Available online at www.sciencedirect.com

SciVerse ScienceDirect

Cognitive Systems Research 21 (2013) 65–74

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Stigmergic dimensions of online creative interaction

Action editors: Margery Doyle and Leslie Marsh

Jimmy Secretan

Korrelate, Inc., 4700 Millenia Blvd., Suite 320, Orlando, FL 32839, United States

Available online 5 July 2012

Abstract

This paper examines the stigmergic dimensions of online interactive creativity through the lens of Picbreeder. Picbreeder is a web-based system for collaborative interactive evolution of images. The Picbreeder applet starts by randomly generating several images, which are then mated and mutated based on the user's selections. The user can then publish the image to the Picbreeder website where other users can download and continue the image's evolution. Within this process, users collaboratively create images with significant complexity, all without explicit communication. In short, Picbreeder encourages a new form of stigmergic collaborative creation. The most surprising result of the Picbreeder experiment during more than 3 years of operation has been the quality of the resulting images, despite the limited ways of interacting with other users. This fact challenges some commonly held notions of creativity, both online and offline. While current cognitive research in creativity places significant emphasis of the personal traits and cognitive structures that give rise to creative thought, Picbreeder highlights the potential for the emergence of creativity through stigmergic interaction. Picbreeder offers a rich data set for analysis of collaborative interaction with over 155,000 inputs from hundreds of users combined to create over 7500 images. It is hoped that the insights offered in this paper will influence both the understanding of collaborative creativity and the development of new modes of online creative interaction.

Keywords: Stigmergy; Creativity; Interactive Evolutionary Computation; Collaborative interactive evolution; Memetic evolution

1. Introduction

The web has experienced a recent proliferation of design expert communities in domains from software engineering (e.g. Sourceforge and Github) to art (DeviantArt and others). These communities have become hotbeds of creative interaction, with users posting their projects, closely interacting on new endeavors, and engaging in spirited discussion about their craft. With users in these communities constantly generating out new software, images, music and any other artifact imaginable, it is hard to deny that there is significant creative interaction happening. Members of these communities often possess widely varying degrees of proficiency, but more often than not, they have some baseline amount of talent that allows them to enter the community.

Enter Picbreeder. Picbreeder is a web-based system for collaborative interactive evolution of images. The Picbreeder applet starts by randomly generating several images, which are then mated and mutated based on the user's selections. The user can then publish the image to the Picbreeder website where other users can download and continue the image's evolution. Within Picbreeder, one need not have artistic talent to contribute to the community, although good taste typically helps. As in more traditional design, new innovations are typically small modifications to the existing structure, which can change the design incrementally or effect a larger shift. Even though users followed their individual interests when evolving this phylogeny, new interesting directions emerged. Many users contributed repeatedly to an evolving lineage, using the design itself to encourage and facilitate collaboration.

Successful collaborative design in Picbreeder does not require shared intentions, suggesting that effective collabo-

E-mail address: jimmy@korrelate.com

1389-0417/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.cogsys.2012.06.006

ration may be emergent rather than planned from the top down. The surprising result of this emergent process is the gradual discovery by untrained users of hidden treasures within a vast uncharted space. Picbreeder also serves as a fascinating, though initially unintentional, experiment in stigmergic creativity.

The concept of stigmergy was first introduced by Pierre-Paul Grassé, a zoologist, who used it to describe the activities of the termite mound. As he described it, "[s]tigmergy manifests itself in the termite mound by the fact that the individual labour of each construction worker stimulates and guides the work of its neighbour" (Grassé, 1982). The concept of stigmergy can be extended to human endeavors if one expands the notion of the mound to human venues, and replaces "construction worker" with any type of worker. If such an extension is permitted to human creative communities, this description becomes even more apt. Part of the excitement inherent in creative pursuits, whether it is visual art, music or creating open source software, is the moment when the work of a colleague "stimulates and guides" ones own work. Theraulaz and Bonabeau (1999) add that "[in] an insect society individuals work as if they were alone while their collective activities appear to be coordinated." This description too can apply to creative communities. Gabora (1997) points out that "[s]tudies on creativity... have focused on the *individual*, obscuring the fact that creativity is a collective affair. The ideas and inventions an individual produces build on the ideas of others (the ratchet effect)." It is very easy to focus on individual creative luminaries, while forgetting the environment and social milieu that are a large part of their creative interaction.

The results of Picbreeder not only demonstrate the truth of creativity as collaboration, but that a large component of creativity can be stigmergic. By abstracting out almost all direct communication and collaboration, and allowing users to be stimulated only by their work and the work of others, Picbreeder demonstrates the extent to which stigmergic processes can yield astounding results. This paper expounds on this point by first describing in detail what Picbreeder is and how it works (Section 2). Next, the paper casts creativity in general and Picbreeder specifically into the context of memetic evolution, a model of how ideas spread, change, evolve and die out (Section 3). The point is then made in Section 4 that these collaborative creative environments draw a great deal of their effectiveness from stigmergic interaction facilitated through creative artifacts. In Section 5, an analysis of the Picbreeder data is described that shows, despite the fact that Picbreeder users engage in almost no direct communication, it shares numerous properties with other collaborative creative environments. Finally, some conclusions and recommendations are made in Section 6.

2. Picbreeder

Picbreeder, at its core, is based on the principle of Interactive Evolutionary Computation (IEC). IEC (Takagi, 2001) describes the process of artificial evolution guided by human direction, pioneed by Dawkins in his book, *The Blind Watchmaker* (Dawkins, 1986). IEC applications generate an initial population of individuals, from which the user selects those that are most appealing. The selected individuals then become the parents of the next generation. As this process iterates, the population evolves to satisfy the user's preferences. Fig. 1 shows what a user session with the Picbreeder IEC client looks like. IEC is well-suited to domains in which success and failure are subjective and difficult to formalize. For example, traditional evolutionary algorithms would struggle to determine whether an image is "attractive" or not, yet humans can easily perform such evaluations. IEC can thus generate a variety of digital artifacts including images (Hart, 2007; Sims, 1993, 1997; Unemi, 1999), music (Biles, 1994; Collins, 2002; Hoover & Stanley, 2009; Johanson & Poli, 1998; Nelson, 1993), three dimensional models (Husbands, Germy, McIlhagga, & Ives 1996; Nishino, Takagi, Cho, & Utsumiya 2001), movies (Unemi, 1999) particle systems (Hastings, Guha, & Stanley, 2007), dancing avatars (Balogh, Dubbin, Do, & Stanley, 2007), and collages (Unemi, 1999), to name a few. It is worthwhile to note that Picbreeder's operation is complex and has numerous facets. A complete description of Picbreeder and how it operates can be found in Secretan et al. (2011).

Picbreeder, like other collaborative interactive evolution (CIE) systems before it (Anonymous (2006–2007); Langdon, 2005; Sims, 1993, 1997) attempts to improve IEC by involving multiple users in the evolutionary process. These systems ultimately aim to increase the variety and quality of solutions that can be evolved. But the user preferences and goals that drive the output of these systems are often in conflict. To solve this difficulty, Picbreeder allows users to collaborate in a unique way. Users can start with images in the system they admire or find interesting, branch a copy, and continue the evolution in a direction of their choosing. Note that, ultimately, whether the resulting images are "art" or not, is a deeply philosophical question, which is not addressed in this paper. The question of whether or not this is a human creative process is also deferred to more philosophical works. With respect to the evolutionary process, all that matters is if users respond positively (e.g. users give the image a high rating, or branch off their own images).

Picbreeder users can begin evolving in one of two ways: In the traditional option, users start from a random population of images and select those that they like, which spawn a new generation. When the user is satisfied with an image, he or she *publishes* the image, making it visible to others. The key idea in Picbreeder is that other users can alternatively begin evolving from an *already* published image instead of from scratch by *branching* the image, thereby continuing its evolution. Fig. 2 shows the tree of progeny associated with a skull image evolved by a user. The children of the image represent sessions where other users, starting with the skull image as a seed, evolved their own images.

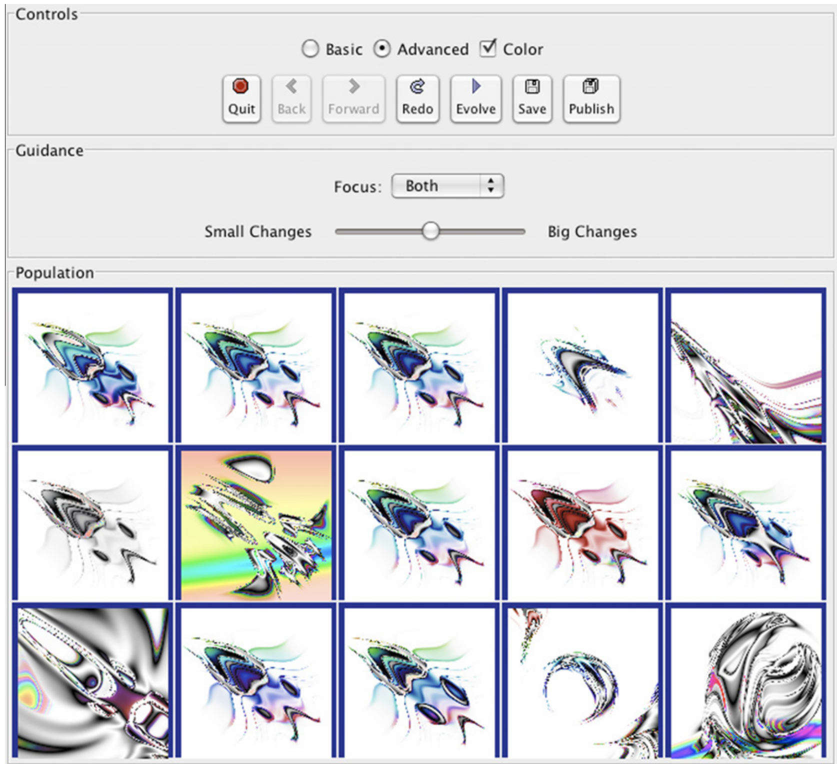

Fig. 1. The Picbreeder client program. The interface can spawn or re-spawn a new generation, move back and forth through the generations, and either publish the image or save it for later editing.

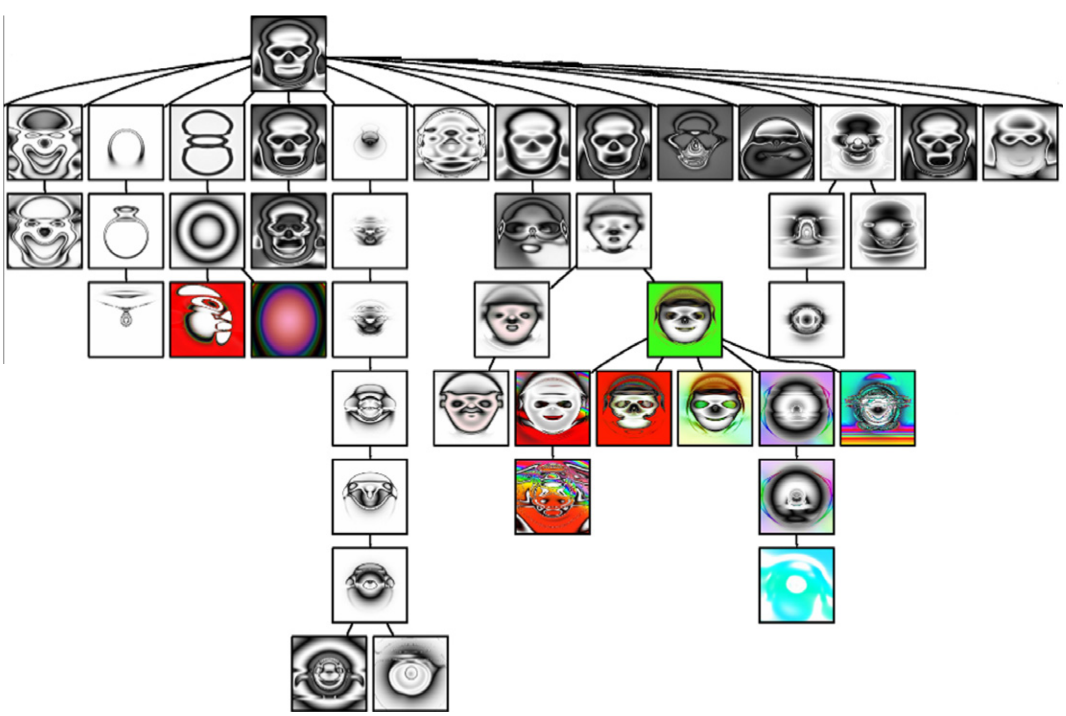

Fig. 2. Creative lineage. This view from Picbreeder demonstrates how collaborative evolution is analogous to collaborative creative processes in other fields, despite its stigmergic interaction. The parent image on the top is branched by several other users. The branching process copies the genetic code making up the image, and the branching users can take the image in whatever evolutionary direction they wish. As with concepts in other creative communities, the images can look closely related, or can quickly diverge.

The images in Picbreeder are drawn by *compositional pattern producing network* (CPPN) (Stanley, 2007), a structure with close kinship to neural networks. A CPPN is a function of n Cartesian dimensions that outputs a pattern in space. For example, a two-input CPPN produces a two-dimensional image. For each (x, y) coordinate in that space, its color is output by the CPPN that encodes the image.

Internally, a CPPN is represented as a connected-graph (i.e. a network) of functions chosen from a canonical set including sine, Gaussian and sigmoid. The structure and connection weights of the graph represents how the functions are composed to process each input coordinate. While the representation is intentionally biased toward exploiting the geometric regularities (e.g. symmetry) that commonly occur in images through its set of canonical functions (Stanley, 2007), it can represent any function (Cybenko, 1989), and therefore any image. In Picbreeder, three CPPN outputs represent colors in Hue, Saturation, Brightness (HSB) color space.

The evolutionary algorithm at the heart of Picbreeder's interactive process is called NeuroEvolution of Augmenting Topologies (NEATs). It addresses several fundamental challenges in evolving connected-graph structures like CPPNs (Stanley & Miikkulainen, 2002, 2004). NEAT starts evolution with a population of small, simple CPPNs and elaborates on them over generations by adding new nodes and connections. Each new component expands the search space, allowing evolution to gradually discover increasingly complex images.

Fig. 3 show qualitatively that human input clearly guides the evolutionary process. The four images at the top of the figure are part of a randomly generated initial generation. To the left are "randomly evolved" images. That is, each was evolved with random selections at each iteration. The number of generations to evolve the image was chosen according to a similar distribution as the images evolved in Picbreeder. Although these images exhibit more structure than their initial generation, they still do not resemble anything recognizable.

Because the CPPN-NEAT algorithm driving Picbreeder evolves structures similar to neural networks, instead of the images themselves, one can draw some interesting parallels between a single-user evolutionary session and individual creative thought. To begin with, many of the ideas from IEC are reminiscent of cognitive theories of creativity. Similar to the Geneplore model (Finke, Ward, & Smith, 1992), individual concepts (images in Picbreeder's case) are generated and then selected for further exploration. In addition, the process where individual users select interesting ideas to further modify and refine them, in an open-ended evolutionary process is similar to concepts that are described in memetic evolution, which is discussed in the next section.

3. Memetic evolution

The *meme*, as introduced in Richard Dawkin's book, *The Selfish Gene* (Dawkins, 1989) is a unit of cultural transmission. It is an idea that, much like the analogous gene for which it is named, can be replicated and mutated. There are many real world examples of this process. Bennett, Li, and Ma (2003) uses an evolution framework to study a corpus of 33 chain letters. The fact that it was easy to notice mutations in the transmission of the chain letter made them ideally suited for studying the memetic evolutionary process. Using bioinformatics techniques, the authors were able to create a phylogenetic tree of the letters. At each stage of replication, there was a chance for new ideas, paragraphs, phrases, and details to be introduced, making the final letters substantially different than the initial letters. Carrying this type of analysis over a longer term, Wood (2010) delineates the evolutionary journey of the folk song *John Barleycorn* over 250 years. The story, words and music to this song slowly changed over time, owing both to conscious changes and random errors through its oral transmission. For songs like this, is difficult to ascribe authorship because so many individuals may have had a hand in changing it. Also related is the evolution of language itself. Clearly, many useful and expressive languages have evolved, and continue to evolve to this day. Languages themselves can be studied from the perspective of information flow across complex networks (Gong & Wang, 2005), tracing the emergence of words, syntax and usage. Because of the availability of large user bases across the Internet, and crowd sourcing mechanisms like Amazon's Mechanical Turk, many of these same properties can be studied in online language evolution.

While a distinction is sometimes made between culture and creative artifacts, the distinction appears arbitrary. Most art museums now contain sections where one can find artifacts from prehistorical periods. These artifacts represent the terminal expression of a succession of constantly changing ideas commonly thought of as culture. The meme may be embodied in different ways, depending on the type of collaborative creative endeavor. Its evolution may proceed in different ways as well, and may involve a great deal of active discussion (e.g. concepts in string theory) or may not (e.g. folk songs).

Picbreeder serves as a vivid illustration of the process of memetic evolution. Changes to the CPPNs underlying the evolved images are reflected as structural changes to the child images. In some cases, the images change imperceptibly. In others significant structure is added or removed. Still, in some child images, some mutations cause the images to look completely different. All of these types of changes are visible in the phylogenetic tree shown in Fig. 2. The parent image of the skull is largely preserved in many of the child lineages, with small variations (e.g. changing the highlighting, increasing the size of the mouth, modifying the shape of the eye sockets, etc.). The samurai-like image of the far right of the image clearly shares many similarities with the parent image, but makes a significant change to the mouth portion of the structure. Some of the images are entirely different with some lineages devolving into noisy multicolored blobs and other lineages

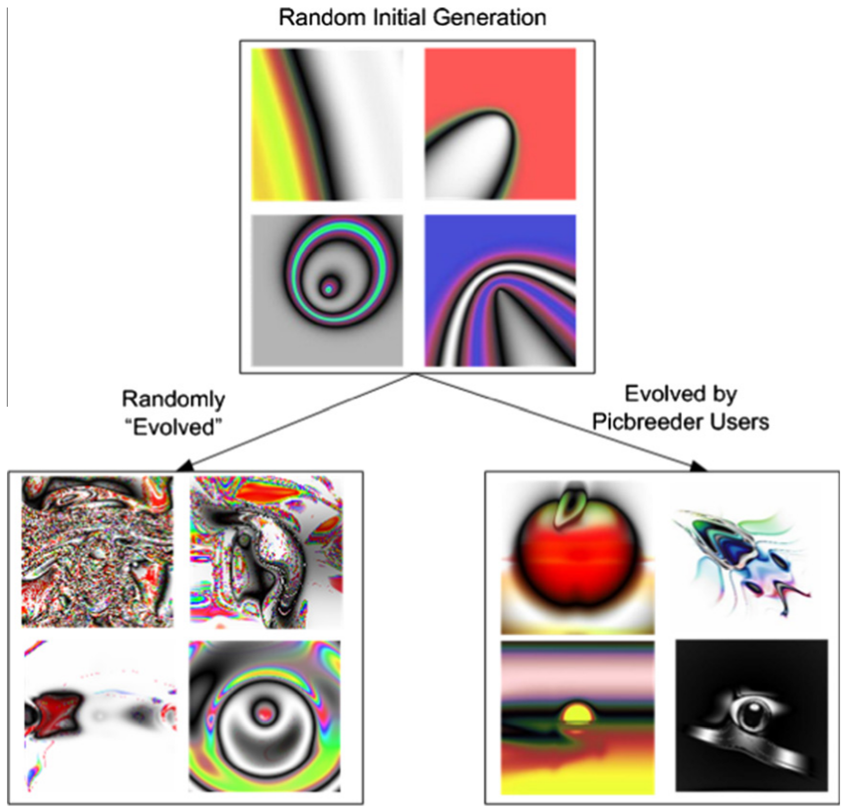

Fig. 3. Evolved versus Random. From a random initial generation, a comparison of randomly evolved images and those evolved by Picbreeder users.

becoming completely different faces. Picbreeder clearly shares much in common with other substrates for memetic evolution. Because of its online format, with almost no direct communication, Picbreeder provides a unique opportunity to isolate the portion of collaborative creativity that is emergent and not centrally planned.

4. Creativity as stigmergy

If stigmergy happens when an agent's effect on the environment "stimulates and guides" the work of others, then certainly creative communities must be subject to some kind of stigmergy. No creative endeavor exists in a vacuum, and being inspired and stimulated by the work of another is so fundamental to creative communities of artists, academics, engineers, etc., that it is difficult to imagine these communities functioning any other way.

Closely related to the concept of stigmergy is the concept of self-organization. The reason that it is remarkable that one user's work stimulates another's is the emergence of patterns that appear as if that they could be centrally controlled. Often, a mix of direct communication and control as well as emergent properties of the social structure give rise to collaborative creative activities. Fig. 4 suggests an informal ordering of the amount direct communication and coordination involved in several different types of creative processes, with emergent creative processes on the left end, and highly coordinated processes on the right.

The right end of the continuum begins with large-scale production of movies and video games. Each of these industries generate billions of dollars in revenue every year, and rely on many centralized, tightly coordinated projects. Movies often begin with a screenplay picked up by producers and directors. These lead figures dictate everything from the visual effects of the film to its production schedule. Many of these same constraints apply to the video game industry, which subsumes a large body of artists, developers, story writers and the like to create a final product that frequently follows a centralized artistic direction. Clearly while the process is tightly coordinated, almost all of the workers are indirectly influenced by art, music, code and other media that they have seen. In this sense, their activities may still be stigmergic as well, in that they are positively stimulated by the work of others either by artifact or in person.

Often less coordinated than large scale media projects, academic publishing and research draws from both emergent and direct means of collaborative creativity. Researchers, who frequently collaborate, are also influenced by the cumulative body of research papers and demonstrations, culled from hundreds or thousands of researchers whom they have never met.

Open source software engineering, often facilitated by sites such as Sourceforge and Github, are fascinating examples of non-traditional creative collaboration. These projects are placed in the middle of the proposed continuum

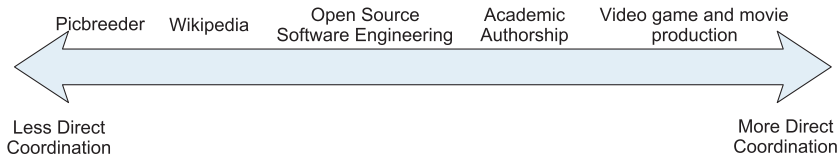

Fig. 4. Continuum of creative coordination.

because they frequently involve a balanced mix of indirect and directly coordinated interaction. Web systems like Sourceforge help to coordinate direct interaction and centralized planning between team members of small and moderately sized projects. The Linux kernel, on the other hand, has seen thousands of collaborators, many of whom are submitting patches of code being integrated by others with whom they have little if any direct contact (Kroah-Hartman, Corbet, & McPherson, 2009). Heylighen (2007)), expounds on this to describe the stigmergic elements of open source software design:

> *While people are of course much more intelligent than social insects and do communicate, open access development uses essentially the same stigmergic mechanism... any new or revised document or software component uploaded to the site of a community is immediately scrutinized by the members of the community that are interested to use it. When one of them discovers a shortcoming, such as a bug, error or lacking functionality, that member will be inclined to either solve the problem him/herself, or at least point it out to the rest of the community, where it may again entice someone else to take up the problem.*

At the level below open source software is the Wikipedia, which generally persists without central coordination, only general governance. By far the most frequent way of collaborating in Wikipedia involves users editing the same article with little or no direct communication (Voss, 2005). Wikipedia pages do have associated "talk" pages, which are host to discussions of procedure and collaborative rewriting, but this is the smaller part of their content. It should be noted that Wikipedia entails a convergence of fact and description which attempt to most accurately describe the subject in question. This is in contrast Picbreeder, which seeks to explore and expand the space of artifacts instead of converging on one.

Finally, in the continuum of creative coordination, Picbreeder is the example given with the lowest potential for centralized coordination and planning. While for most of the active life of Picbreeder, the site has hosted a forum for members to speak, interaction is almost exclusively through the branching mechanics of the site. While the forum in its lifetime has seen on the order of tens of posts, the site itself has seen thousands of branching actions. Picbreeder users seem to work "as if they were alone while their collective activities appear to be coordinated" (Theraulaz & Bonabeau, 1999), hence fitting into the commonly observed emergent self-organization of stigmergic processes. It is because of these unique dynamics that Picbreeder is well suited to study the specific contributions of stigmergy to creative collaboration.

5. Analysis of picbreeder interaction data

The Picbreeder experiment generated a huge amount of data about stigmergic creative interaction. The Picbreeder website opened to the public on August 1, 2007, and by August 1, 2010, it was visited over 85,000 times. The website catalogs over 7500 publicly available images resulting from over 155,000 user image selections. The images were contributed both anonymously and by over 500 registered users. Both branching statistics and 1–5 star ratings were collected for the images.

First, it is worthwhile to note that Picbreeder was a success over traditional, single-user IEC. While published images in Picbreeder are evolved for an average of 20.31 generations (sd = 29.94, median = 11) during a single user session, each published image has an average of 151.16 cumulative generations (sd = 108.77, median = 135) from an average of 9.75 ancestor images (sd = 9.18, median = 7). Users are often choosing to branch from existing images to create new images of increasing complexity and interest. The images generated by the process are also better, as judged by the community. Of the 763 images in the top-rated category (a special designation for images with high enough average ratings and a minimum number of votes), only 1.97% of them were evolved within the 20 generation limit typically reported for single users in Takagi (2001); the other 98.03% of the top-rated images in the system took many more cumulative generations to evolve (an average of 153.92). That means that, in all likelihood, if left entirely to a single-user IEC process, many of the community's favorite images simply would not exist.

Stigmergy, at its core, is a feedback process that determines how the work of one agent influences and stimulates the work of another. These stimuli to work happen over relatively short periods, but accumulate to a long term stigmergic effect (e.g. building a termite mound, authoring a Wikipedia article, or designing a piece of open source software). Therefore, if Picbreeder followed a stigmergic process, one would expect that on days where more users were participating, there would be greater productivity per user. These effects may not be long lasting, because the community may only retain interest in images while they remain "new." At some point, no matter how many images are generated, they will simply lose their novelty. In the short term, users would see the creation of more images which would in turn draw more interest. And indeed, Fig. 5 demonstrates just such a relationship. A histogram was generated for the amount of user productivity on days where different amounts of users were active. Values for number of users where there were not at least 5 representative days were filtered out because of too few observations. The trend line shows a small, but overall positive relationship (slope = 0.079) between number of users who published an image during a day and the productivity of all users for that day (adjusted R^2 equal -0.0368).

This stimulus to generate more images is just one of the positive feedback processes in Picbreeder, giving rise to stigmergic behavior. The community in Picbreeder is able to generate output that could be considered novel and appropriate in much the same way as other creative communities. Therefore, Picbreeder serves as an example of how creative collaboration can exist through almost entirely stigmergic means. In the remainder of this section, Picbreeder image publishing data is analyzed for common patterns found in other collaborative creative networks. Quantitative comparisons are made with academic paper authoring, because the bibliometric is freely available and richly analyzed.

Because of how these other creative communities work, it is difficult to isolate and study their emergent versus their centrally coordinated properties. It is hoped that by showing parallels between Picbreeder and these other creative communities, it will be possible to shed light on how stigmergy might influence these communities as well.

5.1. Properties of the collaboration graph

Because there is generally no direct communication among users in Picbreeder, users can only influence each other through their published images. To measure this interaction, an influence graph was built from Picbreeder user data. An undirected graph was generated with the edge between user *A* and user *B* representing the total number of interactions (i.e. branching actions) in either direction between user *A* and user *B*.

Fig. 6 is a visualization of the Picbreeder influence network. The network is strongly connected, with relatively few collaborators connected to many other users. This graph is closely related to academic co-authorship graphs. For instance, the co-authorship graph in Börner, Dall'Asta, Ke, and Vespignani (2005) visualizes the co-authorship of 20 years of information visualization research. The largest component of the collaboration graph strongly resembles the Picbreeder influence graph: a few, highly connected collaborators connected to many of the other authors. However, one important difference is that the information visualization authorship graph contains numerous small disconnected components, meaning that the co-authorship is more fragmented than Picbreeder influence. The authors mention that there are typically strong connections between authors at the same physical research centers. This fact may explain the relative lack of fragmentation in the Picbreeder interaction graph, because it has a "flat" organizational structure where physical collocation does not matter.

To compare to a typical academic publishing graph, an analysis framework from (Albert & Barabási, 2002) was used. Table 1 compares Picbreeder data with two sets of academic co-authorship data, one from neuroscientific authors (with analysis in Barabasi et al. (2002)), and another from *MEDLINE* (with analysis in (Newman, 2001)). The number of nodes, the average degree ($<k>$), the average shortest path length (l) and the graph clustering coefficient (C) are given in the table. Each entry is also compared with the average shortest path length (l_{rand})

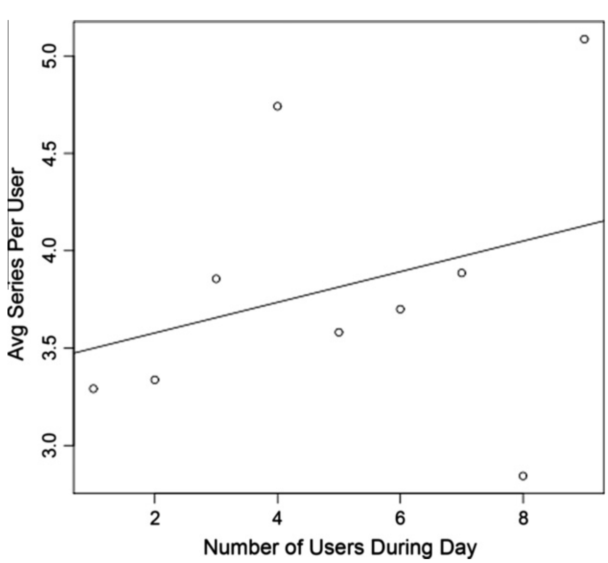

Fig. 5. Creative productivity per user on days of varying traffic. Days with larger numbers of users publishing may drive other users to be more productive.

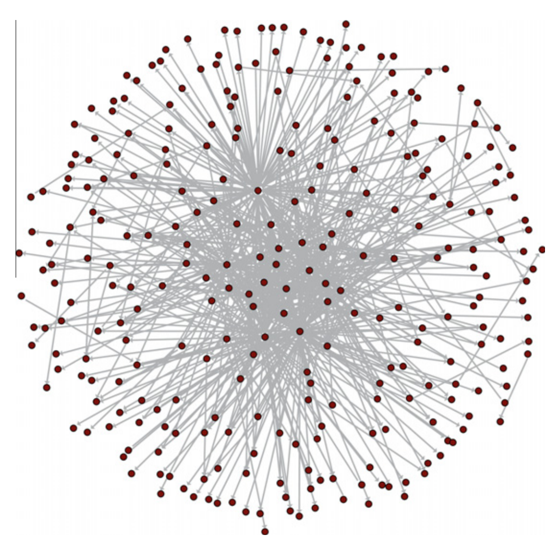

Fig. 6. Visualization of Picbreeder collaboration graph. It is easy to see a few highly connected members in the center.

Table 1
Comparison of network metrics for Picbreeder graph and Co-authorship graphs.

Graph	Nodes	<k>	l	l_{rand}	C	C_{rand}
Picbreeder	286	4.986	2.628	3.52	0.681	0.0174
Neurosci. co-authors	209293	11.5	6	5.01	0.76	5.5×10^{-5}
MEDLINE co-authors	1520251	18.1	4.6	4.91	0.066	1.1×10^{-5}

and cluster coefficient (C_{rand}) of a randomly generated graph with similar number of nodes and average degree. As is characteristic with the co-authorship graph, the Picbreeder influence data has a large clustering component value compared to the randomly generated graph.

In analyzing the degree distribution of the Picbreeder collaboration graph, it fits a power law with coefficient 1.79 (KS = 0.0927). While surveying the degree exponents of collaboration graphs (Albert & Barabási, 2002) found degree exponents in the range of 1.2–2.5, which are in agreement with the Picbreeder data. Fig. 7 graphs the degree distribution on a log–log scale, which shows a relatively typical power law relationship.

One underlying mechanism that can generate these kinds of network characteristics is known as *preferential attachment*. This is a scientific formulation of a "rich get richer" effect, where nodes that are already popular, tend to gain more influence because of that popularity. For instance, in the case of the academic authorship, well known academics are likely to attract even more coauthors, funding, better Ph.D. students etc. In Picbreeder, any preferential attachment must be stimulated by the images themselves, because users have no other dimension of reputation. Therefore, despite the fact that almost all interaction in Picbreeder is stigmergic, the influence data follows similar distributions to co-authorship data, which can be influenced by non-stigmergic communication and interaction.

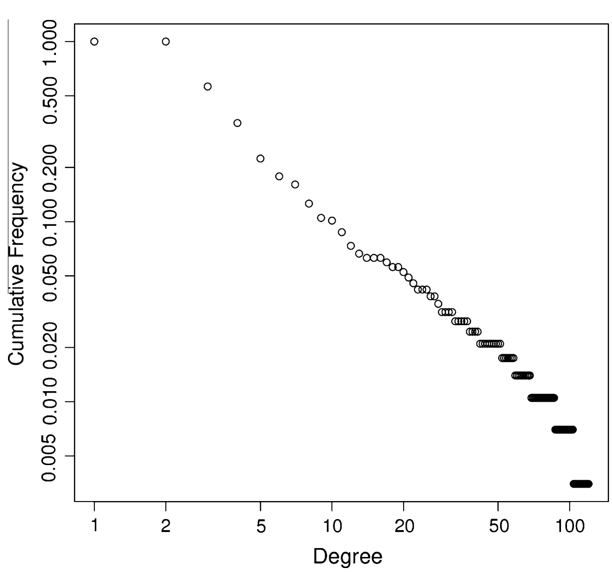

Fig. 7. Degree distribution for Picbreeder collaboration. The graph on a log, log scale shows a clear power law relationship.

5.2. Distribution of creative productivity

The concept of stigmergy not only describes ways that an agent stimulates the work of other agents, but also the way an agent is stimulated by its own work. In Picbreeder, users who already had highly rated images were more likely to contribute more highly rated images (Secretan et al., 2011). There is likely another stigmergic positive feedback loop at work here: a user who has seen his own positively received creative output is spurred on to create more.

One of Picbreeder's central tenets is to enable users who do not necessarily have any artistic expertise to participate in a creative community. Because the evolutionary algorithm handles the generation of artifacts, users rely more on good taste and willingness to explore the design space than on their technical proficiency in art. However, Picbreeder still appears to be subject to some common patterns that emerge in other collaborative creative endeavors.

Especially in regard to academic paper authorship, the creative output of an individual tends to follows Lotka's Law, a special application of Zipf's law (Lotka, 1926). The number of authors with n contributions is about $1/n^a$ of those making one contribution. While typically in academics, $a = 2$, in the case of Picbreeder, $a = 1.58$ (KS = 0.137), indicating that the distribution is less drastically pronounced in the most productive members. This reflects the more democratized creativity of Picbreeder. The histogram of creative output distribution is shown in Fig. 8. This result again demonstrates that despite the fact that Picbreeder's interaction is stigmergic, this kind of power law productivity is still present. Because Picbreeder has a very low barrier to contribution, its coefficient more closely reflects this. In academic publishing however, a significant amount of talent is also required, reflecting a steeper drop off in less productive users.

5.3. Properties of the artifact lineage

Part of the operation of the Picbreeder site involves promoting popular images, as measured by both the number of branched children, and the average star rating, to a more prominent position. The top of the Picbreeder front page contains panels featuring these images. The front page also contains panels for showing random images, as well as newer "up and coming" images. In this way, every image has a reasonable chance of climbing to the top of the rankings. Analyzing the number of children per image reveals that branching activity is not evenly spread across images. As is seen in many places within the Picbreeder data, the

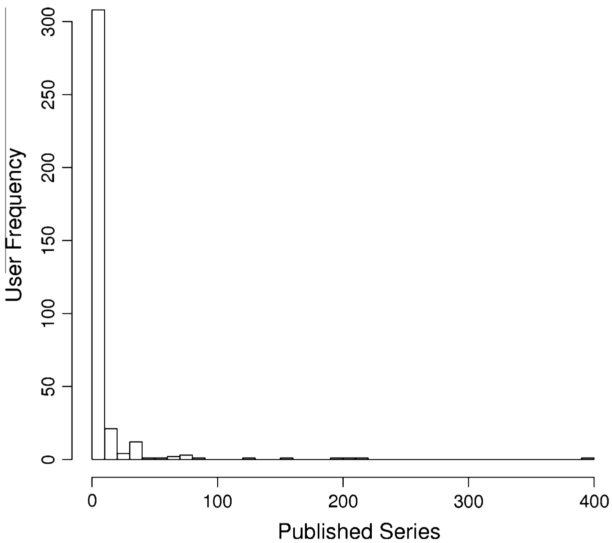

Fig. 8. Distribution of creative output. Similar to other domains of creative output, Picbreeder users have a power law productivity graph.

number of children per image follows a power law distribution with a coefficient of 2.30 (KS = 0.063). The distribution is illustrated in Fig. 9, confirming its power law nature. In a stigmergic process of positive feedback, images that are more interesting quickly amass more children, further promoting the image and resulting in another “rich get richer” effect that underpins many power law processes. This effect is almost certainly dependent on displaying highly rated and highly branched images in prominent positions.

With regard to the world of academic publishing, this branching graph is closely akin to a citation graph. (Redner, 1998) studied citation graphs of papers in the *Physical Review D* and found that the probability a paper is cited

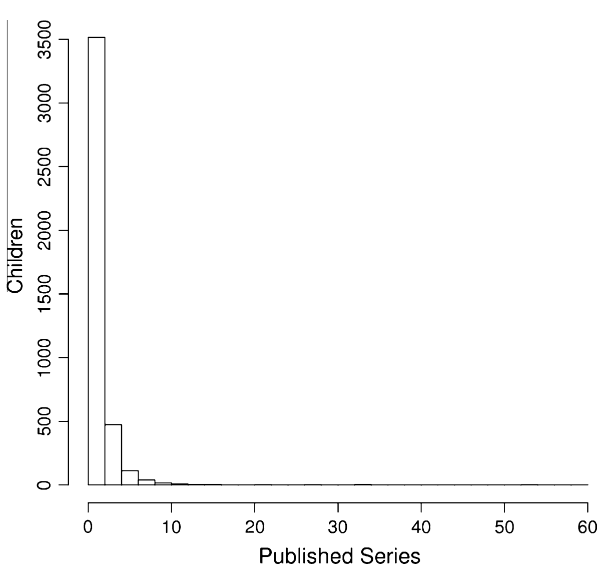

Fig. 9. Distribution of branched children per image. The number of children branched from each image also follows a power law distribution.

follows a power law distribution with an exponent of 3. Its higher exponent may be related to the fact that in Picbreeder, an image can only have one parent, where in academic publishing, numerous papers may be part of a particular paper’s “ancestry.” It is also not clear which papers are cited because they are genuine influences on research, and which are cited for completeness. The power law coefficients between the Picbreeder data and citation data are close, suggesting a relatively similar effect to other creative communities. No doubt, in the world of academic publishing, once a paper begins to garner attention, it is likely to gain further acclaim. Its more frequent citation may cause other authors to cite a paper in their own, establishing it as part of the standard canon.

6. Conclusion

This paper has shown that Picbreeder, an almost fully stigmergic means of collaborative creative interaction, follows many of the same patterns as other collaborative creative networks. Picbreeder demonstrates that it is possible to facilitate creative collaboration through entirely stigmergic means, and this paper explored the mechanisms that gave rise to that stigmergy. Because in other creative communities, stigmergic and non-stigmergic components of creative interaction are difficult to separate, Picbreeder provided an ideal opportunity to study this dimension. It is hoped that future studies will be able to isolate and study the contribution of stigmergic components in other creative communities.

It is also hoped that more quantitative analysis will be done on other creative communities. Academic publishing bibliometrics were used because they are plentiful and easy to access. While it is difficult to trace influence in similar way in musical or visual arts communities, developing techniques to analyze these communities is a worthwhile pursuit. This analysis may provide answers of real economic value. For instance, to answer the question, what will create a broader, more economically viable base of musical development, a US style system in which music distribution is dominated by a few large gatekeepers to the music industry, or a Canadian style system which frequently uses government sponsored incentives to encourage development in musical communities?

There is a great deal of analysis left to be done and questions to be answered with respect to the dynamics of creative communities. For instance, how can Axelrod’s model of cultural diffusion (1997) explain creative influence? Also, how can Friedkin’s analysis of weak ties versus strong ones in organization flows (1982) inform the analysis of how creativity develops within and between organizations. Picbreeder is currently a “flat” community, which does not fully represent the wide variety of social creative arrangements. The addition of this dimension to analysis will hopefully yield additional insight.

Stigmergy is clearly involved in creativity. It is no accident that Silicon Valley is well known for technical innovation and Paris is a well known muse of artists. These physical locations host large collaborative and competent communities for one, but also frequently display and demonstrate the results of their interaction, to “stimulate and guide” other participants. Other creative communities might benefit by explicitly taking advantage of stigmergic concepts to improve their efficacy. Imagine a paint studio where artists paint in a circle, with the paintings facing inward. Or a research lab where everybody’s latest work in progress is posted to a highly visible electronic board. The more we understand the role of stigmergy in creativity, the better we can shape and guide the process. Ultimately, every creative discipline, along with humanity itself, will be the beneficiaries of this advancement.

References

Albert, R., & Barabási, A.-L. (2002). Statistical mechanics of complex networks. *Reviews of Modern Physics, 74*, 47–97.

Anonymous (2006–2007). Living Image Project. <http://www.w-shadow.com/li/>.

Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *The Journal of Conflict Resolution, 41*(2), 203–226.

Balogh, J., Dubbin, G., Do, M., & Stanley, K. O. (2007). Dance evolution. In *Proceedings of the twenty-second national conference on artificial intelligence (AAAI-07) AI video competition*. Menlo Park, CA. AAAI Press.

Barabasi, A., Jeong, H., Neda, Z., Ravasz, E., Schubert, A., & Vicsek, T. (2002). Evolution of the social network of scientific collaborations. *Physica A Statistical Mechanics and its Applications, 311*, 590–614.

Bennett, C. H., Li, M., & Ma, B. (2003). *Chain letters & evolutionary histories* (pp. 76–81). Scientific American.

Biles, J. (1994). Genjam: A genetic algorithm for generating jazz solos. In *International computer music conference* (pp. 131–137). Aarhus, Denmark.

Börner, K., Dall’Asta, L., Ke, W., & Vespignani, A. (2005). Studying the emerging global brain: Analyzing and visualizing the impact of co-authorship teams: Research articles. *Complex, 10*, 57–67.

Collins, N. (2002). Interactive evolution of breakbeat cut sequences. In *Proceedings of cybersonica*. London, England: Institute of Contemporary Arts.

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems, 2*(4), 303–314.

Dawkins, R. (1986). *The blind watchmaker*. Norton & Company Inc.

Dawkins, R. (1989). *The selfish gene* (3rd ed.). Oxford, UK: Oxford University Press.

Finke, R., Ward, T., & Smith, S. (1992). *Creative cognition: Theory, research and applications*. Cambridge: MIT Press.

Friedkin, N. (1982). Information flow through strong and weak ties in intraorganizational social networks. *Social Networks, 3*, 273–285.

Gabora, L. (1997). The origin and evolution of culture and creativity. *Journal of Memetics, 1*.

Gong, T., & Wang, W. S. (2005). Computational modeling on language emergence: A coevolution model of lexicon, syntax and social structure. *Language and Linguistics, 6*, 1–41.

Grassé, P. P. (1982). *Termitologia: Anatomie, physiologie, biologie, systematique des termites*. Masson.

Hart, D. A. (2007). Toward Greater Artistic Control for Interactive Evolution of Images and Animation. In *Proceedings of the 2007 Evoworkshops on Evocomnet, Evofin, Evoiasp, Evointeraction, Evomusart, EvoSTOC and Evotranslog: Applications of Evolutionary Computing* (Vol. 4448, pp. 527–536). Valencia, Spain. Springer-Verlag.

Hastings, E., Guha, R., & Stanley, K. O. (2007). NEAT Particles: Design, representation, and animation of particle system effects. In *Proceedings of the IEEE symposium on computational intelligence and games (CIG-07)*, Honolulu, HI. IEEE Press.

Heylighen, F. (2007). Why is open access development so successful?: Stigmergic organization and the economics of information. In B. Lutterbeck, M. Bдrwolff, & R. A. Gehring (Eds.). Open Source Jahrbuch.

Hoover, A., & Stanley, K. O. (2009). Exploiting functional relationships in musical composition. *Connection Science Special Issue on Music, Brain, and Cognition, 21*(2 and 3), 227–251.

Husbands, P., Germy, G., McIlhagga, M., & Ives, R. (1996). Two applications of genetic algorithms to component design. *Evolutionary Computing: LNCS, 1143*, 50–61.

Johanson, B. & Poli, R. (1998). GP-music: An interactive genetic programming system for music generation with automated fitness raters. In *Proceedings of the third annual conference: Genetic programming* (pp. 181–186).

Kroah-Hartman, G., Corbet, J., & McPherson, A. (2009). Linux kernel development. <http://www.linuxfoundation.org/sites/main/files/publications/whowritesl%inux.pdf>.

Langdon, W. B. (2005). Pfeiffer – A distributed open-ended evolutionary system. In B. Edmonds, N. Gilbert, S. Gustafson, D. Hales, & N. Krasnogor (Eds.), *AISB’05: Proceedings of the Joint Symposium on Socially Inspired Computing (METAS 2005)* (pp. 7–13). UK: Hatfield.

Lotka, A. J. (1926). The frequency distribution of scientific productivity. *Journal of the Washington Academy of Sciences, 16*(12), 317–324.

Nelson, G. L. (1993). Sonomorphs: An application of genetic algorithms to growth and development of musical organisms. In *4th Biennial Art and Technology Symp*. (pp. 155–169). New London, CT.

Newman, M. (2001). The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences, 98*(2), 404–409.

Nishino, H., Takagi, H., Cho, S., & Utsumiya, K. (2001). A 3D modeling system for creative design. In *Proceedings of the 15th international conference on information networking* (pp. 479–487). Beppu City, Oita, Japan: IEEE Press.

Redner, S. (1998). How popular is your paper? An empirical study of the citation distribution. *European Physical Journal B, 4*, 131–134.

Secretan, J., Beato, N., D’Ambrosio, D. B., Rodriguez, A., Campbell, A., Folsom-Kovarik, J. T., et al. (2011). Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary Computation, 19*(3), 373–403.

Sims, K. (1993). Genetic images interactive exhibit. *Centre Georges Pompidou Revue Virtuelle Notebook*, 5.

Sims, K. (1997). Galapagos interactive exhibit, decordova museum, Lincoln, MA. <http://www.genarts.com/karl/geneticimages>.

Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines Special Issue on Developmental Systems, 8*(2), 131–162.

Stanley, K. O., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation, 10*, 99–127.

Stanley, K. O., & Miikkulainen, R. (2004). Competitive coevolution through evolutionary complexification. *JAIR, 21*, 63–100.

Takagi, H. (2001). Interactive evolutionary computation: Fusion of the capacities of EC optimization and human evaluation. *Proceedings of the IEEE, 89*(9), 1275–1296.

Theraulaz, G., & Bonabeau, E. (1999). A brief history of stigmergy. *Artificial Life, 5*(2), 97–116.

Unemi, T. (1999). SBART 2.4: breeding 2D CG images and movies and creating a type of collage. In *Knowledge-Based Intelligent Information Engineering Systems, 1999. Third International Conference* (pp. 288–291). Adelaide, SA.

Voss, J. (2005). Measuring Wikipedia. In *Proceedings of ISSI*, Stockholm.

Wood, P. (2010). *John barleycorn revisited: Evolution and folk song*. The Magazine for Traditional Music Throughout the World.