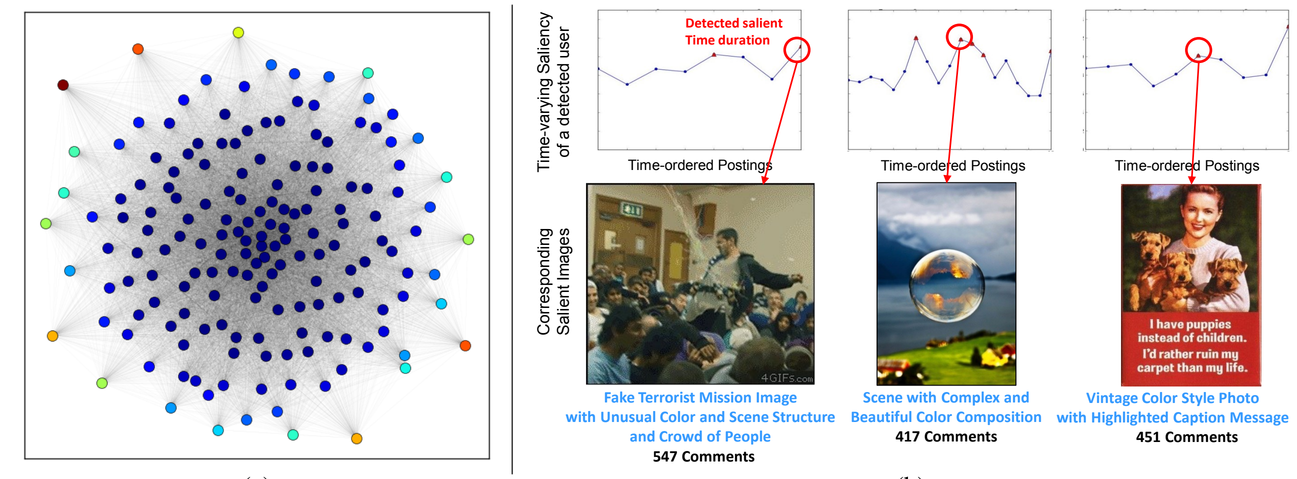

(a) (b)

Fig. 4. (a) Detected salient users (bright colors) with respect to the whole user group. Salient users are detected using kernel density approximation technique from high-dimensional profile space. The detected users appear at boundaries which showcase the usefulness of the proposed approach. (b) Illustration of our approach to identify salient key images. (b,top) Examples of time-varying saliency from three salient users where detected salient durations are marked by red triangles. (b,bottom) Example images sampled from detected salient durations, which include fake terrorist act, supernatural scene, and vintage-style photo with controversial caption. Images show unique visual characteristics, which increases their saliency, and receive large amount of engagement.

the users are within such radius, on average. Accordingly, our methodology assumes that two profiles do not match at the p-th percentile level if the consumption profile of a user is not within the p percent neighborhood of the corresponding supply profile. To provide additional insight regarding the sensitivity of our analysis with respect to the total number of image clusters C, Fig. 3 shows analysis for multiple values of C, where the overall progression patterns are fairly similar while the difference in the match ratio (y-axis) at any particular threshold stays at approximately 20%.

An interesting finding from the analysis shown in Fig. 3 is that, at the 5-th percentile threshold t_5, only ˜60% users show matched profiles (for C=100 or 250), leaving the rest of the 40% showing different profile patterns. It is worth noting that even the threshold t_5 is fairly large in the sense that it will generally contain about 5% of the rest of the users within its radius. This finding contradicts a baseline hypothesis that users are likely to post and comment on similar image categories, and suggests that personalization of social multimedia services needs to be optimized from both angles, rather than pursuing a single model for each user.

IV. Detecting Salient Users and Images

In this section, we introduce our approach to identify salient users along with the key images which substantially contributed to the users' saliency. In particular, our work focuses on identifying salient 'suppliers' and images posted by them, because supply is arguably a more important contributing factor towards the success and failure of social multimedia services, more so than consumption. A novel aspect of our approach is that a user's saliency is measured based on the visual content they supply, i.e., their supply profile, while most existing work [13], [14] address this problem by analyzing abnormal structural patterns in social networks such as properties of friendship connectivity. In addition, we also show that users' saliency may be time-varying – a user who exhibits regular behavior in general may still show salient patterns during certain time intervals, and by detecting such intervals, we can identify key images that contribute significantly to increase the user's saliency.

A. Detecting Salient Users using Kernel Density Estimation

Users of social multimedia forums are diverse in nature, accordingly, the distribution of user profiles tends to include numerous local modes. By definition, a salient user exhibits a unique profile, which tends appear at low-density regions. To detect salient users under such scenarios, we propose a detection method based on non-parametric kernel density estimation techniques [9]. Our approach estimates the likelihood of a user's profile as a weighted sum of its similarity towards all the rest of the users. In detail, for the k-th user profile y_k, the likelihood $p(y_k)$ is defined as:

$$p(y_k) = \frac{1}{|\{i \backslash k\}|} \sum_{\{i \backslash k\}} \alpha \exp\left\{ -\frac{(y_k - y_i)^2}{\sigma^2} \right\},$$

where $\{i \backslash k\}$ indicates the set of all indices except k and α and σ are parameters of Gaussian kernel function. Low-likelihood corresponds to high saliency. Accordingly, we define saliency $s(y_k)$ to be an inverse exponential function of likelihood as follows: $s(y_k) = e^{-p(y_k)}$.

The effectiveness of the proposed approach can be observed from the visualization of detected salient users (bright color) along with the rest of the users (dark color) shown in Fig. 4(a). Each node corresponds to a user and the length of the edges between users correspond to the distance between user profiles. For this visualization, the coordinates of user nodes are computed using the graph layout technique in [18] which incorporates edge distance constraints. It can be observed that users with salient profiles identified by the proposed method appear mostly at the boundary as expected [1] .

[1]Note that the illustration in Fig. 4(a) is not optimized to show the multi-modal nature of full distribution on 2D. A PCA-based 2D projection is an approach that is applicable, however, we found that the diversity of user profiles are not characterized by 2 dimensions effectively.

B. Detecting Key Images from Salient Users

Once salient users are identified, it is important to identify the key images which contributed the most towards saliency. For this purpose, we compute the time-varying saliency of the detected users from time-ordered images based on their posting history. A time-window was used to compute the saliency at every few postings, effectively creating detailed snapshots of temporal user profiles using the identical technique described in Sec. III-B. For example, Fig. 4(b, top) shows the varying saliency of three example salient users (y-axis) as a function of time (x-axis). For these results, we used a time window of size 10, and a stride size of 5 images.

In Fig. 4(b, top), it can be observed that the saliency of users tend to be varying substantially across postings. Subsequently, using thresholds, certain durations are detected to be significantly salient, which are marked as red triangles on the graph (best viewable in zoom mode) in Fig. 4(b).

To qualitatively assess the characteristics of salient images, we sampled images from each flagged duration, which are shown at the bottom of Fig. 4(b). For example, these images show a fake terrorist scene, a super-natural scenery, and a vintage style image with a highlighted caption which contain controversial argument about family value. Overall, these images showcase that the proposed method can identify images with interesting and unusual visual characteristics. We have also observed that such images tend to be consumed heavily by diverse users, as can be observed from the number of comments each image received.

We believe that the approach of detecting salient users and images presented in this section provides a useful tool to identify information which can be used to improve the quality of social multimedia services, such as detecting and promoting viral contents early, or down-weighting postings deemed to be taboo or not fit for the social communities. Automatic approaches to determine such follow-up actions have not been explored yet, and remain to be addressed in our future work.

V. Conclusion

In this work, we have addressed the problem of characterizing individual users' supply and consumption profiles with respect to extremely diverse images on social forums. Our analysis shows that there are a significant number of users who show fairly different behavior in terms of the images they supply and consume, which suggests that personalization of social multimedia services needs to be optimized from both of the angles. Additionally, we have proposed an approach to identify salient users and images they engage with. We propose to identify salient users using a kernel density estimation and identify key images from users' temporal profiles. Our experimental results demonstrate the usefulness of the proposed approaches.

Acknowledgment

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR0011-10-C-0112. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or the U.S. Government. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the DARPA or the U.S. Government.

References

[1] M. Cha, A. Mislove, and K. Gummadi, "A Measurement-driven Analysis of Information Propagation in the Flickr Social Network," in *International World Wide Web Conference (WWW)*, 2009.

[2] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon, "I Tube, You Tube, Everybody Tubes: Analyzing the Worlds Largest User Generated Content Video System," in *Usenix/ACM SIGCOMM Internet Measurement Conference (IMC)*, 2009.

[3] S. Petrovic, M. Osborne, and V. Lavrenko, "RT to Win! Predicting Message Propagation in Twitter," in *International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2011.

[4] E. F. Can, H. Oktay, and R. Manmatha, "Predicting retweet count using visual cues," in *Proceedings of ACM International Conference on information and knowledge management (CIKM)*, 2013.

[5] H. Lakkaraju, J. McAuley, and J. Leskovec, "Whats in a name? Understanding the Interplay between Titles, Content, and Communities in Social Media," in *AAAI International Conference on Weblogs and Social Media (ICWSM)*, 2013.

[6] I. Uysal and W. B. Croft, "User Oriented Tweet Ranking: A Filtering Approach to Microblogs," in *Proceedings of ACM International Conference on information and knowledge management (CIKM)*, 2011.

[7] H. Sundaram, L. Xie, M. De Choudhury, Y.-R. Lin, and A. Natsev, "Multimedia semantics: Interactions between content and community," *Proceedings of the IEEE*, vol. 100, no. 9, pp. 2737–2758, 2012.

[8] C. Danescu-Niculescu-Mizil, R. West, D. Jurafsky, J. Leskovec, and C. Potts, "No Country for Old Members: User Lifecycle and Linguistic Change in Online Communities," in *International World Wide Web Conference (WWW)*, 2013.

[9] V. A. Epanechnikov, "Non-parametric estimation of a multivariate probability density," *Theory of Probability and Its Applications*, vol. 14, no. 1, pp. 153–158, 1969.

[10] L. Hong, O. Dan, and B. D. Davison, "RT to Win! Predicting Message Propagation in Twitter," in *International World Wide Web Conference (WWW)*, 2011.

[11] A. Kovashka and K. Grauman, "Attribute Adaptation for Personalized Image Search," in *International Conference on Computer Vision (ICCV)*, 2013.

[12] J. McAuley and J. Leskovec, "Image Lablineg on a Network: Using Social-Network Metadata for Image Classification," in *European Conference on Computer Vision (ECCV)*, 2012.

[13] N. A. Heard, D. J. Weston, K. Platanioti, and D. J. Hand, "Bayesian anomaly detection methods for social networks," *The Annals of Applied Statistics*, vol. 4, no. 2, pp. 645–662, 2010.

[14] Y. Altshuler, M. Fire, E. Shmueli, Y. Elovici, A. Bruckstein, A. S. Pentland, and D. Lazer, "Detecting anomalous behaviors using structural properties of social networks," in *Social Computing, Behavioral-Cultural Modeling and Prediction*. Springer, 2013, pp. 433–440.

[15] K. E. A. van de Sande, T. Gevers, and C. G. M. Snoek, "Evaluating Color Descriptors for Object and Scene Recognition," vol. 32, no. 9, pp. 1582–1596, 2010.

[16] J. Sivic and A. Zisserman, "Video Google: A text retrieval approach to object matching in videos," in *International Conference on Computer Vision (ICCV)*. IEEE, 2003, pp. 1470–1477.

[17] Y.-G. Jiang, C.-W. Ngo, and J. Yang, "Towards optimal bag-of-features for object categorization and semantic video retrieval," in *Proceedings of the ACM international conference on Image and video retrieval (CIVR)*. ACM, 2007.

[18] T. Dwyer, "Scalable, versatile and simple constrained graph layout," in *Computer Graphics Forum*, vol. 28, no. 3. Wiley Online Library, 2009, pp. 991–998.