Mutual Outcome of Content Excellence And Social Ties on User Administration

A. Vasudeva Rao
Dept of Computer Science and Engineering
Dadi Institute of Engineerng And Technology

Abstract- *User participation in online communities is driven by the intertwinement of the social network structure with the crowd-generated content that flows along its links. These aspects are rarely explored jointly and at scale. By looking at how users generate and access pictures of varying beauty on Flickr, we investigate how the production of quality impacts the dynamics of online social systems. We develop a deep learning computer vision model to score images according to their aesthetic value and we validate its output through crowd sourcing. By applying it to over 15B Flickr photos, we study for the first time how image beauty is distributed over a large-scale social system. Beautiful images are evenly distributed in the network, although only small cores of people get social recognition for them. To study the impact of exposure to quality on user engagement, we set up matching experiments aimed at detecting causality from observational data. Exposure to beauty is double-edged: following people who produce high-quality content increases one's probability of uploading better photos; however, an excessive imbalance between the quality generated by a user and the user's neighbors leads to a decline in engagement. Our analysis has practical implications for improving link recommender systems.*

Keywords- Content quality, Image aesthetics, Network effects, Causal inference, Influence, Matching, Flickr.

I. INTRODUCTION

The user experience in online communities is mainly determined by the social network structure and by the user generated content that members share through their social connections. The relationship between social network dynamics and user experience [1], [2], as well as the influence of quality of content consumed on user engagement [3], [4], [5] have been extensively researched. However, the relationship between network properties and the production of quality content remains largely unexplored. This interplay is key to reach a full understanding of the user experience in online social systems. Learning how people engage with a platform in relation with the content they produce and consume is crucial to prevent churning of existing users, keep them happy, and attract newcomers. The growing availability of interaction data from social media, along with the development of increasingly accurate computational methods to evaluate quality of textual and visual content [6], [7], [8], [9], has recently provided effective means to fill this knowledge gap. We tap into this opportunity and we aim to advance this research direction by providing the first large-scale study on the production and consumption of quality in online social networks.

Key findings from the analysis include the following:

- Unlike popularity, quality is evenly distributed across the network. The resulting mismatch between talent and attention received leaves large portions of the most proficient users with little peer recognition. Users who produce high-quality content but receive little social feedback tend to stay active only for short periods.
- The level of user-generated quality is correlated with individual social connectivity, which causes a majority

illusion effect: users are exposed to images whose average beauty is considerably higher than the average beauty of photos in the platform.

The outcomes of our study have practical implications in the domain of recommender systems. We sketch a simple proof-of-concept of a social link recommender algorithm that maximizes the beauty flow while limiting the beauty imbalance between friends (x7). Simulations show that this simple strategy balances beauty supply and demand, increasing the level of social inclusion in the class of talented yet unpopular users.

II. RELATED WORK

Computational Aesthetics With this work, we build on recent literature exploring the possibility of measuring the intrinsic visual quality of images. Previous related work belongs to the research field of computational aesthetics, a domain in which computer vision is used to estimate image beauty and quality. Traditional aesthetic prediction methods are based on handcrafted features reflecting the compositional characteristics of an image. Datta et al. [10] and Ke et al. [11] were pioneers in this field, with their early work on training classifiers to distinguish amateur from professional photos. Researchers have produced increasingly more accurate aesthetic models by using more sophisticated visual features and attributes [12], [13], looking at the contribution of semantic features [14], [15], and applying topic-specific models [16], [17] and aesthetic-specific learning frameworks [18]. Similar hand-crafted features have successfully been employed to predict higher-level visual properties, such as image affective value [19], image memorability [20], video creativity [21], and video interestingness [22], [23]. Such hand-engineered features are of crucial importance for computer vision frameworks requiring interpretability.

Media Content Quality and User Experience. Similar to our work, several user studies in controlled lab settings have evaluated how quality affects user experience in relation to different types of media content. Gulliver et al. [5] found that video frame rate and network characteristics such as bandwidth and video topic impact user perception of information quality. Bouch et al. explored the importance of contextual and objective factors for media quality of service [3], and Ceaparu et al. found causes of user frustration in web browsing, e-mail, and word processing [4]. In this work we explore the impact of visual aesthetic quality in online social networks..

Networks and Media Diffusion. Bakshy et al. examined the role of social networks in information diffusion with a large-scale field experiment where the exposure to friends' information was randomized among the target population [30]. They found that users who are exposed to friends' social updates are significantly more likely to spread information and do it sooner than those who are not exposed. They further examine the relative role of strong and weak ties in information propagation, showing that weak ties are more likely to be responsible for the propagation of novel information. Social exposure, assortative mixing, and temporal clustering are not the only factors that drive information diffusion and influence. Aral et al. studied the effect of homophily in explaining such evidence [31]. They developed a dynamic matched sample estimation framework to distinguish influence and homophily effects in dynamic networks, and they applied it to a global instant messaging network of 27.4 million users.

III. DATA SET

Flickr is a popular photo-sharing platform on which users can upload a large number of pictures (up to 1 TB), organize them via albums or free-form textual tags, and share them with friends. Users can establish directed social links by following other users to get updates on their activity. Since its release in February 2004, the platform has gathered almost 90 million registered members who upload more than 3.5 million new images daily1.

We collected a sample of the follower network composed of the nearly 40M public Flickr profiles that are opted-in for research studies and by all the 570M+ following links incident to them. For each profile in the sample, we get the complete information about the photos they upload (around 15B in total), the favorites their photos receive from other users, and the groups they are subscribed to. Every piece of information is annotated with timestamps that enable the reconstruction of the full temporal profile of a user's public activities. The whole data spans approximately 12 years starting from the debut of the service in 2004 until March 2016.

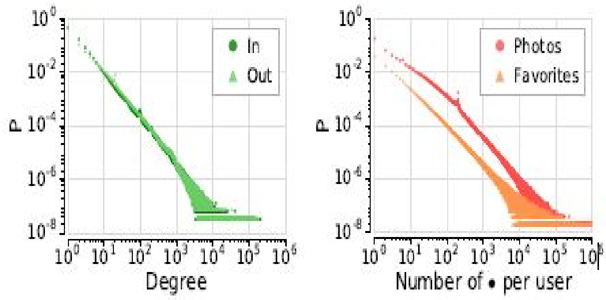

IV. SCORING IMAGE BEAUTY

Training vs Fine-tuning. Deep neural network architectures are essentially layers of artificial neurons that progressively abstract the input data (the image pixels) into an output network response (the predicted category of the input image). In the training phase, network parameters are tuned in order to maximize metrics such as category prediction accuracy. Given the number of parameters involved in such complex architectures, effectively training neural networks is typically a long, expensive process.

Training on Object Detection. We start with a network pretrained for object detection. The architecture and training process for this network are similar to the reference model proposed by Krizhevsky et al. [40]. However, we introduce a few fundamental changes. We doubled the size of the fc6 (second-last) layer from 4096 to 8192. We also used a final fc8-layer consisting of 21841 units (instead of 1000), corresponding to the complete collection of annotated objects in the ILSVRC ImageNet dataset [24].

Fine-Tuning on Aesthetic Scoring. After pre-training on the ImageNet classification task, we fine-tune the network for the aesthetic scoring task. The training set for the aesthetic quality classification task is an internal dataset created using a proprietary social metric of image quality based on Flickr's user interaction data, that has proved to correlate closely with subjective assessments of aesthetic quality. We rank all images from the YFCC100MM dataset [41] according to this metric and then create buckets of "low quality", "median quality", and "high quality" by sampling images from the bottom 10-percentile, the middle 10-percentile, and the top 5-percentile respectively.

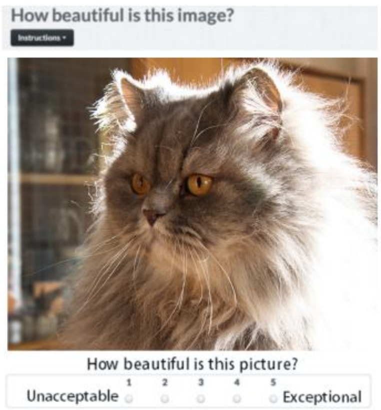

Classification vs. Regression. We tested the possibility to predict a continuous aesthetic score using regression: we obtained a continuous aesthetic score for each sample in our training set by placing the categorical annotations ona continuous scale and normalizing in the range [0,1]; we designed the output layer to contain one single neuron predicting the aesthetic score; we trained to minimize Euclidean loss.

V. CROWD SOURCING BEAUTY ASSESMENT

In addition to the standard performance test on benchmarking datasets, we further evaluate the effectiveness of the aesthetic network with a crowd sourcing experiment. We ask people to evaluate pictures in terms of their beauty, and then compare the human judgments to the aesthetic score predicted by our framework. To design our experiment, we draw inspiration from the image beauty assessment crowd sourcing experiments conducted by Schifanella et al. [29]. Crowd sourcing tasks are complex and can be influenced by unpredictable human factors [42]. Modern crowd sourcing platforms offer control mechanisms to tune the annotation process and enable the best conditions to get high quality judgments. To annotate the beauty of our images, we use CrowdFlower2, a popular crowd sourcing platform that distributes small tasks to online contributors in an assembly line fashion.

Data selection. To help the contributor to assess the image beauty more reliably, we build a photo collection that represents the full popularity spectrum, thus ensuring a diverse range of aesthetic values. To do so, we identify three popularity buckets obtained by logarithmic binning over the range of number of favorites f received. Crowd sourcing task setup. The task consists in looking at a number of images and evaluating their aesthetic quality. At the top of the page we report a short description of the task and we ask to answer the question "How beautiful is this image?" . The contributor is invited to judge the intrinsic beauty of the image and not the appeal of its subject; for example, artistic pictures that capture non conventionally beautiful subjects (e.g., a spider), should be considered beautiful. Quality control. To maximize the quality of human judgments, we apply several controls on the contributors' input. First, we open the task only to Crowd flower contributors with an "excellent" track record on the platform (responsible for the 7% of monthly Crowd Flower judgments). We also limit the task to contributors from specific countries3, to ensure higher cultural homogeneity in the assessment of image beauty [47], [48], [49], [50], [51].

VI. NETWORK EFFECTS

While previous work has studied beauty at the picture level, our large-scale rating of image beauty further enables us to analyze the how beauty is produced over a large social network. In the following, we will characterize the beauty_b (i) of a user i as the average beauty of all of i's public photos. We will refer to this score as user beauty or user quality, for brevity.

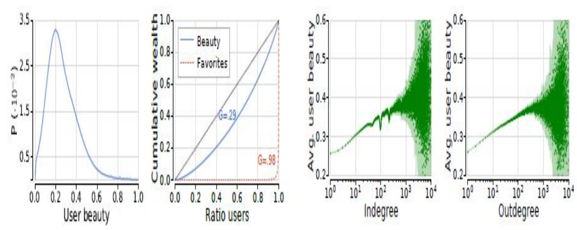

Distribution of quality over the social network

Different activity indicators of social media users tend to be correlated. This has been verified in multiple social mediaplatforms, including Flickr, on a wide range of indicators, especially in relation to nodal degree [55], [56]. We areinterested in verifying whether the level of user quality is correlated to social connectivity or other activity indicators.

Network effects on user retention and quality production

The assortative mixing of quality in the social network could be ascribed mainly to homophily or influence [61]. On onehand, users might preferentially connect to accounts that publish pictures with a similar quality to their own. This would seem natural in a platform like Flickr that hosts a heterogeneous user-base: semi-professional photographers might be interested in following users who are well-versed in the use of photographic techniques, whereas casual users might be following each other mostly for social reasons, unconcerned about aesthetic photo quality. On the other hand, pairs of users might be imbalanced in terms of their quality at the time they connect and close their quality gap later on, over time. For example, amateur photographers could follow professionals and learn new skills from them, thus improving the quality of their pictures.

VII. BEAUTY BASED LINK RECOMMENDER

Classic link recommendation approaches based on the graph structure (e.g., common neighbors and all its variations) tend to suggest popular and very connected users [70], thus increasing the linkage to—and consequently the level of attention on—already well-regarded individuals, keeping potential new talents away from the spotlight. However, since connectivity and user quality are largely orthogonal, algorithms that favor highly-connected users won't necessarily provide adequate visibility to high-quality content.

The clustering results confirm that the talent of a large portion of the user-base —more that 1=4th of the overall population— remains largely untapped, despite its high skill level (as evidenced by the high average beauty value). This group of users is associated with a lower time on platform, measured as the number of weeks with at least one photo upload (Table 3). This gives further support to the intuition that photographers who do not receive adequate recognition for their contributed value tend to churn out sooner. Furthermore, their activity in terms of number of pictures uploaded is limited (the lowest compared to other user classes), thus reducing the flow of incoming highquality content in the platform.

	%users	beauty	fav/photo	connects
Low quality	41.2%	0.17	0.00	0.06
Forlorn beauty	28.1%	0.42	0.01	0.10
Regular user	22.1%	0.25	0.01	0.21
Superstar	8.6%	0.42	0.15	0.35

	Low	Forlorn	Regular	Superstar
photo count	1060	200.4	1869	822.4
time on platform	104.4	84.68	187.0	198.3

Implications

Adopting popularity-driven policies to promote content and users in social networks is a fallacious way of growing healthy online communities [72]. Nevertheless, for several years popularity has been one of the core elements of several online services including search, promoted content, and recommendations. For the first time, we have shown that it is possible to run at scale a reliable profiling of users that captures their contributed quality rather than their popularity. This can have direct practical impact not only in recommender systems, but in any application that need to retrieve, rank, or present images. Furthermore, our study about the notion of quality in combination to the network structure yields important theoretical implications in the domain of social network analysis and, more broadly, network science.

Limitations and future work

Our analysis scratches only the surface of this mostly unexplored research area. Our causal inference analysis groups together similar users to get a balanced matching between control and treatment sets. That is convenient to measure causal effects globally but does not directly allow for a fine-grained analysis of how meaningful user groups (e.g., newcomers vs. professional users) are impacted. The extent to which the exposure to content quality has a different impact on those user categories is an interesting extension of this work. The deep learning algorithm we use is very powerful but lacks explainability: in contrast with classic image aesthetic frameworks based on compositional features, it is not possible to determine why a picture has a given beauty score. Research in explainability in deep learning is still at an early stage, also in the sub-field of image aesthetics.

REFERENCES

[1] M. J. Halvey and M. T. Keane, "Exploring social dynamics in online media sharing," in WWW. ACM, 2007. [Online]. Available: http://doi.acm.org/10.1145/1242572.1242804

[2] A. Susarla, J.-H. Oh, and Y. Tan, "Social networks and the diffusion of user-generated content: Evidence from youtube," Information Systems Research, vol. 23, no. 1, 2012.

[3] A. Bouch, A. Kuchinsky, and N. Bhatti, "Quality is in the eye of the beholder: Meeting users' requirements for internet quality of service," in CHI. ACM, 2000. [Online]. Available: http://doi.acm.org/10.1145/332040.332447

[4] I. Ceaparu, J. Lazar, K. Bessiere, J. Robinson, and B. Shneiderman, "Determining causes and severity of end-user frustration," International journal of human-computer interaction, vol. 17, no. 3, 2004.

[5] S. R. Gulliver and G. Ghinea, "Defining user perception of distributed multimedia quality," ACM TOMM, vol. 2, no. 4, 2006.

[6] X. Lu, Z. Lin, H. Jin, J. Yang, and J. Z. Wang, "Rapid: Rating pictorial aesthetics using deep learning," in Multimedia. ACM, 2014. [Online]. Available: http://doi.acm.org/10.1145/2647868. 2654927

[7] X. Jin, J. Chi, S. Peng, Y. Tian, C. Ye, and X. Li, "Deep image aesthetics classification using inception modules and fine-tuning connected layer," in Wireless Communications & Signal Processing (WCSP), 2016 8th International Conference on. IEEE, 2016, pp. 1–6.

[8] S. Kong, X. Shen, Z. Lin, R. Mech, and C. Fowlkes, "Photo aesthetics ranking network with attributes and content adaptation," in European Conference on Computer Vision. Springer, 2016, pp. 662–679.

[9] L. Mai, H. Jin, and F. Liu, "Composition-preserving deep photo aesthetics assessment," in CVPR, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 497–506.

[10] R. Datta, D. Joshi, J. Li, and J. Z. Wang, "Studying aesthetics in photographic images using a computational approach," in ECCV. Springer, 2006. [Online]. Available: http://dx.doi.org/10.1007/11744078 23

[11] Y. Ke, X. Tang, and F. Jing, "The design of high-level features for photo quality assessment," in CVPR. IEEE, 2006.

[12] M. Nishiyama, T. Okabe, I. Sato, and Y. Sato, "Aesthetic quality classification of photographs based on color harmony," in CVPR. IEEE, 2011.

[13] S. Dhar, V. Ordonez, and T. L. Berg, "High level describable attributes for predicting aesthetics and interestingness," in CVPR. IEEE, 2011. [Online]. Available:http://dx.doi.org/10.1109/CVPR. 2011.5995467

[14] L. Marchesotti, F. Perronnin, D. Larlus, and G. Csurka, "Assessing the aesthetic quality of photographs using generic image descriptors," in ICCV. IEEE, 2011.

[15] N. Murray, L. Marchesotti, and F. Perronnin, "Ava: A large-scale database for aesthetic visual analysis," in CVPR. IEEE, 2012.

[16] Y. Luo and X. Tang, "Photo and video quality evaluation: Focusing on the subject," in ECCV. Springer, 2008.

[17] P. Obrador, X. Anguera, R. de Oliveira, and N. Oliver, "The role of tags and image aesthetics in social image search," in WSM. ACM, 2009. [Online]. Available: http: //doi.acm.org/10.1145/1631144.1631158

[18] O. Wu, W. Hu, and J. Gao, "Learning to predict the perceived visual quality of photos," in ICCV. IEEE, 2011.

[19] J. Machajdik and A. Hanbury, "Affective image classification using features inspired by psychology and art theory," in Multimedia. ACM, 2010.

[20] P. Isola, J. Xiao, A. Torralba, and A. Oliva, "What makes an image memorable?" in CVPR. IEEE, 2011.

[21] M. Redi, N. O'Hare, R. Schifanella, M. Trevisiol, and A. Jaimes, "6 seconds of sound and vision: Creativity in micro-videos," in CVPR. IEEE, 2014.

[22] M. Redi and B. Merialdo, "Where is the beauty?: Retrieving appealing videoscenes by learning flickr-based graded judgments," in Multimedia. ACM, 2012. [Online]. Available: http://doi.acm.org/10.1145/2393347.2396486

[23] Y.-G. Jiang, Y. Wang, R. Feng, X. Xue, Y. Zheng, and H. Yang, "Understanding and predicting interestingness of videos." in AAAI, 2013.