

And Now for Something Completely Different: Visual Novelty in an Online Network of Designers

Johannes Wachs
Central European University
Budapest, Hungary
wachs_johannes@phd.ceu.edu

Bálint Daróczy
Institute for Computer Science and
Control, Hungarian Academy of
Sciences (MTA SZTAKI)
Budapest, Hungary
daroczyb@ilab.sztaki.hu

Anikó Hannák
Central European University
Budapest, Hungary
hannaka@ceu.edu

Katinka Páll
Institute for Computer Science and
Control, Hungarian Academy of
Sciences (MTA SZTAKI)
Budapest, Hungary
pall.katinka@sztaki.mta.hu

Christoph Riedl
Northeastern University
Boston, MA
Harvard University
Cambridge, MA
c.riedl@neu.edu

ABSTRACT

Novelty is a key ingredient of innovation but quantifying it is difficult. This is especially true for visual work like graphic design. Using designs shared on an online social network of professional digital designers, we measure visual novelty using statistical learning methods to compare an image's features with those of images that have been created before. We then relate social network position to the novelty of the designer's images. We find that on this professional platform, users with dense local networks tend to produce more novel but generally less successful images, with important exceptions. Namely, users making novel images while embedded in cohesive local networks are more successful.

KEYWORDS

Novelty; image analysis; neural networks; Fisher information; social networks

ACM Reference Format:

Johannes Wachs, Bálint Daróczy, Anikó Hannák, Katinka Páll, and Christoph Riedl. 2018. And Now for Something Completely Different: Visual Novelty in an Online Network of Designers. In *WebSci '18: 10th ACM Conference on Web Science, May 27–30, 2018, Amsterdam, Netherlands*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3201064.3201088>

1 INTRODUCTION

High-quality creative design work can create tremendous value for organizations. It helps technical products gain acceptance [26] and it often serves as the basis for competition in cultural markets [59].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebSci '18, May 27–30, 2018, Amsterdam, Netherlands

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5563-6/18/05...\$15.00

<https://doi.org/10.1145/3201064.3201088>

Consequently, there has been mounting interest in the use of designers by organizations as a source of value creation [43, 46, 47]. One important ingredient to successful designs is novelty: the degree to which a design is new, original, or unusual relative to what has come before. One reason for this is that derivative work is frowned upon in creative fields [6]. Indeed novelty is the prime ingredient of innovation and the production of new things [19]. Economists have long known that innovation is the driving influence behind economic growth and development [50], and recent studies suggest that successful companies make 80% of their revenue with products younger than five years [34].

Despite its importance novelty is difficult to measure, especially in the context of creative design. In this paper we investigate three research questions related to novelty in design: (1) how can we measure novelty in digital design? (2) who produces novel work? and (3) what is the relationship between novelty and success? We develop and compare different mathematically-grounded measures of novelty or distinctiveness of digital images to better understand its antecedents and subsequent effect on success in a community of professional designers.

To investigate these questions we collect roughly 40,000 images posted by over four thousand professional designers on an online community over a period of about four years. We propose and evaluate a measure of novelty for digital design at the image level using two feature sets: one capturing content and structure defined using an Inception neural network, the other capturing visual aesthetics using classical compositional features. We visualize the distributions of images in low dimensional projections of these feature spaces to better understand what these features capture and how they may capture novelty of an image.

We calculate novelty by comparing an image with prior images in terms of these derived features using information theoretic methods. This focus on temporal order distinguishes novelty from more “timeless” notions like beauty or appeal [18]. Calculating novelty using the compositional features yields a measure of aesthetic or style novelty based on colors, spatial arrangement, and symmetry,

while using Inception features results in a measure of content novelty. We validate our measures by showing that the earliest images annotated with emerging labels or “tags” for new kinds of designs are indeed more content-novel.

With these measures of novelty for digital design in hand, we ask two questions: who produces novel images? How does novelty relate to success? The social networks literature makes two suggestions. Individuals with open, diverse social networks have access to diverse sources of information, which they may synthesize in novel ways [10, 24]. But individuals in cohesive, closed networks have greater access to trust and social support, allowing them to more easily take the risk inherent in the creation of novelties [12, 35]. The literature suggests that when the domain is quickly changing and when the space of possible novelties is large, it is rather cohesive networks that facilitate novelty [3]. We argue that our topic of study is such a domain: design evolves quickly and new trends can be drastically different, and so we hypothesize that cohesive local networks do more to facilitate novelty in this domain than diverse ones.

Using a regression framework to analyze our panel data, we find a positive relationship between the local cohesion of a user’s network on the site and the novelty of her images. Users in the global center of the network make less novel images. We suggest one possible explanation: that standing out is a form of risk-taking and that local network density facilitates this behavior. Furthermore, we find that novel images are on average less successful, but can be successful when originating from the right network position. Finally, we demonstrate that our novelty measures add explanatory power to a machine learning model predicting success, above and beyond a user’s network position. This suggests that network position does not entirely mediate the relationship between novelty and success.

Our paper makes three contributions to the literature. First, we qualitatively compare the data encoded in different feature sets that can be derived from images. Second, we define a statistically-sound measure of the novelty of images, applicable to either set of features. Third, we provide empirical evidence for relationships between novelty, network position, and success, showing that novelty and network position together can predict success.

2 RELATED WORK

In this section we first survey research on the quantification of novelty. Next, we overview literature on the relationship between novelty, social network position, and success, and finally, introduce studies that have looked at design in an online setting.

2.1 Quantifying Novelty

As novelty is a complex construct with various dimensions, many different measures of it have been proposed [9, 21, 44]. One key notion underlying the measurement of novelty is the concept of recombination: that novelty is the result of reconfiguration of old ideas [58]. Novelty is distinguished from aesthetic quality or beauty because it carries an intrinsic temporal property. For example, it is difficult to judge in retrospect how novel a product was at the time of its release. Previous studies on the beauty of images utilize the

fact that crowdsourced judgments of beauty are relatively stable over time [49].

Recent models of novelty frame it in terms of the “actual” and the “possible”. Models consider what it means for something to be new in terms of a path of discovery in an evolving complex space [37]. When something new is done for the first time, the space of the “adjacent possible” grows, making new things possible. In this framework, novelties are discrete, binary events.

Previous work from the data mining community on novelty of images has mostly been concerned with the detection of outliers or anomalies within images, rather than across images [8, 52]. Most applications concern the detection of verifiable facts about an image: the presence of specific objects in satellite images, detecting biological abnormalities like cancer, etc. One commonality across these efforts, and indeed our own, is that features need to be extracted from an image to make computational analysis tractable.

Several recent data-driven studies quantify novelty in creative fields. In a study of popular music, Askin and Mauskopf compare songs with their predecessors using cosine similarity of a set of derived features like danceability and tempo [4]. Past work has quantified the creativity of visual art as a combination of both novelty and influence using visual features [18]. Redi et al. quantify the novelty of short video clips using a similar approach to ours [44], while Khosla et al. use image features to predict engagement on social media [32]. Natural language processing has also been applied to measure the novelty of textual content including scientific article abstracts [9, 20] and equity crowdfunding campaigns [29]. We do not define novelty of a thing in terms of success [51] or surprise [5]. Novelty of a thing as we consider it says nothing intrinsically about its impact, influence, or outcomes. At the same time, we acknowledge that any attempt to measure novelty or distinctiveness can only capture a small facet of the phenomenon.

2.2 Network Position, Novelty, Success

Psychological research emphasizes that creativity is a demanding enterprise, requiring focus and concentration [13]. Given the apparent difficulty of creative endeavors, it is perhaps no surprise that social network structure plays a significant role in both facilitating novelty and shaping its reception. In fact, recent studies of creativity emphasize that novel products, even nominally created by a single author, can sometimes be understood as “products of a momentary collective process” [27]. How the networks that synthesize creative products fit together have strong predictive power of their eventual success [16].

So what kind of network position facilitates novelty? Creators embedded in a cohesive social network can hope to benefit from high amounts of social capital and support [12]. Strong ties represent avenues of trust, which greatly facilitates the kind of risk-taking inherent in making a novel product in a professional, creative environment [35]. One study indicates that central actors in a network of research scientists produce more creative outputs, indicating that established actors can feel the freedom to experiment more broadly [42].

It is also true that diversity of social connections has been shown to foster creativity. Weak ties in social networks tend to bridge groups and provide an actor with access to novel information [24].

Indeed the same study of research scientists cited above shows that creativity increases with the number of weak ties [42]. This line of thought is built on the idea that bridging actors occupying “structural holes” can create their own social capital by leveraging their unique access to diverse information [10]. Whether open or closed networks better support novelty creation in our context is therefore an empirical question.

Besides the relationship of network position and novelty, the perception of novelty is also of interest to the research community. What ratio of traditional and novel maximizes success? Work across many disciplines find an inverse-U shaped relationship between novelty and success [4, 9]. One prolific strand of the literature models novelty as the recombination of known ideas in new ways, and that the key to successful novelty is the combination of many conventional ingredients with relatively few new ones [56].

2.3 Online Design Communities

Closest to our work empirically are studies on online design communities, like Dribbble, Behance, or Threadless. These studies generally focus on the question of how users or products become successful, and how different groups of users fare [17, 45]. For instance several studies find significant differences in the behavior and success of men and women on these sites [33, 57].

Dribbble has received attention from researchers because of its importance to the professional design community and its exclusive, invitation-only nature. In an interview-based study researchers found that users leverage the site and its social network to gather inspiration, learn skills by reverse engineering examples, anticipate trends in the marketplace, and to gather feedback [40]. The study also found that users invested significant effort in developing a professional identity through the site. As in many other online communities, users reported the status importance of having many followers and collecting likes.

More recently, machine vision researchers have taken an interest in learning from image data taken from online design communities, as they offer substantially different opportunities to develop machine vision than, say, photographs [60]. Similarly, the dual functions of online digital communities as places to post and places to be inspired offer interesting opportunities for bespoke recommendation systems [48].

3 DATA

In this section we describe the Dribbble platform, our data collection method, and outline the extracted features at the image, user, and network levels.

3.1 Dribbble

Dribbble, founded in 2009, is an online community where designers share their work by posting images. It is a highly-visited site, with an Alexa rank of 1104, the second most popular website for design sharing after Behance. Unlike most content-sharing platforms, the site operates on an invitation-only basis: though the site can be viewed by anyone, only invited users can post images. Active users are occasionally given invitations which they can use to invite other designers. Moreover, the number of images a user can post in a given time frame is capped. All together, this leads to high-quality

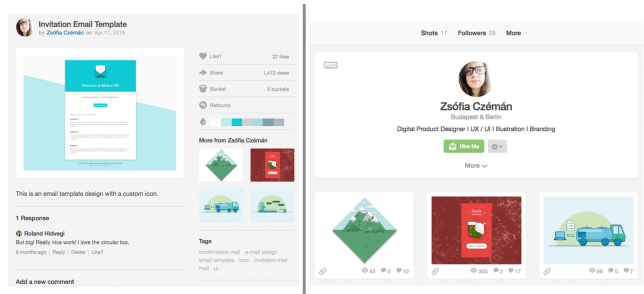


Figure 1: Shot (Image) and User Pages on Dribbble.

content and the feeling of belonging to an “elite” community among users.

The stakes on Dribbble are high. Interviews with users on the site reveal that individuals use the site to develop their professional identities [40]. Indeed most users use their real names, post photographs of themselves for their account image, and link to their accounts on other online platforms including LinkedIn and Twitter. Users build their portfolio of designs over many years. They accumulate reputation by gathering views and likes (engagement) on their images, called shots on the site, and followers on their account. The social network aspect of the site facilitates continued interactions as users see more and more of each others’ work. Success on Dribbble has impact outside the site itself, as it can bring significant employment opportunities and influence. The platform has recently added a job board and special recruiter accounts.

3.2 Data Collection

Our data sample consists of all Dribbble users who were members of a team at the time of the data collection. Typically companies form teams on Dribbble as paid umbrella accounts that users can join. We select this sample in order to gather a comparable set of users who are both active and committed members of the site. We then crawled the profiles of 6,215 users identified as team members. Next, we crawled all 60,406 images made by these users. In subsequent analysis, we discard users making fewer than five images¹. We also discard images posted by the team account with identifiable individual author. We share examples of an image and a user page in Figure 1. Data collection took place between September and November 2016 and observed listed rate-limits on the Dribbble API.

3.3 Extracted User features

At the shot level we first record the image itself, the date it was made, and the identity of the author. We also note the tags the author annotated the shot with. Tags are free-form key words that say something about the image. Others can search for images listing specific tags. Tags therefore serve a dual purpose: to describe what the author is doing, and to help others find the image. Each shot has a count of the likes that it received, which can be thought of as the main success measure in the community.

At the user level we collect the name of the author, whether the author has a “pro-badge”, and the author’s tenure on the platform

¹Our results are robust to including these users.

(in days). A pro-badge is a sign that the user has paid for a premium account, which facilitates job search features on the site and lifts the cap on the number of shots a user can make in a given amount of time. We consider pro-badges as a proxy for buy-in on the platform. At the shot level we calculate how many shots a user has made before to quantify their experience. Finally, we also estimate the gender of each user. Since the profiles do not directly list gender, we infer them from the users' first names using the US baby name data set [1]. For any user with a name not in the database or an ambiguous gender score (i.e. greater than 10% and less than 90%) we manually check their self-portrait on Dribbble and on linked social media accounts.

3.4 Network Features

Like many other online communities, Dribbble is built on top of a social network. When a user follows another user, the second user's future shots are included in the default newsfeed of the first user and so following a user has bandwidth costs. We collect a list of all following relationships amongst our users and when they were created. These timestamped edges allow us to recreate the social network of our users at the time when an image was submitted. For each image we calculate several network measures quantifying the position of the user at the time of creation.

- *In-degree*: How many followers the user has.
- *Out-degree*: How many other users the user follows.
- *Closeness centrality*: One over the average distance of the user from all other nodes [7]. This measures how close the user is to the center of network.
- *Constraint*: Burt's measure of the extent to which a user's outgoing connections are redundant [10].
- *Density*: The ratio of observed ties to possible ties among the users the user follows.

In- and out-degree quantify the simple connectivity of a user. Closeness centrality is a global network measure which increases as the user is closer to the center of the network. Constraint and density of the user measure the cohesiveness of his local social network.

4 EXTRACTING IMAGE FEATURES

In this section we describe two sets of images features upon which we calculate an image's novelty. First we calculate *compositional features*. Then we use a neural network framework to extract a set of unsupervised features. We compare the two feature spaces by projecting them to a low-dimensional space in which similar images are placed closer to one another. We examine what kind of images are similar according to the two feature sets, finding that the compositional features capture color and style while the neural network features capture content.

4.1 Compositional Features

Imitating precisely previous work on the qualitative features of images [49], we define 47 compositional features for each image. These features are derived from aesthetic considerations and have proven to have significant predictive power of the beauty or attractiveness of images. Previous work groups the features into the following categories: colors, spatial arrangements, and texture.

Color features include contrast (defined in terms of luminance) and the averages of hue, saturation, and brightness across both the whole image and a subset in its center [15]. We also include three "emotional" features which are linear combinations of saturation and brightness: pleasure, arousal, and dominance [39]. Binning hue, saturation, and brightness yield *Itten Color Histograms* and taking their standard deviations yields *Itten Color Contrasts* after a careful segmentation. Spatial features include symmetry and salience [30], the distribution of which describes how attention-grabbing different regions of the image are. Finally, Haralick's texture features quantify image complexity: entropy, energy, homogeneity, and contrast [25].

4.2 Neural Network Features

Feedforward-based neural networks have made tremendous strides in object-in-image classification tasks in recent years. Many such networks have penultimate layers which reduce images input for classification into a feature space for the classification layer. It is possible to extract these features from pre-trained neural networks. We harness one such network: the Inception v3 [54], originally constructed to optimally classify a large dataset of images into 1000 categories. We acknowledge here that there are many alternative specifications to generate similar sets of features. Passing our images through the network we generate 2048 features that encode highly discriminating facets of the data.

4.3 Visualizing Image Features

Before proceeding, we pause to visualize and inspect our data in the two visual feature spaces. We reduce the 47 and 2048 dimensional spaces to two-dimensions using t-SNE, a popular dimensionality reduction method that uses information theoretic methods to minimize distances between data points in the projection as a function of their similarity [38]. In Figure 2, we visualize the 2-D t-SNE projections of a random sample of 200 images a year from 2012 to 2016 using the Inception and compositional features, respectively.

In both projections we observe the clustering of images into groups. The qualitative attributes that define the clustering, however, are quite different. As highlighted in Figure 2, clustering on compositional features is based on color and aesthetic style, as expected. In the projection based on Inception features, however, we observe that images cluster based on their content. In other words, images with highly similar Inception features are likely to represent similar concepts, be they logos, mobile phone interfaces, icons, wireframes, etc. This is perhaps not surprising given Inception's origin as an object-in-image classification tool. This characterization of the two features sets as describing style and content is important for understanding their novelty.

5 NOVELTY MEASURES

In this section we define a reference novelty based on user annotations or tags of an image by defining the relative surprise of seeing a set of tags on image, compared with the tags that came before. We then define a measure of novelty for our visual feature spaces using Gaussian mixtures and Fisher information.




Figure 2: Visualizing sample images using t-SNE dimensionality reduction of Inception and compositional features. We highlight three example groups of images. Images from the gold group are close together in the compositional feature space but spread out in the Inception feature space. The teal group is close in both feature spaces, with one exception in the compositional space. Images from the purple group are close in Inception space but scattered in compositional space. The gold group consists of a logo, a collection of icons, a web page design, and an email flier: they are likely clustered in compositional space because of their color. The members of the purple group are all mobile phone screens. Members of the teal group are likely clustered in both spaces because they share both structural and compositional qualities.

5.1 Tag Novelty

Before calculating novelty using visual features, we create a novelty measure using the tags an author gives an image. Following [53], we calculate the “surprise” of each tag of an image. That is, given all the images and their tags posted before the image, we define the probability of observing a tag t as $P(t)$, the proportion of previous images listing that tag. The log of $P(t)$ is our measure of the surprise of a tag. As we are especially interested in completely new tags, we also include the focal image and its tag when we calculate $P(t)$, to avoid taking the log of 0. We then define tag novelty N_i of an image i with tags $t_1, t_2, \dots, t_n \in T_i$ as the aggregate the surprise of an image’s tags:

$$N_i = -\frac{1}{|T_i|} \sum_{t \in T_i} \log P(t)$$

In order to make our measure robust to the order of the images, we scale each image’s tag novelty by the maximum possible novelty. Namely, if I is the number of images made before image i , we normalize the equation above by $-\log(|I|)$.

5.2 Visual Novelty via Fisher Information

To study the visual novelty of images, we define a parametric model for images in terms of their position in a given feature space. Given a new image, we consider the distribution of previous images in a feature space and approximate them using Gaussian mixture models. We calculate the likelihood of the focal image relative to these

distributions using its Fisher information, an information theoretic measure which we prefer to alternatives such as the Akaike information criterion because of its reparametrization-invariance. Specifically we define novelty as one minus the norm of the Fisher vector of an image over the Gaussian mixture models. This approach is similar in style to a recent method to calculate novelty using a data point’s distance to the centroids of a k-means clustering [44].

Formally, let be $x \in \mathbb{R}^d$ a finite d-dimensional real representation of an image and a parametric model $p(x|\theta)$ where θ is the parameter of the density function. If the model is a Gaussian mixture model (GMM) with N Gaussians, the pdf is $p(x|\theta) = \sum_i^N \omega_i g_i(x)$ where the $g_i(x)$ is the density function of the i -th Gaussian. The continuously evolving model changes the parameters of the probabilistic model with the emergence of new images in time. We consider two different likelihood measures to apply to the probabilistic model:

- Akaike’s information criterion (AIC) [2]: we measure the AIC per image according the actual state of our generative model.
- Fisher information: after calculating the Fisher score [31] of for each image according to the shape of the model we can measure the similarity of images x and y with the Fisher kernel, as

$$K_\theta(x, y) = \nabla_\theta \log p(x|\theta)^T F_\theta^{-1} \nabla_\theta \log p(y|\theta) \quad (4)$$

where F_θ is the Fisher information matrix. The gradient of the likelihood indicates how the model may change to fit the actual point, in our case an image. Our choice was driven by

the unique invariance properties (e.g. reparametrization invariance) of the Fisher information matrix and the Fisher kernel [11, 36, 55]. Applying Cholesky decomposition, the kernel can be defined as a simple scalar product, as $K_\theta(x, y) = \mathcal{G}_\theta(x)^T \mathcal{G}_\theta(y)$ where $\mathcal{G}_\theta(x) = \nabla_\theta \log p(x|\theta) F_\theta^{-1/2}$ is the normalized Fisher score or the Fisher vector of image x . We note that the Fisher vector has dimension $O(d|\theta|)$.

On account of its reparametrization invariance we choose to continue with the Fisher information as our measure of likelihood. Although estimation of the Fisher information matrix is difficult, there are known closed form approximations for both Gaussian mixture models [41] and special classes of Markov random fields [14]. We suggest two potential definitions of novelty measures based on the Fisher information:

- Norm of the Fisher Vector over Gaussian Mixture (FVGMM): as the Fisher score highlights how the model parameters should change to best fit the focal image, our first novelty measures the norm of the Fisher vector for each image as

$$N_{FV}(x) = \|\mathcal{G}_\theta(x)\| = \|\nabla_\theta \log p(x|\theta) F_\theta^{-1/2}\|. \quad (2)$$

In case of Gaussian Mixtures the pdf is $p(x|\theta) = \sum_i^N \omega_i g_i(x)$ where θ consists of the mixture weights, mean, and covariance parameters of the Gaussian mixture. In practice we observe that the Fisher score for both compositional and Inception features is very sparse because of the “peakness” property of the *membership probability*, defined as the probability that a point is generated from one of the Gaussians. In comparison with [4] this method puts the most weight on the most similar images that came before the focal image.

- Similarity graph over the Gaussian Mixture (FVMRF): one approach to overcoming the “peakness” property while still capturing the temporal distribution is to define a Markov random field following [14] with the mean of the Gaussian mixture as the sample set. The main idea is to define an undirected random field, which is a graph with N nodes consisting of random variables and sample points, connected to our image as a separate random variable in a star. The probability density function of the new distribution can be factorized over the maximal cliques in the resulting graph. In our case the edges and therefore the pdf are:

$$p(x|\alpha, \theta) = \frac{e^{-\sum_i \alpha_i \|x - \mu_i\|}}{\int_{x \in \mathcal{X}} e^{-\sum_i \alpha_i \|x - \mu_i\|} dx} \quad (3)$$

where μ_i is the mean vector of the i -th Gaussian, α is the relative importance of the cliques. The Fisher vector can be approximated in this context with a simple formula [14]:

$$N_{FVMRF}(x) = \left\{ \frac{d_i(x) - \mathbf{E}[d_i(x)]}{\text{Var}^{-\frac{1}{2}}(d_i(x))} \right\}$$

where $d_i(x) = \|x - \mu_i\|$ and $i \in 1, \dots, N$.

Given the relatively high complexity of the random field approach, we define novelty using the norm of the Fisher vector². As the method returns a similarity score, we subtract one to define visual novelty. For the rest of the paper we refer to this novelty score

²In applications where the aforementioned peakness issue is more pronounced, we recommend using the random field approach




Figure 3: Kernel density estimated distributions of tag, compositional, and Inception novelty.

as *Inception novelty* when it is calculated using Inception features, and *compositional novelty* when it is calculated using compositional features.

5.3 Comparison of Novelty Scores and Validation

We visualize the distribution of tag, Inception, and compositional novelty scores in Figure 3. We correlate the two novelties with tag novelty and several user-level features in Table 1. We find that both visual novelties are weakly correlated with tag novelty. The correlation is roughly twice as strong for Inception novelty than compositional novelty. This suggests that tags are used to describe images in a conceptual rather than stylistic manner. The two visual novelties are significantly correlated, and, together with tag novelty, are negatively correlated with engagement. We note that the platform’s design may explain the trade-off between engagement and tag novelty: users can search for images by tags.

5.3.1 Validation of Visual Novelty. As discussed, novelty is an ephemeral quality of a cultural product and its measurement implicitly requires comparison, more so than, for example, its beauty. We cannot, for instance, ask someone to evaluate the novelty of a four-year-old mobile phone application layout. In this case success and perceptions of novelty are likely anti-correlated: success breeds familiarity.

One approach to validate our measures of visual novelty, besides the correlations with tag novelty noted above, is to identify a population of images which are likely to be covering a new kind of product that emerges in the middle of our dataset. We identify emerging product types by finding tags which are used only after 2013, yet still are among the 200 most used tags. We find two such tags³ which we can verify as representing truly emerging novelties: “material” and “principle”.

³Other examples of tag fitting our quantitative criteria are tags used by groups of designers to indicate group membership. Though these tags certainly merit further study, they do not capture the emergence of a new design approach or method

| | (1) | (2) | (3) | (4) |
|---------------------------|--------|--------|--------|-------|
| Tag Novelty (1) | | | | |
| Inception Novelty (2) | 0.123 | | | |
| Compositional Novelty (3) | 0.067 | 0.274 | | |
| Likes (Log) (4) | -0.138 | -0.082 | -0.014 | |
| Views (Log) (5) | -0.138 | -0.114 | -0.058 | 0.927 |

Table 1: Correlation matrix of novelty and success features.

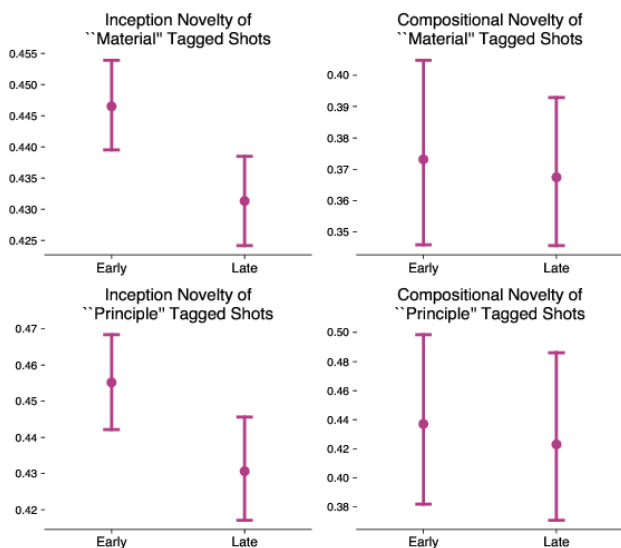


Figure 4: Comparison of visual novelty scores of images with the tags “material” and “principle”. We consider those images in the first 10% and most recent 10% of all images created using the tags. We find that Inception novelty is significantly higher for images listing these “emerging” tags.

Material design⁴ is a design language or vocabulary created by Google, announced to the public in June 2014. Like other design languages, it has guidelines and principles that shape the design process, resulting in a consistent look with certain qualities. Material design was created especially for use in digital and technological areas. It emphasizes the use of print design best practices together with motion. Material or “material design” appears as a tag in 748 images in our dataset.

Principle⁵ is a new software design tool for creating interactive and dynamic user interfaces. Released in August 2015, it is a popular tool for designers to prototype UIs. 243 images in our dataset include a “principle” tag.

For both tags we compare the distributions of novelty for the first 10% of images using the tag, with the most recent 10% of images using the tag. In figure 4 we plot the resulting distributions. We find that Inception novelty is significantly higher for the earliest images tagged with “material” (Mann-Whitney $U = 1897$, $p < .01$) and “principle” (Mann-Whitney $U = 190$, $p < .01$) compared with the most recent ones. Though the average compositional novelty is higher for the earliest images in both cases, the differences are not statistically significant (resp. $U = 2465$, $p = .26$; $U = 288$, $p = .32$).

6 NOVELTY, NETWORKS, AND SUCCESS

In this section we investigate which users are more likely to create novel images and whether novel images are more or less likely to be successful. We consider both Inception and compositional novelty.

⁴<https://material.io/>

⁵<http://principleformac.com/>

First we use hierarchical linear regression [23] on data at the image level with user random-effects and controls to predict novelty. Our aim is understand who makes novel images. Then we predict success using novelty and network position. In both cases we control for gender, the (log) number of shots made previously, the (log) number of days the user has been active on the site at the time of the shot, and whether the user has a paid account. In other words we control for gender, productivity/experience, tenure, and investment into the site.

6.1 Who makes novel shots?

We find several significant predictors of Inception novelty, both among our control variables and network variables. Interestingly, the network features we consider do not impact compositional novelty. We summarize these findings in Table 2.

For both compositional and Inception-based measures we find that pro-users are less likely to make novel images. One interpretation is that users who take the site more seriously are more risk-averse and less likely to experiment. Users making more shots in the past make slightly more novel shots. There is mixed evidence that users active for a longer period of time make less novel shots. We detect no gender disparity.

The two novelty measures diverge when we consider the impact of network features. The Inception-based measure of novelty is significantly lower for users closer to the core of the network, and higher for users with cohesive local networks defined by density and constraint. This supports our hypothesis that cohesion facilitates novelty. We find no significant relationship between network position and compositional novelty.

6.2 When are novel shots successful?

We now turn to the question of predicting engagement, measured by likes, using novelty. We find that novel shots are generally less successful. We summarize our findings in Table 3. Pro users are more successful, as are those who have many followers. We find that constrained users are less successful. Finally, novel images are in general less successful.

We find an interesting interaction between constraint and Inception novelty. Namely, users embedded in highly constrained networks making novel images do better than those in unconstrained networks making novel images. To better interpret this finding we visualize this relationship in Figure 5. In other words, the least constrained users have a penalty for novelty while the most constrained users have a bonus for novelty. We also find a significant interaction between inception novelty and closeness centrality: novelty has an increasingly negative relationship with success as a user is more central in the network, but no relationship between local density and either novelty measure.

Finally, using a machine learning framework, we check how well our features can predict success binned into three separate class labels: less than ten likes, between ten and one hundred likes, and more than one hundred likes. As an initialization we used the first year as the first training period and for every consecutive quarter thereafter we consider the previous year. We found that the random field approach to calculating the Fisher vector (FVMRF) was most effective in predicting engagement. Using the area under

| | Dependent variable: | | | | | |
|---------------------|---------------------|-------------------|-------------------|---------------------|-------------------|-------------------|
| | Inception Novelty | | | Composition Novelty | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Days Active (log) | -0.015*** (0.002) | -0.016*** (0.002) | -0.017*** (0.002) | -0.000 (0.002) | 0.000 (0.002) | -0.000 (0.002) |
| nShots Previous | 0.020*** (0.007) | 0.022*** (0.007) | 0.022*** (0.007) | 0.013* (0.007) | 0.013* (0.007) | 0.013* (0.007) |
| Male | -0.002 (0.013) | 0.001 (0.013) | 0.001 (0.013) | -0.000 (0.013) | -0.000 (0.013) | -0.000 (0.013) |
| Pro | -0.036*** (0.010) | -0.034*** (0.010) | -0.034*** (0.011) | -0.034*** (0.011) | -0.034*** (0.011) | -0.034*** (0.011) |
| In-Degree (log) | 0.004 (0.004) | 0.002 (0.004) | 0.001 (0.004) | 0.004 (0.004) | 0.004 (0.004) | 0.004 (0.004) |
| Closeness | -0.042*** (0.008) | | | 0.001 (0.008) | | |
| Constraint | | 0.060** (0.025) | | | 0.018 (0.026) | |
| Density | | | 0.046** (0.023) | | | -0.005 (0.024) |
| Constant | 0.002 (0.025) | -0.009 (0.026) | -0.002 (0.026) | -0.048* (0.026) | -0.053** (0.026) | -0.047* (0.026) |
| Observations | 37,799 | 37,799 | 37,799 | 37,799 | 37,799 | 37,799 |
| Log Likelihood | -25,740.880 | -25,749.400 | -25,750.350 | -25,731.900 | -25,730.540 | -25,730.860 |
| Bayesian Inf. Crit. | 51,576.620 | 51,593.660 | 51,595.570 | 51,558.660 | 51,555.950 | 51,556.580 |

User random effects

*p<0.1; **p<0.05; ***p<0.01

Table 2: Predicting novelty with network position.

| | Dependent variable: | |
|--------------------------|---------------------|-------------------|
| | Log Likes | |
| | (1) | (2) |
| Days Active (log) | -0.006** (0.003) | -0.004* (0.003) |
| nShots Previous | 0.086*** (0.023) | 0.090*** (0.023) |
| Male | -0.046 (0.044) | -0.047 (0.045) |
| Pro | 0.196*** (0.037) | 0.196*** (0.037) |
| In-Degree (log) | 0.371*** (0.009) | 0.373*** (0.009) |
| Out-Degree (log) | -0.046*** (0.013) | -0.046*** (0.013) |
| Constraint | -0.234*** (0.059) | -0.233*** (0.059) |
| Incep. Nov. | -0.108*** (0.009) | |
| Incep. Nov. × Constraint | 0.084** (0.039) | |
| Comp. Nov. | | -0.025*** (0.010) |
| Comp. Nov. × Constraint | | 0.017 (0.040) |
| Constant | 2.930*** (0.088) | 2.908*** (0.089) |
| Observations | 37,799 | 37,799 |
| Log Likelihood | -36,353.290 | -36,450.650 |
| Bayesian Inf. Crit. | 72,833.060 | 73,027.780 |

User random effects

*p<0.1; **p<0.05; ***p<0.01

Table 3: Predicting success with novelty and network position.

the receiver operating characteristic curve (AUC), we find that a gradient boosted trees model [22] on network and content features has significant predictive power. As we can see in Figure 6, even though the network features are the best indicators of success, the content and novelty of the images, encoded using the Inception-based Fisher vectors offer additional predictive power. This suggests that it is possible to use image features to predict success on the




Figure 5: Relationship between success and novelty as constraint varies. Low constraint users have less success when making novel shots. High constraint users have more success with novel shots.

site. It is likely possible to do better if features are extracted with the aim of predicting success.

7 CONCLUSIONS

In this paper we developed, evaluated, and compared measures of novelty of images using data from an online community of digital designers. We first compared different feature sets of images, noting that compositional features like entropy, contrast, and brightness capture qualitatively different facets of an image and features derived from an Inception neural network learning framework capture qualitatively different facets of an image. Specifically, compositional features seem to capture stylistic aspects and while Inception features capture content, in line with their origins.




Figure 6: Average AUC of quarter to quarter success prediction. We predict success of the images using gradient boosted trees on the visual image features, novelty scores, and network position of the users. We find that novelty scores extracted image features increase the predictive power of the model including the network features, This suggests that network position does not entirely mediate the relationship between novelty and success.

Next, we created a mathematical framework to compare images with all images that came before in terms of either set of image features. To calculate the novelty of an image, we estimate the distribution of previous shots in the given feature space using a Gaussian mixture model. We then calculate the likelihood of the the image - in other words we quantify how statistically similar the image is to those that came before. We define novelty of an image as one minus this similarity score.

We find that both novelties calculated from the Inception features and compositional features are significantly correlated with a measure of novelty based on author text annotations or “tags” of their images. We also found that Inception novelty was significantly higher for images created in the early stages of an emerging tag compared with images using the same tag later.

Attempting to understand the profile of a user who makes more novel shots, we turned to the site’s social network. Using temporal following data, we related social network position at the time of the creation of an image to its novelty. We found that users with cohesive local networks (quantified by density or Burt’s constraint measure) tend to post images with higher Inception novelty.

We also find that users close to the center of the network, in a global sense, make less novel shots. Users with a “pro-badge” (paid account) likewise make less novel images. Given the professional atmosphere of the site, including for example its invitation-only participation, the presence of significant players and companies in the field, and the potential for economic opportunities, it seems reasonable that established designers may have reason to make more conventional images. That Dribbble is an online community

only compounds the potential costs of creating unsuccessful novelty: though a designer’s support system and network of strong ties cannot vastly grow, her audience can scale drastically. The underestimated permanence of online identities makes this asymmetry all the more important when we consider what it means for a designer to take a risk with a distinctive image.

Indeed professional online communities present a dilemma for users in general. Though the feelings of anonymity and distance may facilitate bold experimentation, members of online communities who wish to leverage their investment of time and effort into professional advancement must credibly link their online identities to their real ones. Even users who want to stay anonymous often have a hard time doing so [28]. Once this identification has occurred, the individual must consider that anything they share online is widely broadcast and more consistently recorded and preserved than what they may say or share offline. We claim that as the labor market becomes increasingly digital, online social networks merit closer study.

Turning to the relationship between novelty and success, we find that novelty is related to worse outcomes. We also find that users in highly constrained positions are less successful. On the other hand, the interaction between constraint and novelty is positive: users with cohesive local networks of strong ties making novel images find more success. We argue that these relationships merit further study. Are these embedded designers better positioned to take risks? Can we interpret images with high novelty score, according to our definition, as being risky? The negative relationship between novelty and network centrality raises even more questions.

Our study has several limitations. Given the transient nature of novelty, we have only limited tests of validity for our measures. Given the ubiquity of digital technology, a highly novel digital design from five years ago likely looks highly outdated now. Moreover, the networking behavior of designers on this platform is highly tailored to the situation. For example, users adopting a strategy of aggressive following anticipating reciprocity, may end up in highly dense networks. All at once, Dribbble serves as a social network, professional portfolio, information network, and status hierarchy for the field. Any attempt to infer causal relations between social network structure and the creation of new ideas on this platform must disentangle the complicated layers driving interactions. We also concede that novelty is multi-faceted: no single measure can totally capture such a broad concept. In future work we aim to better understand influence and spreading of novelty.

8 ACKNOWLEDGEMENTS

The authors wish to thank Zsófia Czémán, Anna May, and anonymous referees for their helpful suggestions. This research has been funded in part by NSF grant IIS-1514283. D.B was supported by the Momentum Grant of the Hungarian Academy of Sciences (LP2012-19/2012).

REFERENCES

- [1] 2016. Baby Names from Social Security Card Applications-National Level Data. data.gov. (2016). <https://catalog.data.gov/dataset/baby-names-from-social-security-card-applications-national-level-data>.
- [2] Hirotugu Akaike. 1981. Likelihood of a model and information criteria. *Journal of econometrics* 16, 1 (1981), 3–14.

- [3] Sinan Aral and Marshall Van Alstyne. 2011. The diversity-bandwidth trade-off. *Amer. J. Sociology* 117, 1 (2011), 90–171.
- [4] Noah Askin and Michael Mauskopf. 2017. What Makes Popular Culture Popular? Product Features and Optimal Differentiation in Music. *American Sociological Review* 82, 5 (2017), 910–944.
- [5] Andrew Barto, Marco Miroli, and Gianluca Baldassarre. 2013. Novelty or surprise? *Frontiers in psychology* 4 (2013).
- [6] Julia Bauer, Nikolaus Franke, and Philipp Tuertscher. 2016. Intellectual property norms in online communities: How user-organized intellectual property regulation supports innovation. *Information Systems Research* 27, 4 (2016), 724–750.
- [7] Alex Bavelas. 1950. Communication patterns in task-oriented groups. *The Journal of the Acoustical Society of America* 22, 6 (1950), 725–730.
- [8] Giacomo Boracchi, Diego Carrera, and Brendt Wohlberg. 2014. Novelty detection in images by sparse representations. In *Intelligent Embedded Systems (IES), 2014 IEEE Symposium on*. IEEE, 47–54.
- [9] Kevin J Boudreau, Eva C Guinan, Karim R Lakhani, and Christoph Riedl. 2016. Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science. *Management Science* 62, 10 (2016), 2765–2783.
- [10] Ronald S Burt. 2004. Structural holes and good ideas. *American journal of sociology* 110, 2 (2004), 349–399.
- [11] LL Campbell. 1986. An extended Čencov characterization of the information metric. *Proc. Amer. Math. Soc.* 98, 1 (1986), 135–141.
- [12] James S Coleman. 1988. Social capital in the creation of human capital. *American journal of sociology* 94 (1988), S95–S120.
- [13] Mihaly Csikszentmihalyi. 1996. *Flow and the psychology of discovery and invention*. New York: Harper Collins.
- [14] Balint Daroczy, David Siklois, Robert Palovics, and Andras A Benczur. 2015. Text Classification Kernels for Quality Prediction over the C3 Data Set. In *Proceedings of the 24th International Conference on World Wide Web*. ACM, 1441–1446.
- [15] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z Wang. 2006. Studying aesthetics in photographic images using a computational approach. In *European Conference on Computer Vision*. Springer, 288–301.
- [16] Mathijs De Vaan, David Stark, and Balazs Vedres. 2015. Game changer: The topology of creativity. *Amer. J. Sociology* 120, 4 (2015), 1144–1194.
- [17] Biplob Deka, Haizi Yu, Devin Ho, Zifeng Huang, Jerry O Talton, and Ranjitha Kumar. 2015. Ranking designs and users in online social networks. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 1887–1892.
- [18] Ahmed Elgammal and Babak Saleh. 2015. Quantifying Creativity in Art Networks. In *Proceedings of the Sixth International Conference on Computational Creativity June*. 39.
- [19] Maria-Isabel Encinar and Felix-Fernando Munoz. 2006. On novelty and economics: Schumpeter’s paradox. *Journal of Evolutionary Economics* 16, 3 (2006), 255–277.
- [20] James A Evans and Jacob G Foster. 2011. Metaknowledge. *Science* 331, 6018 (2011), 721–725.
- [21] L. Fleming. 2001. Recombinant uncertainty in technological search. *Management Science* 47, 1 (2001), 117–132.
- [22] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.
- [23] Andrew Gelman and Jennifer Hill. 2006. *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press.
- [24] Mark S Granovetter. 1973. The strength of weak ties. *American journal of sociology* 78, 6 (1973), 1360–1380.
- [25] Robert M Haralick. 1979. Statistical and structural approaches to texture. *Proc. IEEE* 67, 5 (1979), 786–804.
- [26] A. Hargadon and R.I. Sutton. 1997. Technology brokering and innovation in a product development firm. *Administration Science Quarterly* 42, 4 (1997), 716–749.
- [27] Andrew B Hargadon and Beth A Bechky. 2006. When collections of creatives become creative collectives: A field study of problem solving at work. *Organization Science* 17, 4 (2006), 484–500.
- [28] Emőke-Ágnes Horvát, Michael Hanselmann, Fred A Hamprecht, and Katharina A Zweig. 2012. One plus one makes three (for social networks). *PloS one* 7, 4 (2012), e34740.
- [29] Emőke-Ágnes Horvát, Johannes Wachs, Aniko Hannak, and Rong Wang. 2018. The Role of Novelty in Securing Investors for Equity Crowdfunding Campaigns. In *The 6th AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*. AAAI.
- [30] Xiaodi Hou and Liqing Zhang. 2007. Saliency detection: A spectral residual approach. In *Computer Vision and Pattern Recognition, 2007. CVPR’07. IEEE Conference on*. IEEE, 1–8.
- [31] Tommi S Jaakkola, David Haussler, et al. 1999. Exploiting generative models in discriminative classifiers. *Advances in neural information processing systems* (1999), 487–493.
- [32] Aditya Khosla, Atish Das Sarma, and Raffay Hamid. 2014. What makes an image popular?. In *Proceedings of the 23rd international conference on World wide web*. ACM, 867–876.
- [33] Nam Wook Kim. 2017. Creative Community Demystified: A Statistical Overview of Behance. *arXiv preprint arXiv:1703.00800* (2017).
- [34] W Chan Kim and René Mauborgne. 1997. *Value innovation: The strategic logic of high growth*. Harvard Business School Pub.
- [35] David Krackhardt. 2003. The strength of strong ties. *Networks in the knowledge economy* (2003), 82.
- [36] Guy Lebanon. 2004. An extended Čencov-Campbell characterization of conditional information geometry. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*. AUAI Press, 341–348.
- [37] Vittorio Loreto, Vito DP Servidio, Steven H Strogatz, and Francesca Tria. 2016. Dynamics on expanding spaces: modeling the emergence of novelties. In *Creativity and Universality in Language*. Springer, 59–83.
- [38] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9, Nov (2008), 2579–2605.
- [39] Jana Machajdik and Allan Hanbury. 2010. Affective image classification using features inspired by psychology and art theory. In *Proceedings of the 18th ACM international conference on Multimedia*. ACM, 83–92.
- [40] Jennifer Marlow and Laura Dabbish. 2014. From rookie to all-star: professional development in a graphic design social networking site. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, 922–933.
- [41] Florent Perronnin and Christopher Dance. 2007. Fisher kernels on visual vocabularies for image categorization. In *Computer Vision and Pattern Recognition, 2007. CVPR’07. IEEE Conference on*. IEEE, 1–8.
- [42] Jill E Perry-Smith. 2006. Social yet creative: The role of social relationships in facilitating individual creativity. *Academy of Management journal* 49, 1 (2006), 85–101.
- [43] D. Ravasi and G. Lojacono. 2005. Managing design and designers for strategic renewal. *Long Range Planning* 38 (2005), 51–77.
- [44] Miriam Redi, Neil OHare, Rossano Schifanella, Michele Trevisiol, and Alejandro Jaimes. 2014. 6 seconds of sound and vision: Creativity in micro-videos. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*. IEEE, 4272–4279.
- [45] Christoph Riedl and V Seidel. 2018. Learning from Mixed Signals: Evidence from a Contest-based Online Innovation Community. *Organization Science* (2018).
- [46] V. Rindova, E. Dalpiaz, and D. Ravasi. 2011. A cultural quest: A study of organizational use of new cultural resources in strategy formation. *Organization Science* 22, 2 (2011), 413–431.
- [47] V.P. Rindova and A.P. Petkova. 2007. When is a new thing a good thing? Technological change, product form design, and perceptions of value for product innovations. *Organization Science* 18, 2 (2007), 217–232.
- [48] Maja R Rudolph, Matthew Hoffman, and Aaron Hertzmann. 2016. A joint model for who-to-follow and what-to-view recommendations on behance. In *Proceedings of the 25th International Conference Companion on World Wide Web*. International World Wide Web Conferences Steering Committee, 581–584.
- [49] Rossano Schifanella, Miriam Redi, and Luca Maria Aiello. 2015. An Image Is Worth More than a Thousand Favorites: Surfacing the Hidden Beauty of Flickr Pictures. In *Ninth International AAAI Conference on Web and Social Media*.
- [50] Joseph Schumpeter and Ursula Backhaus. 2003. The theory of economic development. *Joseph Alois Schumpeter* (2003), 61–116.
- [51] Roberta Sinatra, Dashun Wang, Pierre Deville, Chaoming Song, and Albert-László Barabási. 2016. Quantifying the evolution of individual scientific impact. *Science* 354, 6312 (2016), aaf5239.
- [52] Alex Smola, Le Song, and Choon Hui Teo. 2009. Relative novelty detection. In *Artificial Intelligence and Statistics*. 536–543.
- [53] Sameet Sreenivasan. 2013. Quantitative analysis of the evolution of novelty in cinema through crowdsourced keywords. *Scientific reports* 3 (2013).
- [54] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.
- [55] N. N. Čencov. 1982. Statistical decision rules and optimal inference. *American Mathematical Society* 53 (1982).
- [56] Brian Uzzi, Satyam Mukherjee, Michael Stringer, and Ben Jones. 2013. Atypical combinations and scientific impact. *Science* 342, 6157 (2013), 468–472.
- [57] Johannes Wachs, Aniko Hannak, András Vörös, and Balint Daroczy. 2017. Why Do Men Get More Attention? Exploring Factors Behind Success in an Online Design Community. In *Eleventh International AAAI Conference on Web and Social Media*.
- [58] M.L. Weitzman. 1998. Recombinant growth. *Quarterly Journal of Economics* 113, 2 (1998), 331–360.
- [59] N.M. Wijnberg and G. Gemser. 2000. Adding Value to Innovation: Impressionism and the Transformation of the Selection System in Visual Arts. *Organization Science* 11, 3 (2000), 323–329.
- [60] Michael J Wilber, Chen Fang, Hailin Jin, Aaron Hertzmann, John Collomosse, and Serge Belongie. 2017. BAM! The Behance Artistic Media Dataset for Recognition Beyond Photography. *arXiv preprint arXiv:1704.08614* (2017).