Mind Soc (2013) 12:137–149
DOI 10.1007/s11299-013-0113-9

Flow of innovation in deviantArt: following artists on an online social network site

Alkim Almila Akdag Salah · Albert Ali Salah

Received: 1 July 2012 / Accepted: 6 December 2012 / Published online: 22 February 2013

Abstract Computer and communication technologies created new modes of creating and sharing arts. In this paper, we apply 'diffusion of innovation' theory to investigate how artistic content travels in an online social network site called deviantArt, a site designed for sharing user-generated artworks. We first define what innovation corresponds to in such a context, and then discuss how it can be measured with the help of network, image and text analysis methods. We propose to use user-shared resources as relatively easy targets of tracking innovation.

Keywords Art market · Social network sites · Complex networks · Image analysis · Text analysis

1 Introduction

Rogers's (1995) theory of diffusion of innovations started an avalanche of studies about how recognizable patterns in social systems travel. Especially, and most strikingly, the processes of diffusion for novel ideas, for new products, and the results of medical research have been scrutinized in early 1960s–1970s with

A. A. Akdag Salah (✉)
New Media Studies, UvA, Turfdraagsterpad 9, 1012 XT Amsterdam, The Netherlands
e-mail: a.a.akdag@uva.nl; almila@sabanciuniv.edu.tr
URL: http://home.medewerker.uva.nl/a.a.akdag/

A. A. Akdag Salah
Visual Art and Communication Design, Sabancı University, Orhanlı, Tuzla, 34956 Istanbul, Turkey

A. A. Salah
Department of Computer Engineering, Boğaziçi University, Bebek, 34342 Istanbul, Turkey
e-mail: salah@boun.edu.tr
URL: http://www.cmpe.boun.edu.tr/~salah

different techniques (Rogers 1976; Rogers and Agarwala-Rogers 1976). The concept of diffusion became even more important when it was applied to organizational management problems, i.e. how knowledge should be spread out properly in an organization to keep it dynamic, and to advertising campaigns, i.e. how to persuade users to buy new products.

In this paper, we apply some of the basic ideas of the theory of diffusion of innovation to a new setting, i.e. how innovation travels among a community of artists. In this setting, at least two major aspects differ from previous studies of innovation. First, what is being diffused in the artistic context almost never ends up replacing all its alternatives on a grand scale; even the most successful diffusion remains bounded in its effects, and the adoption of innovation is almost never black and white. The second major difference is that in the spread of innovative products and ideas, the population (broadly speaking) may show skepticism and resistance, whereas artistic communities tend to embrace innovation and novel ideas and to share information on old and new techniques, styles and genres more freely.

If we take a look at the history of Art Academies and how before them art education was handed down from one generation to the next from masters to students, we see that the spirit of sharing has not changed at all. What has changed since the introduction of the Internet is the speed and means of sharing information. Thanks to new tools and theories available to us today, we can follow the process of how artistic content travels, for instance in online social network sites, which build communities around ideas and form some of the most important information sharing platforms of our time (Boyd and Ellison 2007; Boyd 2007; Lewis et al. 2008; Mayer and Puller 2008). Our aim in this paper is to discuss the conceptual and technical tools that are needed to define innovative artistic content in such a social network site (SNS), and to seek quantitative measures for the specification of this content.

This paper is structured as follows: In the next section, we describe our setting, namely deviantArt, our choice of SNS for this study, followed by our research questions and methodology. We exemplify how image, text and network analysis illuminate our case study, and conclude with a brief discussion.

2 Setting

2.1 The deviantArt platform

The biggest art platforms of today are virtual platforms; the Internet has enabled millions of people to collectively pursue artistic concerns, and online communities have emerged to sustain the creation and consumption of arts. The largest of these online communities is *deviantArt* (dA),[1] with close to 22 million members (as of June 2012), an image archive of about 224 million works, and a number of daily visitors that vastly exceeds any renowned museum (Akdag Salah 2010a).

[1] SNS are defined as web-based services for users to construct a public or private profile and to connect with other users in a bounded system (Boyd and Ellison 2007). In this sense, dA resembles SNS as it offers basic services to its users, and it creates a community structure. However, dA works like a blog-sphere as well (see Adar et al. 2004 on blogging), as each dA member is bestowed a website on his/her own.

As an *SNS*, dA offers many services to its members, and enables a platform for sharing not just artworks, but also ideas, techniques, tutorials, reusable images, etc. The social aspect of dA shows itself in various ways as well: beside the ability to ‘favor’ other members’ works, one can leave comments, write journal/news articles, organize contests, curate events and create groups on specific themes. But the main community structure of dA is based on its category system: each artwork that is uploaded into the dA has to be put under one, and only one, category.

A quick look at the categories reveals that the category structure is built predominantly on production techniques (like digital art, traditional art, photography, arts & crafts, film, animation, etc.) as opposed to genres, and quite readily, people using similar techniques form social clusters and communities (Akdag Salah et al. 2012). Nevertheless, the social structure of dA should not just be summarized by categories, as the members’ age, language and geographic location also play major roles in community formation (Liben-Nowell et al. 2005). Furthermore, the dA infrastructure has mechanisms to promote selected content (i.e. shared artworks) and to make these more visible to the members and to newcomers, as well as to keep the social life of the site alive. In the next section, we will discuss one of these features, one that has been a part of dA since its inception, and which we propose to use in this rich context to track down diffusion of artistic ideas.

2.2 Empirical data and daily deviations

One of the promotion mechanisms in dA is the idea of choosing about 20–30 works daily and publishing these on the homepage of the site for the duration of a single day. The selected works illustrate a representative distribution of the dA community: they might belong to the earliest or latest members, to popular or unknown members (Akdag Salah 2010b). Here, a short explanation about the jargon of the site is in order: every member is called a ‘deviant’ and every uploaded work is a ‘deviation’. The selected deviations for each day are called ‘daily deviations’ (DD’s).

Our data stem from these DD’s and galleries of the corresponding deviants and were acquired directly from dA Headquarters. This data set is composed of deviants who have received at least one DD, and all members who are ‘watching’ (i.e. following, or subscribed to) these deviants, or have favored one of their works. One can conceptualize the latter as the ego-network of the deviant, making up its immediate social neighborhood.

As depicted in Fig. 1, every deviant has a ‘home page’ with demographics such as age, gender, location (this information is only available if the user chooses to share it), statistics such as number of artworks, page views, comments, and an interface that leads to the deviant’s gallery, his/her favorites, journal, friends/watchers, etc. Furthermore, each artwork produced by the deviant has its own webpage containing tracked statistics (number of pageviews, downloads, comments).

The DDs data set consists of 30,643 daily deviations created by 21,745 deviants between 2000 and 2011. These deviants had a total of 1,321,264 artworks in the dA archive at the time of data retrieval. Through the individual artwork links, it is

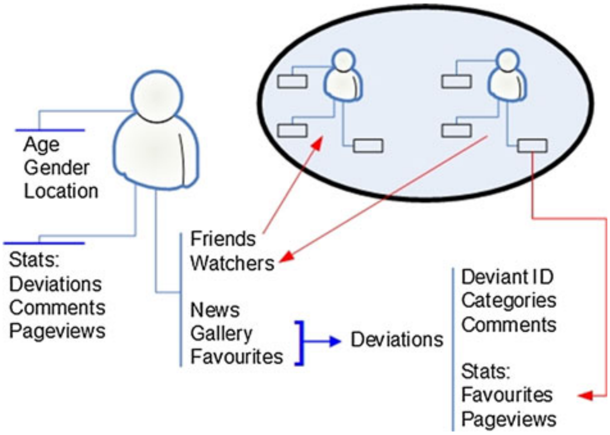

Fig. 1 Diagram depicting the structure of deviantArt

possible to use simple computer scripts to retrieve all comments associated with a particular work, complete with the date of commenting and identity of the commenter.

Given all these data items, our aim is to formulate a series of questions to understand the flow of innovation through this huge artistic community. In the next section, we describe how the diffusion of innovations theory can serve as a framework, and how its possible concerns can be translated into quantitative programs in a way that was impossible for any arts establishment before dA.

3 Research questions

In his classical work, *Diffusion of Innovations*, Rogers (1995:11) defines innovation as “an idea, practice, or object that is perceived as new by an individual or other unit of adoption”. Simple dictionary definitions of *innovation* are (1) the introduction of something new, and (2) a new idea, method or device. However, in the context of a community of artists, the definition of innovation is not so straightforward at all. Each artwork is by definition something new, in the sense that it shows skill and imagination and has novelty beyond what exists up to its creation. In reality, artists are always influenced by other artists, artworks and ideas when creating new works, and it is sometimes possible to trace these influences. Our first question, then, is whether we can use the wealth of data available on deviantArt to trace influences systematically via automated computer analysis.

The second definition of innovation implicates technique, and each new technique can be seen as an innovation. If we look at the history of art, we see that new techniques play important roles in stylistic changes and form foundations for the birth of new genres. Art historians often singled out one artist as the person who initiates the break from an older tradition (be it Giotto, Cezanne or Warhol), and they identified the major innovation as the introduction (or re-introduction) of a technique. Even though it's common to have multiple artists apply and develop similar techniques, usually only one artist (or at most a few) will be associated with it. Subsequently, our second question is whether we can identify the propagation of style and technique across the social network of artists.

In traditional art historical analysis, the 'diffusion of innovation' is pursued through many channels: historical information about an artwork, such as who has commissioned it, where it has been showcased, which gallery owned it, through which dealer houses it traveled. These are details that become important, while each tidbit of private information about the artist itself is evaluated as well: the artist's social circle, friends, masters, colleagues, and historically known sources of inspiration are vital for tracing from where the artist might have received her or his muses. The dA archive offers an incredibly rich reservoir for seeking these answers systematically via computer analysis: It has a fully dynamic temporal aspect, where communications between social contacts are stored with time stamps and are open for text analysis. It also has a hyperlinked structure, and it is common for artists to describe the tools and techniques they use, as well as to provide links for their sources and to their own works. Similarly, 'watchers' of a work leave comments and frequently provide links if they have seen similar work elsewhere or if they have used some of the ideas themselves.

4 Methodology

According to Rogers (1995), "diffusion is the process by which (1) an *innovation* (2) is *communicated* through certain *channels* (3) over *time* (4) among the members of a *social system*." In our study, these four concepts each have their counterparts. Starting from the end, dA is easily characterized as a social system as it has its own rules, etiquette, social hierarchy, lingo, and promotion mechanisms that are recreated daily by interacting members. The time aspect, just as in any SNS, pertains to the actions of the members, like uploading an artwork, commenting on an existing artwork, favoring an artist or an artwork, etc. The comments associated with artworks follow a temporal succession and read like a multilogue. The communication channel is also not a singular medium but rather is comprised of a number of parallel channels that are intertwined, primarily one-to-many mode. The communication occurs by: sharing an artwork, leaving explicit comments under a work, favoring a work, preparing a journal or text entry, or by making explicit instructions for the usage of a technique in the form of a tutorial.

The innovation, the most important and the most elusive concept in the definition, covers a broad spectrum as expressed in the previous section. In this paper, we will focus on a very specific type of innovative content that represents an artist's explicit attempt to diffuse a certain technique that he or she mastered. These are the materials published under the 'resources' category.

As mentioned earlier, the dA website has a category structure that requires its members to assign uploaded artworks into specific categories that are organized in a hierarchical manner. The 'resources' category is among the top categories and is devoted to materials that are essentially used for sharing information, resources, images, and ideas. The 'deviations' put under this category are open-source and free for everyone to download and re-use. The etiquette of the website dictates that one leave a comment under the material one re-uses, and that the re-user link the original material with the resulting deviation. It is also possible to follow these links to collect artworks that borrow themes, features, techniques, or image segments from other artworks.

The main subcategory of the 'resources' is 'stock images', which are images that are prepared for other artists to use as they see fit (see Fig. 2). Other major subcategories include 'tutorials', which contain instructions that are usually described over an example illustrating how to apply a tool or a technique. If it is possible to use computer-based analysis tools to examine the work put under 'resources', as well as the propagation of the work in the network, and how this relates to the social structure of the network, it may be possible to harvest informative cases that will help us in understanding what makes an artistic innovation successful.

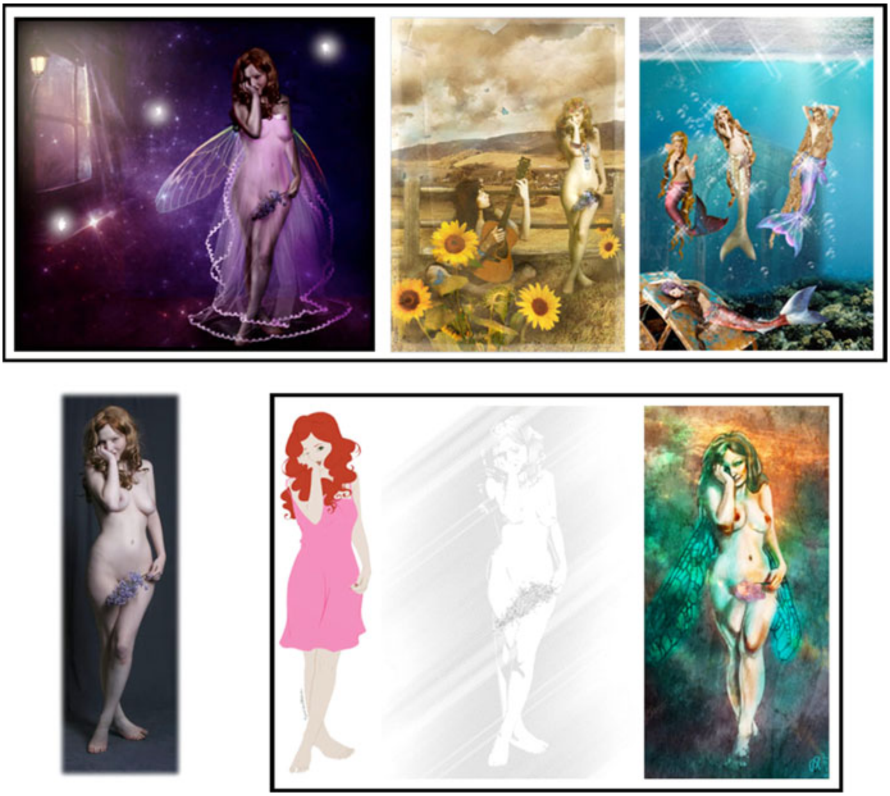

Fig. 2 An example of a stock image (*left*, *bottom*), and images by different artists that use this stock

5 Analysis tools

We next assess existing text, network and image analysis tools briefly to see how resource reuse can be traced through the dA social system.

5.1 Image analysis

The natural way of tracing ideas in visual arts is by visual inspection. In the example given in Fig. 2, even a non-expert human can detect the transfer of visual material from the source (which is a stock image) to the targets, as realized by different individual artists. In the dA archive, we have tens of thousands of resources of this type, and hundreds of millions of images to seek traces of these resources. Nonetheless, it is not the size of the image comparison task that daunts the computers. Ten years ago, a large image dataset meant a hundred thousand images (Smeulders et al. 2000). Today, we have the means of processing hundreds of million images through distributed computing and GPU-powered computers.

The current state-of-the-art in content-based image analysis relies on powerful and generic feature extraction methods that map the image, which is usually a high-dimensional object, to a low-dimensional feature space, thereby enabling fast comparison of many images in short time. An image is a complex entity of many properties (e.g. colors, contours, texture, composition, content), and reusing an image may imply transfer along one or more of these dimensions. To a computer, some of these are easy to compute and compare (like the distribution of colors), but others are extremely difficult. Semantic concepts, for instance, are at the moment mostly beyond the capabilities of computers. Current concept detection methods distinguish small sets of concepts (20 classes in the recent Pascal Visual Object Classes challenge[2]) after extensive training on hand-labeled samples, and special classification methods that are designed to separate a particular set of concepts from each other. These approaches do not readily generalize to recognizing the existence of concepts in arbitrary images. While it is possible to use ontologies and networks of meaningfully related words and concepts (like WordNet[3]) to enhance image analysis by linking related concept detectors to each other (Snoek et al. 2007), the problem of finding arbitrary connections between images is not solved.

To give an illustrative example, we take the stock image case given in Fig. 2. We use a state-of-the-art feature transformation (called Histogram of Oriented Gradients, or simply, HOG), described by Dalal and Triggs (2005). We extract these features globally from each image, and compare the stock image to all the images in the galleries of the few deviants that have actually used this stock image. Figure 3 shows a mapping in 2D where the distances between these images are preserved, and a subset of images is shown as thumbnails. The stock image itself is shown in the near-center with a circle around it. The images that actually use the stock are also circled and connected to the original stock image. When viewed in color, the frames of images as well as the circle indicators are depicted with

[2] http://pascallin.ecs.soton.ac.uk/challenges/VOC/voc2012/index.html.

[3] http://wordnet.princeton.edu/.

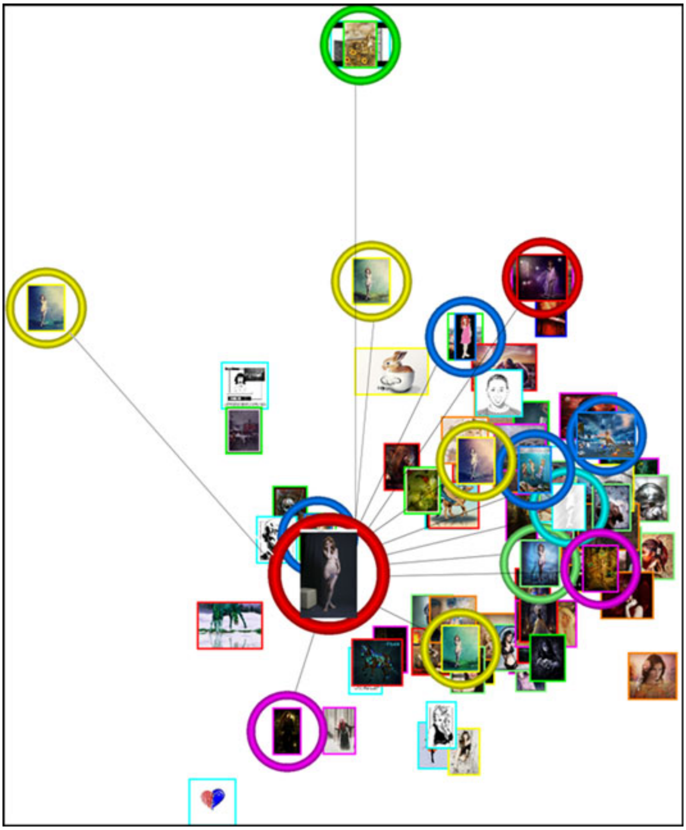

Fig. 3 Images mapped to a 2D feature space. The stock image and images that use the stock are marked and linked

different colors for each artist. Even with this incredibly small number of images, we can observe that neither the closest images to the original stock are derived from it, nor are all the derived images are assigned to the stock's immediate neighborhood.

Given a certain stock image and a set of images that either use or do not use the stock, it may be possible to create a complex feature transform that will indeed separate the derivative images. Buter et al. (2011) have previously developed a tool for processing sets of images to automatically find feature transforms that would achieve such a separation. However, as mentioned before, such a system requires discriminative training, and works with small sets of samples. What we show here is that generic and fast approaches to the problem are completely unusable at the moment for automatically locating an image that uses concepts from another. Image analysis should either be used in broad and informative visualizations of large scale datasets to provide the experts with insights into the data (see (Manovich 2008) for the cultural analytics paradigm that follows this route), or should be combined with network and text analysis that would greatly reduce the amount of visual data to enable more complex image-based inferences.

5.2 Network analysis

The social structure of dA can be represented as a network, where each of the 22 million deviants is a node. The arcs that connect these nodes can be chosen to represent 'watching', 'favoring', or 'commenting' relations. Through network analysis, we can describe the ego-network of particular deviants, define visibility and social attention measures, explore community structures, and evaluate the formation and propagation of new social connections. We can look at connections over the network to see which factors predict increased diffusion probability between individuals.

We have previously described the 'daily deviation' (DD) promotion mechanism. Focusing only on deviants who have received DDs allows us to reduce the data to a manageable amount. We let the arcs in the network represent a deviant 'watching' another deviant. In theory, there is no tangible connection between members who receive a daily deviation promotion,[4] but still the deviants with DDs and all their watchers make up a densely connected network. We further process this network and extract a Resources network of deviants who received DDs with deviations that fall under the 'resources' category. The Resources network is even more densely connected than the DD-watchers network, with 17,015 nodes and more than 1.5 million arcs. For visualization, we pruned the Resources network by snowball sampling to 1,000 nodes (Lee et al. 2006).

Figure 4 depicts this network. In it we identify four main clusters, primarily determined by production technique: Photography, Manga & Anime, Digital Art, and Customization. In our previous study, we found out that Photography was strongly connected to Resources and had a high degree of within-cluster links, whereas subcategories from other categories were intermingled. Here, the strong positions of Digital Art and Manga & Anime are interesting and ask for a further investigation. These two categories are the ones that have the most to gain from Resources in terms of 'adopting' an innovation. In Digital Art, as part of photo-manipulation, stock images are always re-used, and we define 'stock images' as tangible influences. In Manga & Anime, what is adopted is probably the technique rather than the material, since this style asks for stricter technique and less flexibility, especially for beginners.

5.3 Text analysis

Text analysis on the individual page of the artist can provide us with cues for establishing a 'social profile', including demographics such as sex, age and geographical location, as well as artistic cues including production genre/categories, followed genres, and software tool usage. The social profile thus created can serve to determine similarity between individuals, and can be employed to analyze the

[4] See Akdag Salah (2010b) for a description on how the promotion mechanism works.

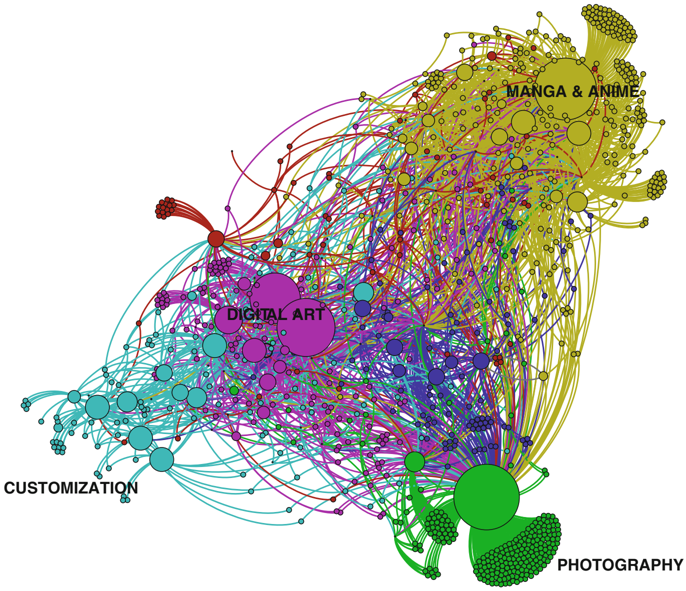

Fig. 4 Resources network, reduced to 1,000 nodes and 2,759 edges through snowball sampling

effect of *homophily* in flow of ideas, which is the principle that similar individuals act alike (Yuan and Gay 2006). More elaborate models can look at all the comments produced by the artist and even seek to determine personality traits by word choice (Minamikawa and Yokoyama 2011).

The artworks in dA are devoted their own webpage where the artist can add various details about his or her creation, including production technique, sources of inspiration, links to related material, etc. Also on these pages, we find comments left by watchers of the artwork and responses of the artist to these comments. For a given artwork, text analysis of this webpage can provide us with a wealth of information, including the genre and category of the work (as determined by the artist). Determining the categories in which the artist produces artworks can help us look for community structures when combined with network analysis.

What is more directly relevant to the analysis of the flow of innovations is the links to other artworks left among the comments. When we apply text analysis to comments retrieved (i.e. *crawled*) from the individual webpages of resources, we can extract links to material generated with the resource since the community etiquette asks for such feedback. We have done this for a small set of popular tutorials and looked at the distribution of reused material as a proof of concept. Figure 5 depicts the (average) number of reuse a tutorial receives as a function of days after its publication. We see a trend that closely resembles the S-shaped curve observed by Rogers in the adaptation of innovations, where the reuse starts slowly (early adopters), has a relatively sharp burst when the social system effectively disseminates the tutorial (early and late majority), and then a leveling off as a certain kind of saturation is reached (laggards).

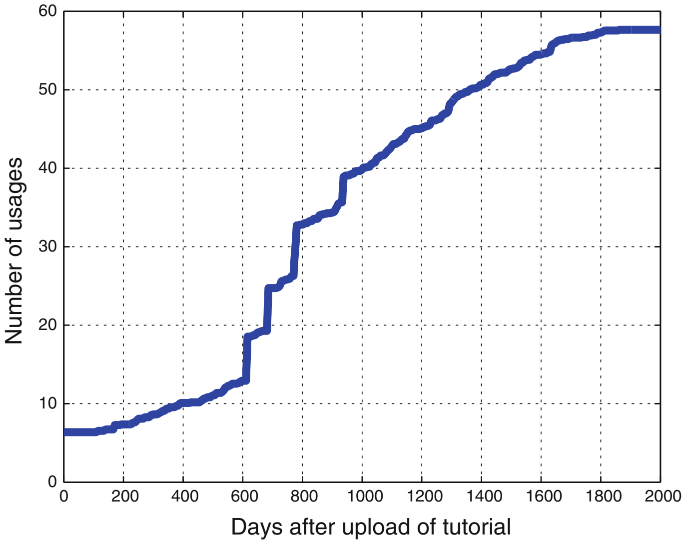

Fig. 5 The average number of usages for popular tutorials as a function of days after the upload of the resource

We are careful not to read too much into this pattern, and apart from the small sample size used to generate it, we would like to point out to an important difference with the pattern proposed by Rogers: In the classical diffusion of innovations graph, the y-axis denotes the population ratio of adaptation for an innovation, and saturation occurs when almost the entire population has adopted the innovation. In our case, the most popular tutorial is perused by hundreds of people, but this cannot be compared to the entire dA population, which is counted by tens of millions.

6 Discussion

In order to trace influences and innovations in an art community, the most important data naturally come from the analysis of the artworks themselves. Usually, historical information and close scrutiny of technique, composition and content are required to establish influences between artworks. However, the sheer size of the dA archive makes it impossible to analyze the artworks by hand in order to detect which artist is influenced by whom. Even with automatic image analysis tools that can process millions of images based on specific features, the influences can exist in so many different semantic levels that it is unreasonable to expect meaningful results. Where individual analysis approaches fail, statistical approaches and massive visualization can provide clues.

In the case of dA, the temporal dimension becomes tangible through processing of the rich information buried in the dA archive. With text and network analysis approaches, we can trace the evolution and propagation of certain ideas over the dA network in time, and characterize the dynamics of diffusion. The channels of communication are laid bare in dA, as messaging, favoring and commenting all occur transparently. In short, visual data contained in images are not easy to analyze, as opposed to the information that can be extracted from links or texts found in the dA website.

There are many questions that can benefit from the proposed approach. We can look for patterns in adoption of tools and styles, as well as for valorization of artists. One of the important steps in this kind of exploration is to find people that are influential. Rogers (1995) defines "opinion leadership" as "the degree to which an individual is able to influence other individuals' attitudes or overt behavior informally in a desired way with relative frequency". There is a strong relationship between opinion leaders' openness to innovations and the degree of innovativeness and rate of innovation adoption of the social group (Valente 1996; Valente and Davis 1999). Of course, one hypothesis for us to test is whether there are opinion leaders in dA and whether they can be automatically located. We can go one step ahead, and try to find artists that are entry points to certain styles. Having access to exact timing of actions over the network allows us to look for statistical traces about people who facilitate the entry to particular styles and genres.

The theory of diffusion of innovations is not the only framework for understanding how new ideas propagate in a social network. The theory of *informational cascades* (Bikhchandani et al. 1992), for instance, seeks to explain the frequently encountered case of people changing their behavior based on seemingly small perturbations. The idea is that individuals commonly conform to decisions made by people before them, leading to the situation that a few people play decisive roles on the behavior of a large number of people. This is also called *localized conformity* in the literature. In dA, as is typical for SNS-type media in general, the amount of people promoting an artist is directly proportional to the attention the artist is receiving, leading to power-law distributions: there are a few artists with an immense number of followers, and an immense number of artists with very few followers. The former group of artists are given a "Deviousness Award", displayed on the main page of the artist, which further corroborates the number of followers for the artist, and increases his or her influence. While it is straightforward to identify artists with Deviousness Awards automatically through text analysis, the 'locality' that is implied in localized conformity does not just pertain to a geographical location but to social communities of the SNS, which can only be analysed through advanced network analysis tools.

Acknowledgments This research is supported by Bogazici University BAP-6531 project and a Nederlandse Organisatie voor Wetenschappelijk Onderzoek NWO-VENI grant.

References

Adar E, Zhang L, Adamic LA, Lukose RM (2004) Implicit structure and the dynamics of blogspace. In: 13th international WWW conference New York: Workshop on the Weblogging Ecosystem

Akdag Salah AA (2010a) The online potential of art creation and dissemination: deviantArt as the next art venue. EVA2010. Computer Arts Society, London

Akdag Salah AA (2010b) Performing curatorial practices in a social network site: The curators of deviantArt. CHArt Conference, London

Akdag Salah A, Salah AA, Buter B, Dijkshoorn N, Modolo D, Nguyen Q, van Noort S, van de Poel B (2012) DeviantArt in spotlight: a network of artists. Leonardo 45(5):486–487

Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. J Pol Econ 100:992–1026

Boyd DM (2007) Why youth (heart) social network sites: the role of networked publics in teenage social life. In: Buckingham D (ed) Youth, identity, and digital media. MIT Press, Cambridge, pp 119–142

Boyd DM, Ellison NB (2007) Social network sites: definition, history, and scholarship. J Comput Mediat Commun 13(1), article 11

Buter B, Dijkshoorn N, Modolo D, Nguyen Q, van Noort S, van de Poel B, Akdag Salah AA, Salah AA (2011) Explorative visualization and analysis of a social network for arts: The case of deviantArt. J Converg 2(2):87–94

Dalal N, Triggs B (2005) Histograms of oriented gradients for human detection. IEEE CVPR WKSHP 886–893

Lee SH, Kim P-J, Jeong H (2006) Statistical properties of sampled networks. PHYS REV E 73

Lewis K, Kaufmann J, Gonzalez M, Wimmer A, Christakis N (2008) Tastes, ties, and time: a new social network dataset using facebook.com. Soc Netw 30:330–342

Liben-Nowell D, Novak J, Kumar R, Raghavan P, Tomkins A (2005) Geographic routing in social networks. PNAS 102(33):11623–11628

Manovich L (2008) Cultural analytics: Analysis and visualization of large cultural data sets. Screen, February: 1–23

Mayer A, Puller SL (2008) The old boy and girl network: social network formation on university campuses. JPBE 92:329–347

Minamikawa A, Yokoyama H (2011) Blog tells what kind of personality you have: egogram estimation from japanese weblog. In: Proceedings of the ACM conference on computer supported cooperative work, pp 217–220

Rogers EM (1976) New product adoption and diffusion. J Consum Res March:290–301

Rogers EM (1995) Diffusion of innovations, 4th edn. Free Press, New York

Rogers EM, Agarwala-Rogers R (1976) Communication in organizations. Free Press, New York

Smeulders AWM, Worring M, Santini S, Gupta A, Jain R (2000) Content-based image retrieval at the end of the early years. IEEE T Pattern Anal 22(12):1349–1380

Snoek CGM, Huurnink B, Hollink L, de Rijke M, Schreiber G, Worring M (2007) Adding semantics to detectors for video retrieval. IEEE T Multimedia 9(5):975–986

Valente TW (1996) Social network thresholds in the diffusion of innovations. Soc Netw. doi:101016/0378-8733(95)00256-1

Valente TW, Davis RL (1999) Accelerating the diffusion of innovations using opinion leaders. Ann Am Acad 566:55–67

Yuan YC, Gay G (2006) Homophily of network ties and bonding and bridging social capital in computer-mediated distributed teams. J Comput Mediat Commun 11(4):1062–1084