

Picbreeder: A Case Study in Collaborative Evolutionary Exploration of Design Space

Jimmy Secretan*, Nicholas Beato, David B. D'Ambrosio, Adelein Rodriguez, Adam Campbell, Jeremiah T. Folsom-Kovarik, and Kenneth O. Stanley
Department of Electrical Engineering and Computer Science, University of Central Florida
Florida

*To whom correspondence should be addressed; E-mail: secretj@gmail.com

Abstract

For domains in which fitness is subjective or difficult to express formally, Interactive Evolutionary Computation (IEC) is a natural choice. It is possible that a collaborative process combining feedback from multiple users can improve the quality and quantity of generated artifacts. Picbreeder, a large-scale online experiment in collaborative interactive evolution (CIE), explores this potential. Picbreeder is an online community in which users can evolve and share images, and most importantly, *continue* evolving others' images. Through this process of branching from other images, and through continually increasing image complexity made possible by the underlying NeuroEvolution of Augmenting Topologies (NEAT) algorithm, evolved images proliferate unlike in any other current IEC system. This paper discusses not only the strengths of the Picbreeder approach, but its challenges and shortcomings as well, in the hope that lessons learned will inform the design of future CIE systems.

Keywords

Interactive Evolutionary Computation, Collaborative Interactive Evolution

1 Introduction

The process of design often requires significant skill and domain knowledge (Graham, 2004; Brown and Birmingham, 1997). A major challenge is that designs are in effect chosen from vast search spaces wherein the most appealing are difficult to find. If, for instance, da Vinci had chosen a design for the Mona Lisa that lacked her enigmatic smile, the painting might not have attained the same iconic status. Searching the vast design space requires significant effort on the part of experts. Thus, making the design process more efficient through automating or augmenting the capabilities of designers is an important research focus (Goldschmidt and Porter, 2004; Shneiderman, 2007). Interactive Evolutionary Computation (IEC) (Takagi, 2001), i.e. artificial evolution guided through human direction, is such a tool that can potentially assist the design process. 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. 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 (Sims, 1993, 1997; Rooke, 2002; Machado and Cardoso, 2002; World,

1996; Hart, 2007; Unemi, 1999), music (Johanson and Poli, 1998; Nelson, 1993; Biles, 1994; Collins, 2002; Hoover and Stanley, 2009; Tokui and Iba, 2000), three dimensional models (Nishino et al., 2001; Husbands et al., 1996), movies (Unemi, 1999) particle systems (Hastings et al., 2007), dancing avatars (Balogh et al., 2007), and collages (Unemi, 1999), to name a few.

The field of Collaborative Interactive Evolution (CIE) attempts to improve IEC by involving multiple users in the evolutionary process (Sims, 1993, 1997; Langdon, 2005; liv, 2007). The hope is to increase the variety and quality of solutions that can be evolved. Yet effectively combining the opinions of multiple users is nontrivial because their preferences and goals are often in conflict. Picbreeder (<http://picbreeder.org>), an online service where Internet users collaborate to evolve pictures, introduces a new approach to this challenge by allowing each user to guide a branch of evolution on its own unique path.

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. Through the compounding effect of branching, and the ability of the underlying NeuroEvolution of Augmenting Topologies (NEAT) algorithm (Stanley and Miikkulainen, 2002b, 2004a) to increase the images' complexity, users succeed in collaboratively searching the space of possible images.

Picbreeder contributes a novel way to generate and maintain a large catalog of user-created content by enabling collaborative search through a vast design space by multiple users. It supports not only evolutionary *optimization*, but also evolutionary *exploration*. That is, the goal of evolution is not necessarily to find a particular image, but to accumulate images in the space of all images that are interesting to at least one user. Picbreeder also empowers users of all experience levels to enjoy recognition for their creative contributions by publishing them. Users thereby experience a new kind of creative social recreation through playful collaborative exploration. While Picbreeder focuses on generating images, it embodies a general framework that can harness the power of a large group of users to search any large design space. Yet images are a good starting point to effectively enable the study of collaborative evolutionary design because they are easy to appreciate intuitively.

This paper comprehensively describes the Picbreeder system and its results, reviewing successes, surprises, and remaining challenges, significantly extending Secretan et al. (2008b,a). In total, Picbreeder users have generated over 7,000 published images from over 140,000 combined generations since opening to the public to become among the largest online repositories of evolved and evolvable content. In part, this study shows how branching enables CIE applications to overcome the limitations of user fatigue, and how several features of the Picbreeder website encourage users to participate. Yet despite these successes, it is clear that Picbreeder specifically and IEC in general are still no match for true expert practitioners. Furthermore, while Picbreeder has received almost 90,000 visits and does have a dedicated core of users, most visitors show little interest or dedication to evolving content. Thus, the results of the Picbreeder experiment suggest that Picbreeder is a first step toward leveraging collaborative online communities, but questions remain on how to fully realize the power of IEC. Nevertheless, as a milestone on the path to augmented and artificial creativity that can rival and

perhaps someday surpass traditional designers, Picbreeder reveals important lessons with the potential to improve future IEC systems.

In the next section, background on groupware, IEC and CIE are discussed. The background also describes the NEAT algorithm and CPPNs, which underpin Picbreeder. Next, the Picbreeder approach is discussed (Section 3). The Picbreeder on-line system itself is described in detail in Section 4. In Section 5, details of the ongoing experiment follow. Then, in Section 6, the successes of the experiment are described, as well as the difficulties encountered and the challenges that remain. Future directions are then explored (Section 7), followed by conclusions (Section 8).

2 Background

Picbreeder represents the confluence of several different fields. Research in Picbreeder is broadly related to groupware and interactive evolutionary computation, as well as collaborative interactive evolution, an area of research at the intersection of these two fields. NEAT, the algorithm that powers the evolutionary process in Picbreeder, also powers several other genetic art programs. In Picbreeder, the NEAT algorithm evolves compositional pattern producing networks, which in turn render the images visible to the user. These research areas and technologies are described in this section.

2.1 Groupware

Groupware, which enables *computer supported cooperative work* (CSCW) (Grudin, 1994a; Carstensen and Schmidt, 1999), coordinates users interacting across a network for a shared purpose. For example, Sourceforge (sou, 2008) hosts open source projects and allows software and documentation to be written by loosely associated groups across the Internet. In the open source community itself, there is an implicit branching of projects whereby many works influence spin-off projects, or enable others. Sourceforge often serves as a hub for this branching process.

Another example, Wikipedia (wik, 2007), is a popular online encyclopedia that allows numerous users across the Internet to edit and add subjects about which they have some expertise. Similarly, in IBM's *Many Eyes* (Viégas et al., 2007), users can create and share graphic visualizations of data, from which the Internet community can draw and share their interpretations. While such services encourage collaboration, they also require specific talents or expertise, thereby limiting participation.

An early collaborative system based explicitly on branching was introduced by Jorda to help musicians collectively produce compositions (Jorda, 1999). In Jorda's *Faust* system, users can branch from a previously saved song and edit it to create a variant that is then saved back into the revision tree. Faust produced a collection of appealing songs, providing precedent for the approach in Picbreeder. However, Faust did not employ an evolutionary algorithm; rather, users had to manipulate the notes of the song directly and therefore required some basic musical knowledge.

In contrast, Picbreeder enables all users, including non-experts, to contribute creatively.

2.2 Interactive Evolutionary Computation

In single-user interactive evolution (Takagi, 2001), the user is presented with a set of alternatives generated by the system. This initial *population* is then evolved over generations through a process similar to domesticated animal and plant breeding: In each generation, the user *selects* the most promising designs, which are then mated and mutated to create the next generation. In effect, IEC assists the user in exploring a poten-

tially vast design space about which he or she may have little knowledge or expertise.

A prominent application of IEC is *genetic art*, in which evolution is driven by subjective appreciation for the evolved digital artifacts (Smith, 1991; Sims, 1991; Todd and Latham, 1992; Dawkins, 1989; Greenfield, 2000; McCormack, 1993; Lund et al., 1995; Baluja et al., 1994). Many genetic art programs follow the original *Blind Watchmaker* paradigm from Dawkins (1986), in which simple genetically-encoded patterns are evolved through an interactive interface. Genetic art can encompass a variety of digital media (Romero and Machado, 2007), including images (Machado and Cardoso, 2002; World, 1996; Hart, 2007), movies (Unemi, 1999), three dimensional models (Nishino et al., 2001), and music (Johanson and Poli, 1998; Nelson, 1993; Biles, 1994; Collins, 2002; Hoover and Stanley, 2009; Tokui and Iba, 2000). Notable examples include SBART (Unemi, 1999, 1994), which is an IEC program that evolves color images rendered by mathematical expression trees. Its capabilities have been extended to include generating movies and collages from source images. Another, GenJam, is an IEC system for evolving improvisational jazz music (Biles, 1994). With rhythm and chord progressions as inputs, the program generates melodies that the user can evaluate by song or measure. GenJam has been employed to produce a commercially-available album, as well as to play concerts with accompaniment (Biles, 2010). The art design system developed by Nishino et. al (Nishino et al., 2001) supports art education by interactively evolving parameters of a three dimensional model. Evolved variations on a model's motif not only help users generate better art, but also help them hone artistic sensibilities (Nishino et al., 2001).

Developments in genetic art have also influenced practical digital graphics tools, such as Kai's Power Tools Texture Explorer (Krause, 1996) and Apophysis (apo, 2007), which interactively evolve image textures and digital flame designs, respectively. IEC has been applied to a wide variety of other practical domains including industrial design (Husbands et al., 1996), police face sketching (Caldwell and Johnston, 1991), correcting speech distortion (Watanabe and Takagi, 1995), database retrieval (Lee and Cho, 1999), and several others (Takagi, 2001).

While IEC is a powerful approach to helping users generate digital artifacts, results are often limited by human fatigue (Takagi, 2001). According to Takagi (Takagi, 2001), a normal IEC process lasts 10–20 generations for a single user. However, the problem is that it may take more generations to produce notable results.

2.3 Collaborative Interactive Evolution (CIE)

CIE (Szumlanski et al., 2006) systems involve multiple users in one IEC application, hoping to create products with broader appeal and greater significance. CIE systems can be physical installations or online services; each has unique methods through which to merge input from multiple users.

Among the first CIE installations were two museum exhibits by Sims (1993, 1997). The Genetic Images (Sims, 1993) exhibit let visitors stand in front of several displays to select pictures to produce the next generation. Other users could select individuals at the same time, or continue the evolutionary run where prior users left off. In the Galapagos (Sims, 1997) exhibit, visitors evolved virtual three-dimensional creatures with a similar interface.

While these exhibits were innovative, the museum environment does not encourage users to frequently return to the installation. *Electric Sheep* (Draves, 2005), a distributed Internet application inspired by Sims' work, enables users to evolve images over a longer time period than is possible with a museum installation. The application

runs as a screen-saver that renders fractal animations called *sheep*. The sheep are distributed to users, who can vote for the ones they like. Both in the museum exhibits of Sims and in Electric Sheep, voting is the primary means of guiding evolution, and the ability of users to choose a starting point for the process is limited.

Pfeiffer (Langdon, 2005) is another pioneering CIE system that allows users to endorse candidates for further evolution in an online, multiuser environment aimed exclusively at evolving snowflakes. Even with this limited scope, Pfeiffer processed over 68,000 user inputs from every continent since 2001. While Pfeiffer demonstrates that users are willing to participate in CIE, it raises the question of whether a broader scope of images would be possible to evolve collaboratively.

Szumanski et al. (2006) introduced a different CIE framework based on conflict resolution. Users log into the system to vote on a particular individual selected by the system. To overcome user fatigue, the system combines these inputs to form a *fitness function*, i.e. a measure of quality, for a traditional genetic algorithm. The genetic algorithm then evolves an individual to meet the combined user requirements. This approach evolved characters for an interactive story. While the system effectively circumvents user fatigue, it does not encourage a proliferation of content because a large collection of user input is combined to reach only a single objective.

Another system, Imagebreeder (Hammond, 2007) also offers an online community coupled with an IEC client for evolving images. Users can save their creations to a general pool to be viewed by the larger community. However, Imagebreeder does not include the ability to continue evolving others' images, which means that the complexity of images evolved is limited to what a single user can evolve before fatiguing.

These systems highlight the intriguing potential for evolution to benefit from the collective input of users across the world. However, they also signal a potential drawback: Users' preferences are often in conflict, resulting in mediocre results that cannot satisfy divided opinions. Furthermore, genetic representations often limit genuine creativity by constraining search to a predetermined set of possibilities (e.g. only snowflakes; Langdon 2005). To evolve a broad class of images, an open-ended representation is needed that can potentially represent anything. Because of its ability to complexify, the NEAT evolutionary algorithm satisfies this requirement, as explained in the next section.

2.4 NeuroEvolution of Augmenting Topologies (NEAT)

The creative process can be constrained for novices by controlling the method through which the space of images is searched. In Picbreeder, this constraint is provided by an evolutionary algorithm called NeuroEvolution of Augmenting Topologies (NEAT; Stanley and Miikkulainen 2002b, 2004a), which addresses several fundamental challenges in evolving complex structures. Although NEAT was originally introduced as a method for evolving artificial neural networks (ANNs), a major appeal of NEAT is its ability to evolve increasingly complex structures of any type, so that evolutionary search is not limited to a fixed space of possibilities.

A significant obstacle to evolving complex structures is that heuristically determining the appropriate number of genes, i.e. the number of dimensions in the solution space, is difficult for challenging problems. For example, how many nodes and connections are necessary for an ANN that draws a picture of a face? The answers to such questions cannot be based on empirical experience or analytic methods, because little is known about the solutions. To address this problem, instead of trying to start evolution in the space of the final solution, NEAT starts evolution with a population

of small, simple genomes and elaborates on them over generations by adding new genes. Each new gene expands the search space, adding a new dimension of variation that previously did not exist. That way, evolution begins searching in a small, easily-optimized space, and adds new dimensions as necessary. This approach is more likely to find high-dimensional solutions than an approach that begins searching directly in the intractably large space of complete solutions. The process of *complexification*, i.e. incrementally adding new genes over generations, also occurs in some lineages in nature, leading to increasing phenotypic complexity (Martin, 1999; Amores et al., 1998; Carroll, 1995). By starting minimally and gradually complexifying over the course of evolution, NEAT was able to solve several difficult control problems (Aaltonen et al., 2009; Stanley et al., 2005a,b; Taylor et al., 2006; Risi et al., 2009; Stanley and Miikkulainen, 2002a; Trujillo et al., 2008; Kohl et al., 2006; Stanley and Miikkulainen, 2004b,a, 2002b).

It should be noted that IEC systems often rely on genetic programming and expression trees for image generation (Sims, 1993, 1997; Rooke, 2002; Machado and Cardoso, 2002; Unemi, 1999). While NEAT shares some features with these genetic programming (GP) techniques, there are important differences that motivate the use of NEAT in Picbreeder. While GP tree genomes are capable of variable-length representation, they do not follow the same incremental approach to gradually increasing complexity as in NEAT (Stanley and Miikkulainen, 2004a, 2002b). Some GP-based IEC systems (Sims, 1993, 1997; Rooke, 2002) employ complex functions (e.g. blurs and gradients) into their basic function sets. These functions can introduce a significant bias into the search space (further discussed in Section 2.7). While this bias can sometimes create appealing results, it can also limit the space of images through which the evolutionary algorithm is likely to search. Picbreeder in contrast aims to keep its representation flexible by foregoing such aesthetic bias.

Although it was originally introduced to evolve ANNs, NEAT is sufficiently general to evolve any variable-length genotype. Thus complexification is now a general tool for evolutionary computation. The next section introduces NEAT-based IEC for art, which is the basis for Picbreeder.

2.5 NEAT-based Genetic Art

Independent researchers have released several NEAT-based genetic art programs, beginning with Mattias Fagerlund's DelphiNEAT-based Genetic Art (DNGA) in 2003 (Fagerlund, 2006, 2005). DNGA was followed by Holger Ferstl's SharpNEAT-based Genetic Art (SNGA) in 2006 (Ferstl, 2006). While realistic-looking objects have been evolved with these applications (see Stanley (2007) for examples), they still can require a considerable number of generations to do so, and therefore can fatigue users.

The *Living Image Project*, which was introduced on the world-wide-web in September, 2006 and terminated in November, 2007, is another online collaborative NEAT-based genetic art system (liv, 2007). Living Image displays a population of color images generated by a NEAT-based genetic art program on a web page where users can *vote* for their favorite candidate. Each user is allowed to cast at most 25 votes in one day. After 300 votes, the next generation replaces the current generation and the process begins again. The idea is to integrate the preferences of a broad user population and thereby choose parent images more intelligently. Unfortunately, evolution is slow in such a system because many users must contribute to a single selection event. For example, after 14 months and 8,346 votes, the project achieved only 25 generations of evolution. Furthermore, it is difficult for the system to evolve toward a recognizable form because users have no way to coordinate their potentially conflicting choices.

The main idea in DNGA, SNGA, and Living Image is to enable NEAT to evolve a special kind of network that represents images. These evolved networks are called *compositional pattern producing networks* (CPPNs) because they produce patterns in space by composing functions (Stanley, 2006, 2007). The next section explains how CPPNs represent arbitrary images.

2.6 Compositional Pattern Producing Networks (CPPNs)

A CPPN is a function of n Cartesian dimensions that outputs a pattern in n -dimensional space. For example, a two-input CPPN produces a two-dimensional image (Stanley, 2007). For each (x, y) coordinate in that space, its color is output by the CPPN that encodes the image. It is also possible to augment the (x, y) coordinate inputs with additional (non-orthogonal) inputs that the CPPN can exploit. Figure 1 shows how a two-dimensional image is generated by a function of x , y and the distance d of each pixel from the center of the image.

Internally, a CPPN is represented as a connected-graph (i.e. a network) of functions chosen from a standard set including sine, Gaussian and sigmoid. The structure and connection weights of the graph represent how the functions are composed to process each input coordinate. The output value of a function node is multiplied by the weight connecting it to the input of the next function node. If multiple input connections feed into the same function node, then that node takes the sum of its weighted input values as its input. While the representation is intentionally biased toward exploiting the geometric regularities (e.g. symmetry) that commonly occur in images through its set of standard 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 HSB color space was chosen over the more common Red, Green, Blue (RGB) color space because it allows the greyscale outline of the image (i.e. brightness) to be evolved separately from the color (which is represented in the hue and saturation channels). Picbreeder also provides the distance from the current pixel to the image center as input to the CPPN, to help it further exploit image regularity. Figure 2 illustrates an example image from Picbreeder and the underlying CPPN that renders it.

While CPPNs are similar to ANNs, they differ in their set of activation functions and how they are applied. While ANNs often contain only sigmoid and Gaussian functions, CPPNs can include both types of functions and many others. The choice of functions for the standard set creates a *bias* toward specific types of patterns and regularities. Thus the architect of a CPPN-based genetic art system can bias the types of patterns it generates by deciding the set of standard functions to include.

Furthermore, unlike typical ANNs, CPPNs are applied across the entire space of possible inputs so that they can represent a complete image. Because they are compositions of functions, CPPNs in effect encode images at infinite resolution and can be sampled for a particular display at whatever resolution is optimal. The next section describes why CPPNs are well suited to represent images.

2.7 CPPN Image Representation

Representation is crucial in both evolutionary computation and artificial intelligence in general. A good representation can both efficiently encode complex information and also organize it effectively for search. This section suggests how the choice of CPPNs for image representation can help evolutionary search.

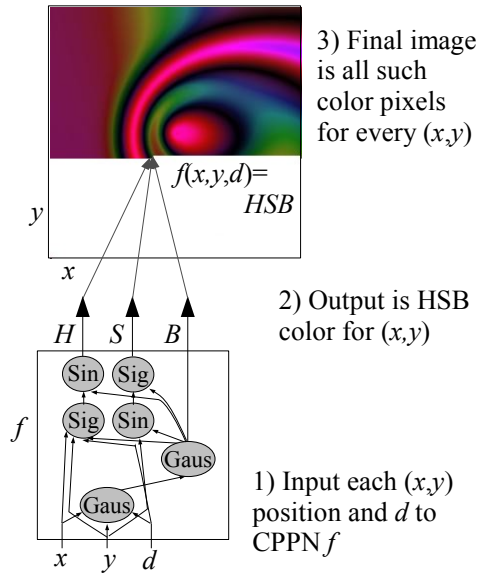


Figure 1: **How CPPNs work.** The diagram shows how CPPNs are internally represented and how a CPPN is converted to an image. An example CPPN f takes arguments x and y , which are coordinates in a two-dimensional space, and d (distance from center) (1). When all the coordinates are drawn with the hue, saturation, and brightness corresponding to the outputs of f (2), the result is an image (3). The CPPN network graph determines which functions connect to which. Note that the topology is unconstrained and can represent any relationship. The depicted functions exemplify the CPPN's ability to compose functions chosen from a standard set. This structure allows the CPPN to represent a large and diverse space of patterns, biased by the set of standard functions.

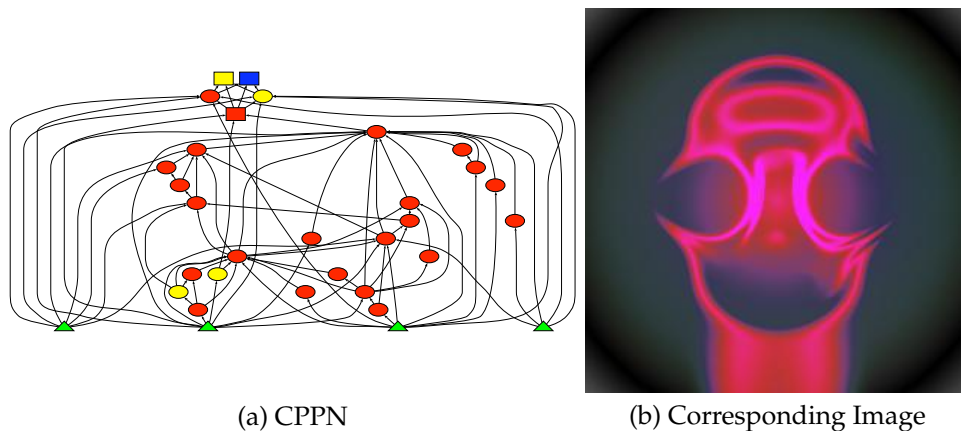


Figure 2: **Example CPPN and image.** The CPPN depicted in (a) renders the alien face in (b); triangles are input nodes, ovals are hidden nodes, and rectangles are output nodes.




Figure 3: **Exploring the space of images.** Images randomly generated pixel by pixel (a) exhibit no meaningful structure, exemplifying a vast and mostly unpromising search space. Images generated by random CPPNs (b) demonstrate the structure afforded by their representation, but do not resemble anything familiar.

McCormack (2008) points out that at 500x400 pixel resolution (200,000 total pixels) and a color depth of 24-bits, the design space of images contains 9.5×10^{1444943} candidates, a set he calls *everyimage*. Furthermore, within this space, McCormack shows that meaningful designs are exceedingly rare needles in a haystack:

While it is difficult to pin down the exact number, it is clear that the fraction of interesting images from the *everyimage* set is extremely Small¹. If you need proof, try randomly generating 500x400 pixel images for a few hours and see how many interesting ones you find.

To illustrate this point, the images in figure 3a are all generated by randomly selecting colors for each pixel. The implication is that the vast majority of this space is without either structure or appeal.

Particularly in evolutionary computation, significant research in recent years has sought to clarify how complex structures can be encoded most efficiently for evolutionary search (Bentley and Kumar, 1999; Hornby and Pollack, 2002; Stanley and Miikkulainen, 2004a, 2003; Angeline, 1995; Hart et al., 1994; Bongard and Pfeifer, 2001). It is now widely recognized that a good encoding allows information in the genotype to be *reused* in producing the phenotype (Gruau et al., 1996; Hart et al., 1994). Encodings with this property are called *indirect encodings* (Stanley and Miikkulainen, 2003). Human DNA, for instance, functions through an indirect encoding, mapping only about 30,000 genes to a body with trillions of cells and 100 trillion neural connections (Zigmond et al., 1999).

It turns out that CPPNs are an indirect encoding with several powerful representational properties that make them particularly suited to encoding and searching for spatial patterns. In particular, they are designed to efficiently encode repetition, repetition with variation, symmetry, and elaboration.

¹The capitalization of *Small* by McCormack borrows Daniel Dennett's notation of capitalizing (e.g. *Vast*) to emphasize the unfathomable scale of the spaces described.

Repetition is essential to many common forms from fish scales to windows on skyscrapers. It is naturally encoded by CPPNs that include periodic functions, such as sine and cosine. These functions generate repeating parts without the need to duplicate the information that encodes each part. *Repetition with variation* is another fundamental motif evident in biology, e.g. among fingers and vertebrae. Repetition with variation means that a pattern is repeated while varying each repeated element a small amount. It is accomplished in CPPNs by combining (i.e. summing) periodic functions with other functions (for instance, sines and Gaussians).

Symmetry, which is fundamental to faces, animals, and vehicles, allows the same information to encode both sides of an object. Symmetry is produced in CPPNs by symmetric functions, such as Gaussian. Furthermore, a Gaussian separated from the x and y inputs by several nodes can yield a symmetric subset of the image, even if the whole image is not symmetric.

Finally, the ability to gracefully elaborate is essential to image evolution. Elaboration encourages increasing complexity by making each image a springboard to the next level of complexity. The NEAT algorithm adds functions and connections to CPPNs as they evolve, thereby elaborating the images they encode. A full review of the capabilities of CPPN image representation can be found in Stanley (2007).

The set of functions used within CPPNs in Picbreeder are cosine, sine, Gaussian, identity, and sigmoid. These functions are chosen to capture regularities that appear frequently in nature (e.g. symmetry, repetition, and repetition with variation) without intentional aesthetic bias.

Each of the images in figure 3b is generated by a randomly evolved CPPN. That is, each was evolved for a different number of generations with a random selection 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 those in figure 3a, they still do not resemble anything recognizable. Thus, while CPPNs do bias the search towards more structured geometry, CPPNs do not automatically generate valuable designs.

3 The Picbreeder Approach

Picbreeder aims to address the primary challenges facing CIE:

1. **Empowering groups, regardless of skill, to collaboratively search a design space.** While groupware often coordinates users by sharing expertise (sou, 2008; wik, 2007), few such projects empower users who may lack such expertise.
2. **Overcoming user fatigue.** In existing single-user IEC applications and CIE services, most users succumb to fatigue before generating significant products (e.g. as in Hammond 2007; liv 2007; Ferstl 2006; Fagerlund 2005).
3. **Proliferating content.** Some CIE systems do not encourage a proliferation of content, but instead concentrate the efforts of many users on single decisions (liv, 2007).
4. **Collaborating without diluting individual contribution.** While existing CIE systems aim to produce more meaningful output by involving many users, they frequently average the contributions of many users to generate an image that is not necessarily pleasing to any (liv, 2007; Sims, 1993, 1997).

5. **Encouraging participation.** CIE systems need to encourage participation through recognizing user achievements and through a flexible interface.
6. **Balancing exploitation with exploration.** Any method for searching design space needs to balance *exploitation* with *exploration*; that is, users should be able to choose to continue evolving less prominent designs in the hope that future generations may improve.

This section describes the innovations designed to allow Picbreeder to address these challenges.

3.1 Empowering Groups to Collaboratively Search Designs

Creating a large corpus of interesting designs can yield significant value, but often requires users with substantial expertise and skill (e.g. Sourceforge; sou 2008). Therefore, there is a strong motivation to automate or increase the efficiency of the design process (Goldschmidt and Porter, 2004; Shneiderman, 2007). However, the design process is difficult to automate, partly because it relies on subjective tasks that are difficult for computers but relatively easy for humans, e.g. visually parsing scenes, understanding speech, and, significantly for Picbreeder, appreciating art. Importantly, even users without special expertise can perform tasks that are too difficult or subjective for a computer. Picbreeder takes advantage of this fact, as well as the generative capabilities of evolutionary algorithms, to tap into the creativity of large online user communities.

By allowing the user to direct what types of artifacts should be proliferated, Picbreeder empowers users, regardless of talent, to search a vast design space. Users simply choose an image that will become the seed for a new evolutionary progression. Within the progression, users then select which of the generated images they find more compelling. Through the evolutionary process, the images are mated and mutated to produce a new generation. Because of the CPPN-NEAT algorithm, the images will gradually become more complex and reflect the selections of the user. Through these mechanisms, users can evolve complex digital content irrespective of their level of experience, unlike in most groupware systems.

The Picbreeder IEC client program that supports evolving images is shown in figure 4. The user breeds images by selecting between one and 14 images from the 15 displayed. The client can operate in either basic or advanced mode. The basic mode only enables the *Evolve* and *Publish* buttons. The user can press *Evolve* to produce the next generation of images from the current selections. When the user decides that an image in the image panel is suitable, he or she selects that single image and presses *Publish*. A publishing interface then allows the user to share it with the community through the Picbreeder website. The advanced mode enables additional client functions. In case the user does not find any images that are worth selecting, the *Redo* button respawns the current generation. The *Back* button allows the user to return to a previous generation and restart the evolutionary progression from there. The user can also navigate back up to the current generation with the *Forward* button, much like in a web browser. The *Save* button stores the evolutionary progression of an image in the user's account as an *unpublished image*, so that the user can continue it later. A slider bar alters the mutation rate, with the limits described to the user as *Small Changes* and *Big Changes*. A checkbox allows the user to specify whether to evolve in color or in greyscale. If color is enabled, a drop down box allows the user to specify whether to evolve the structure (only the parts of the CPPN connected to the *brightness* output), the color (only the parts of the CPPN connected to the *hue* and *saturation* outputs) or both simultaneously.

By providing a simple interface to beginners and more options to advanced users, this client design aims to appeal to the broad Internet community so that everyone’s input can be harnessed to create evolved art.

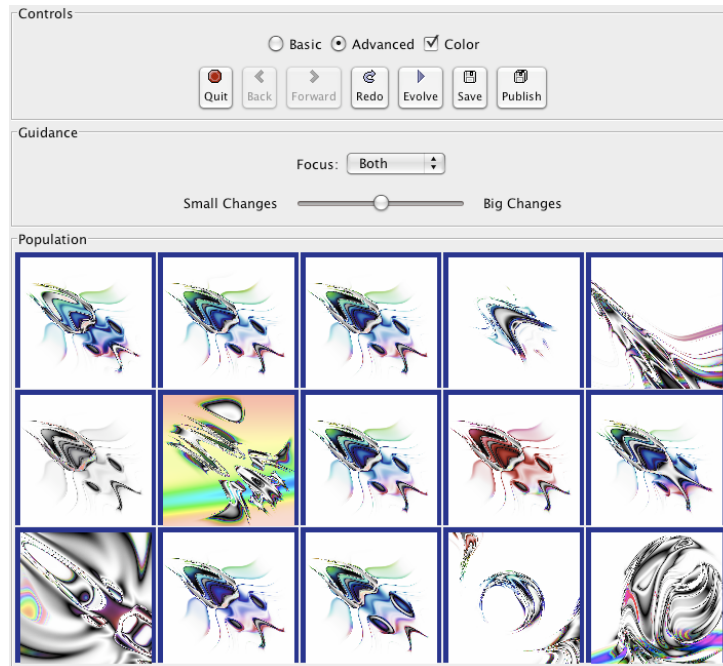


Figure 4: **The Picbreeder client program.** The interface can spawn or re-spawn a new generation, move back and forth through the generations, and either publish the image or save it for later editing.

While some users may aim to evolve a particular image (e.g. an insect), others may simply explore by selecting images that happen to appear compelling. A novel image appearing in one generation may inspire the user to continue in a different direction than initially expected. Even if the user has no concrete goal and is not familiar with the IEC process, the images should become more compelling to the user through the mechanics of the evolutionary process. In this way, the client supports the user’s creative exploration of a design space, allowing any user to continually guide the computer in directions of his or her own interest.

Thus, in contrast to expertise-based groupware (e.g. Sourceforge and Wikipedia), users without specific expertise can contribute images to Picbreeder through the simple IEC client. Picbreeder’s easy access (requiring only a Java-enabled web-browser) encourages wide participation and the online format makes it possible for Picbreeder to scale to larger communities than could be supported by a museum installation.

Picbreeder also recently introduced an *Image DNA* viewer. Through this interactive client (figure 5), users can manipulate an image’s underlying CPPN and observe the consequences. Note in figure 5 that each node displays the *intermediate* pattern that is output at that node. Displaying these intermediate patterns as images helps the user to understand how these patterns are fed into other nodes to compose the final output image. For any node, users can modify the activation function, introduce new connections to and from other nodes, and remove existing connections. Users can also

modify the weights of all connections to see their contribution to the overall image.

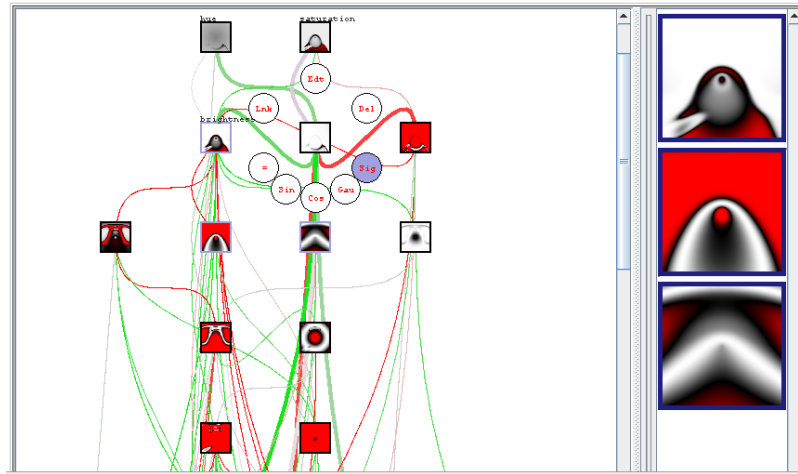


Figure 5: **The image DNA viewer.** The left panel illustrates the structure of the CPPN, and the right panel renders the output of individual nodes. Also, a small version of the intermediate pattern at each node is shown at the location of each such node in the left panel. When the user selects a node, the viewer reveals a set of circular control buttons that can modify the links and activation function. The viewer helps users understand and appreciate what is being evolved.

The Image DNA viewer demonstrates the educational potential of Picbreeder. It intuitively conveys the importance of indirect encodings to evolutionary computation (Stanley and Miikkulainen, 2003) and provides a way to explore the internal dynamics of neural-network like structures (Stanley, 2007). Through observing the effect of changes in network structure, activation functions, connection weights, and by observing intermediate patterns within the network, users gain greater understanding of how images are rendered and how their underlying CPPNs are evolved, thereby enhancing the expertise of Picbreeder users in general.

3.2 Overcoming User Fatigue

It can take many generations of evolution to find interesting designs within a large search space. There is a significant chance that within the typical 10 to 20 generations of IEC (Takagi, 2001), the user will not see anything significant, and may lose interest in exploring further. Even if the user retains interest throughout many more generations, and is willing to continue evolution over several sessions, there are fundamental limits to the number of generations a person can evolve. Without a means to accumulate many more generations of evolution, it is difficult for images to evolve into anything significant. User fatigue is thus a fundamental problem in IEC (Takagi, 2001) that single-user IEC systems often do not address (Fagerlund, 2005; Ferstl, 2006).

Picbreeder addresses user fatigue through a mechanism called *branching*. If the user finds an interesting image on the Picbreeder website, he or she can then choose to branch it, which means continue its evolution. As branches accumulate upon branches, it becomes possible for the complexity of an image to compound for hundreds of generations without fatiguing any single user. Because the user is likely to branch from images that interest him or her and because the IEC process steers images closer to the

user's preferences, conflict over evolving a single image is eliminated. The originating image and the results of its new branches are all stored separately on the Picbreeder website, allowing continued access to all of them.

A typical user session in Picbreeder begins with viewing published images (as seen in figure 6), which can be filtered by different criteria such as *highest rated* and *newest*. Users can choose to branch any image they see, thereby entering the IEC client program (figure 4), which loads a copy of the root image's CPPN. The user then continues the image's evolution through the IEC process, and publishes the branch when satisfied with its appearance.

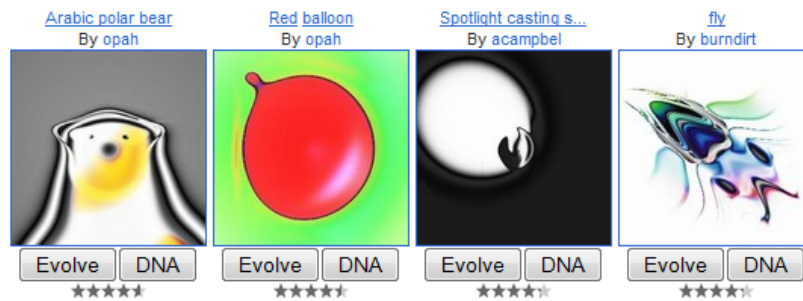


Figure 6: **Example image view from the Picbreeder website.** Any of these images can be branched for further evolution by pressing *Evolve*. The *DNA* button displays the Image DNA viewer (figure 5) and allows the user to manipulate the image's underlying CPPN.

When the user branches, Picbreeder follows the process illustrated in figure 7. The collection of genomes evolved throughout the generations of a single evolution session, along with their associated images, is a *series*. When a series is published, the last individual selected is its *representative*. While Picbreeder retains every image in each series for future analysis, users browsing the site only see representative images. When branched, a representative's genome spawns the first generation of the new branched series. This design accommodates branching while keeping individual series in the chain intact, thereby allowing long chains of content to grow while minimizing the work of each individual user.

Because NEAT lets images complexify throughout evolution, images that have been evolved through a chain from many other users may have already gained significant complexity. Therefore, users can immediately begin with complex structures through branching. Figure 8 illustrates the benefit of branching from an already-complex image. In figure 8a the user required 40 generations to evolve from a random initial starting point to a compelling image. In contrast, it took only 14 generations in figure 8b to evolve the final image as a branch from a prior image because the new image borrows significant structure from its parent (i.e. a shoreline, a skyline, and the general layout of the scene). In this way, user fatigue is overcome.

3.3 Content Proliferation

Some CIE systems (Szumlanski et al., 2006; liv, 2007) combine user input to generate only a few products, which means that the amount of content generated per person in the system is less than in single-user IEC. In contrast, branching in Picbreeder creates a new image with every branch. Importantly, an image may be branched multiple

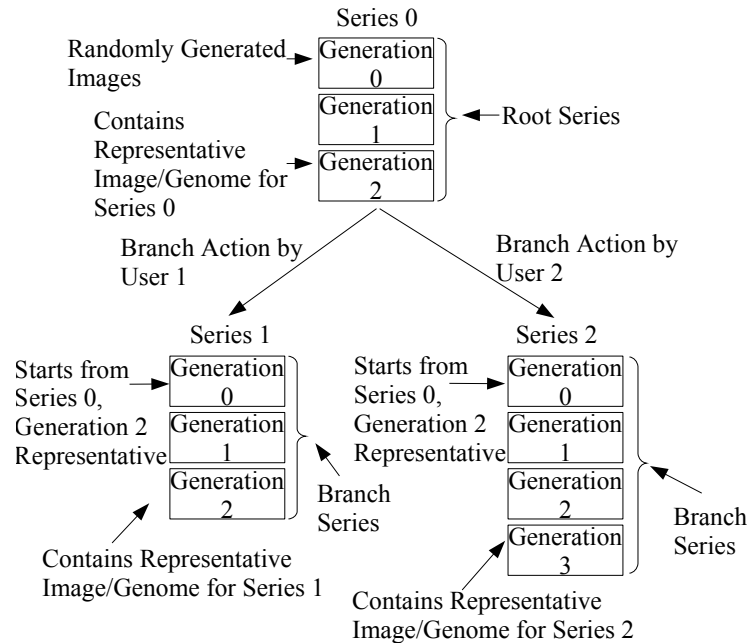


Figure 7: **How branching works in Picbreeder.** The server saves each sequence of generations in a *series*. The representative genome of series 0 seeds the initial generations of other series branched from it.




Figure 8: **Overcoming user fatigue through branching.** The example in (a) was evolved from scratch (i.e. from a random initial image) in 40 generations, beginning with the individual on the left and ending on the right. In (b), the user branched from the individual on the left to yield the individual on the right (another landscape) in only 14 generations by reusing some of the existing image's structure.

times, and all images are preserved indefinitely. Thus, instead of needing many users to generate few images (liv, 2007), Picbreeder allows even a few users to generate many.

Furthermore, content is not displayed on the website from every generation of each series (as in Sims 1993, 1997) but instead from images that users found *worthy of publishing*. Thus, what results is a proliferation of *meaningful* content through ever-expanding branches. Although all branches can be ultimately traced back to an initial series that started from nothing but a completely random population, a diverse and meaningful set of distinct styles and themes nevertheless proliferates.

3.4 Collaboration Without Diluting Individual Contribution

Although systems like The Living Image Project (liv, 2007) are seminal in promoting the idea of CIE, their focus is to merge the artistic sensibilities of several individuals, which can obfuscate the contributions of the individual user. It is possible for users to cancel out each others' contributions by pulling in opposite artistic directions. Furthermore, in most CIE systems (liv, 2007; Sims, 1997, 1993; Szumlanski et al., 2006), it is difficult to determine what contributions each user made to the evolution of a particular image. If users are not recognized for their contributions, they may lack motivation to participate in the system.

Picbreeder's branching also addresses this problem. Each lineage is tracked such that although a branched image is linked to its parent, the user can nevertheless continue evolution in any way desired. Each chain of branches is influenced by every contributor in the chain, yet each individual user takes it on a path chosen solely by that user.

Picbreeder provides a simple interface for browsing the images and users that have contributed to a lineage tree. Recall that each image displayed on the site is a representative of a series that began with its parent series' representative. The most proximate series in an image's lineage can be inspected in a *detailed view* panel (figure 9). In this way, branching creates content that embodies the tastes and ideas of many different users, all while maintaining individual ownership and creative autonomy.

3.5 Encouraging Participation

Without an active user base, images cannot branch and complexify. A major challenge for any groupware is to assemble a critical mass of users to become productive (Grudin, 1994b). An effective method to entice users is to highlight the most appealing content from which to branch. Picbreeder therefore also motivates participation through image ratings and user rankings, and enables users to find interesting images through tagging, browsing, and searching mechanisms.

Users can rate interesting images and thereby credit other users for their creations. The average rating is shown under the image. The front page of the site shows a group of *all time top rated* images, the full set of which includes all images with three or more ratings and an average rating of 3.0 or greater, sorted by average rating. Furthermore, the Picbreeder front page shows the *most branched*, which are images sorted by the number of times they have been branched to evolve new images. The front page also includes a sampling of the *best new* images category, which includes any image evolved with an average rating of 3.0 or greater, sorted by most recent first. This view provides a vehicle for new images to accumulate enough votes to allow entrance into the all time top rated images. All of these views aim to maximize participation by immediately showing users images that have generated the most interest. They also encourage users to publish interesting images, so that they too can be featured on the front page. In addition, Picbreeder assesses overall rankings for each user, based on the number of images published within the last 120 days that have entered the all time top rated category. This ranking encourages users to continuously contribute popular images.

Picbreeder further helps users find interesting images from which to branch through *tagging*, *browsing*, and *searching*. Tags associated with an image during publishing let users find images that match their interests. Users can search for tags with a search-engine-style interface. In addition, tags are automatically grouped into browseable categories and subcategories. The text boxes in which tags are entered during publishing suggest tags as the user types to reduce redundancy. The most frequent

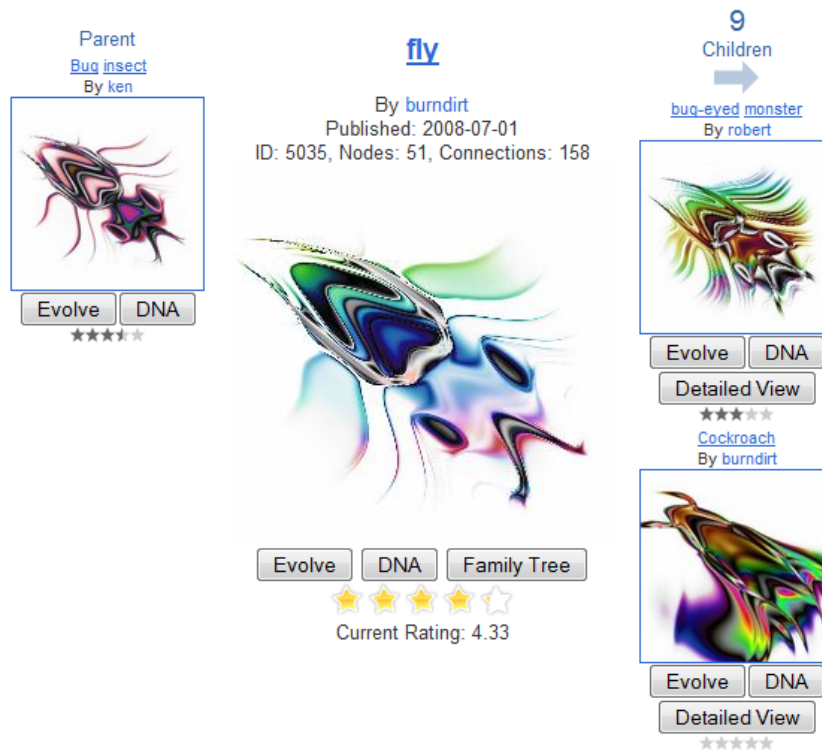


Figure 9: **Browsing children and parents.** This view shows both the children and the parent of an image. Users can explore the tree by clicking on any image, showing them which users participated at every level.

tags form the top-level categories. Images with these tags are queried for their other associated tags, which provide the next level of the categorical hierarchy. This approach creates browseable hierarchies without administrative intervention (figure 10).

As noted above, the user can browse any of the front page categories, or follow an image's children or parents to find more interesting images. Picbreeder also features a *Tree of Life Browser* (pictured in figure 11), a Java applet that allows users to visualize the lineage tree from a high level. The user can interactively expand different branches of the tree, thereby gaining a sense of the motifs through which parents and children are related.

Through the encouragement provided by ratings, and through making images available by tagging, browsing, and searching, Picbreeder makes participation easy and fun, which is essential for CIE to succeed.

3.6 Balancing Exploitation with Exploration

Interactive evolution of digital artifacts, much like its biological analog, can make unexpected leaps. In Picbreeder, images that are not valued by the community may become the roots for other images that are valuable. Therefore, in addition to encouraging users to evolve from images that are popular (i.e. exploitation), Picbreeder also provides several mechanisms that give less popular images the opportunity to evolve (i.e.

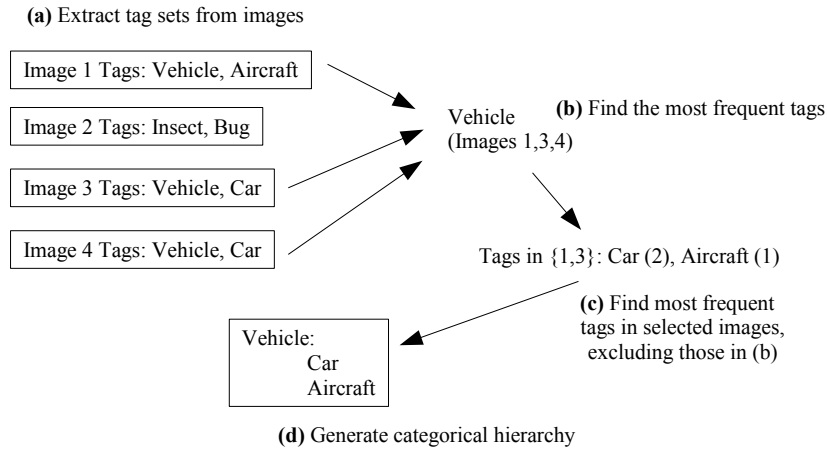


Figure 10: **How the Picbreeder category hierarchy is automatically generated.** The tags are extracted from the published images (a) to find the most frequent tags (b). For images with a particular frequent tag, the tags other than the one currently in focus make the next level of the hierarchy (c), and so on.




Figure 11: **Tree-of-life browser.** Clicking on an image in the interface shown above expands its children. Interactively browsing such a tree helps users to understand how artistic motifs are carried throughout generations.

exploration). For example, all published images are retained independently, providing an opportunity for users to seek out less popular images. Tagging and searching helps users find images that appeal more narrowly. Also, the Picbreeder front page contains a *random* view that shows eight random images on every page load. This view increases the odds that a user will choose to evolve a low-rated or unrated image.

4 System Architecture

Figure 12 illustrates the architecture of Picbreeder, which is a database-driven website. The database stores meta-data about the images including lineage, authorship, ratings, and tags. Images and the CPPNs that generate them are stored on the server filesystem.

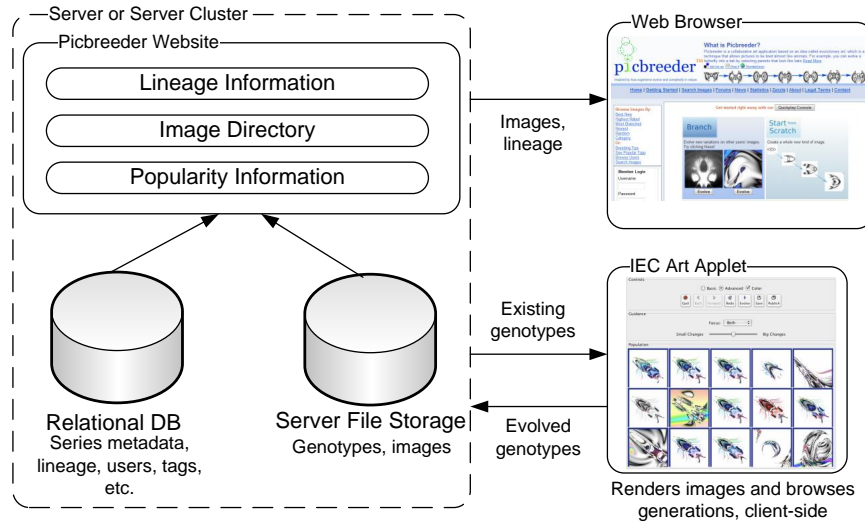


Figure 12: **System architecture.** The main site, including image views and lineage, is viewed through a web browser. The user evolves images through a Java client that sends the results back to the server to be stored on the filesystem and in the database.

The basic functions of the website are implemented as Java web services using Apache Tomcat and Axis. These services wrap interaction with both the relational database and the genomes stored on the server filesystem. The interactive components of the website are PHP scripts that consume SOAP-formatted object information returned by Axis. The Java-based IEC client communicates to the server through web service calls. Java was chosen in part because the client performs the IEC process, including image rendering, on the user's local machine, thereby reducing the load on the server. When the user saves his or her image, the generating CPPN is transmitted to the server in XML format, and saved in the user's account. The user can then tag and publish the image. Picbreeder's service-oriented architecture has made it easy to develop new applications such as the tree-of-life (figure 11) and Image DNA (figure 5) browsers and has facilitated quick, efficient online interaction and dynamic storage.

5 The Picbreeder Online Experiment

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 7,500 publicly available images resulting from over 155,000 user image selections. The images were contributed both anonymously and by over 500 registered users. Statistics that analyze user influence in the text that follows concern registered users only; all other statistics refer to all images on the site. While images evolved by the authors seeded the site, and are therefore not possible to exclude from the reported statistics, they comprised 827 of the over-7,500 images discussed, and are thus significantly outnumbered by non-author contributions. The images published by all users are only the final products that are

willfully published by them; intermediate images in the evolutionary progression and series abandoned by users (e.g. because they did not like the outcome) are not recorded in these statistics. Picbreeder is driven by community interaction: Each active user influences an average of 2.30 other users ($sd = 7.45$, median = 0) through branching and 88.1% of the images are branched from another image.

Statistics collected from Picbreeder suggest that branching and showcasing images allowed the design space to be explored more effectively than a single user could achieve on his or her own. While published images are evolved for an average of 20.31 generations ($sd = 29.94$, median = 11) during a single user session, each image has an average of 151.16 cumulative generations ($sd = 108.77$, median = 135) from an average of 9.75 ancestor series ($sd = 9.18$, median = 7). Because of NEAT's complexification, more cumulative generations yields more structure (i.e. nodes and connections, with correlation coefficients of 0.9462 and 0.9664 respectively) and therefore a greater potential complexity for the images that the evolved CPPNs can represent.

Recall that the *top rated* category includes all images with three or more votes and an average rating of 3.0 or higher. Of the 763 series in the top-rated category, only 1.97% of them were evolved within the 20 generation limit 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), and therefore were unlikely to have resulted from a traditional single-user process. Thus branching is facilitating a deeper exploration of the design space than would otherwise be possible.

Images in the top rated category had an average of 0.37 children who were themselves in the top rated category, compared to an average of 0.09 in the overall population. This difference is significant at $p < 0.001$ according to Student's *t*-test. This fact justifies the promotion of promising designs. However, most images in the top-rated category did not come from top-rated parents, meaning that all images should be retained for the potential to yield highly-rated children. In the Picbreeder community, there is also a correlation between recognition and participation. The ten users who published the most images, while making up about 2% of the user population, published over half of the top rated images.

Ratings also helped encourage participation. While the average image in Picbreeder is branched by 0.27 users ($sd = 0.78$, median = 0), images in the top rated category were branched by an average of 1.09 users ($sd = 1.81$, median = 1), which is significantly more at $p < 0.001$ according to the Student's *t*-test. The rating system is actively used by the community; the average published image is rated 1.37 times ($sd = 1.60$, median = 1) and the average top-rated image is rated 4.41 times ($sd = 2.43$, median = 4).

Tagging is also widely applied: The average published image is tagged 1.45 times ($sd = 0.77$, median = 1). The ten most frequent tags ("Face," "Eye," "Alien," "Creature," "Fish," "Bird," "Cool Pattern," "Dog," "Animal," and "Ghost") include 1,085 (14.2%) of the published images, indicating that the Picbreeder tagging system effectively organizes frequent categories.

A varied sampling of the designs published by Picbreeder users is shown in figure 13. The images represent the gamut of subject matter: faces, animals, structures, landscapes, astronomical bodies and many others that are familiar. While the value of these images is subjective, it is clear that they are markedly different than the random images from CPPNs of similar complexity in figure 3b. Thus complex, recognizable, and, by the community's judgment, valuable images can be evolved without depending on explicit design by instead leveraging the actions of the group.

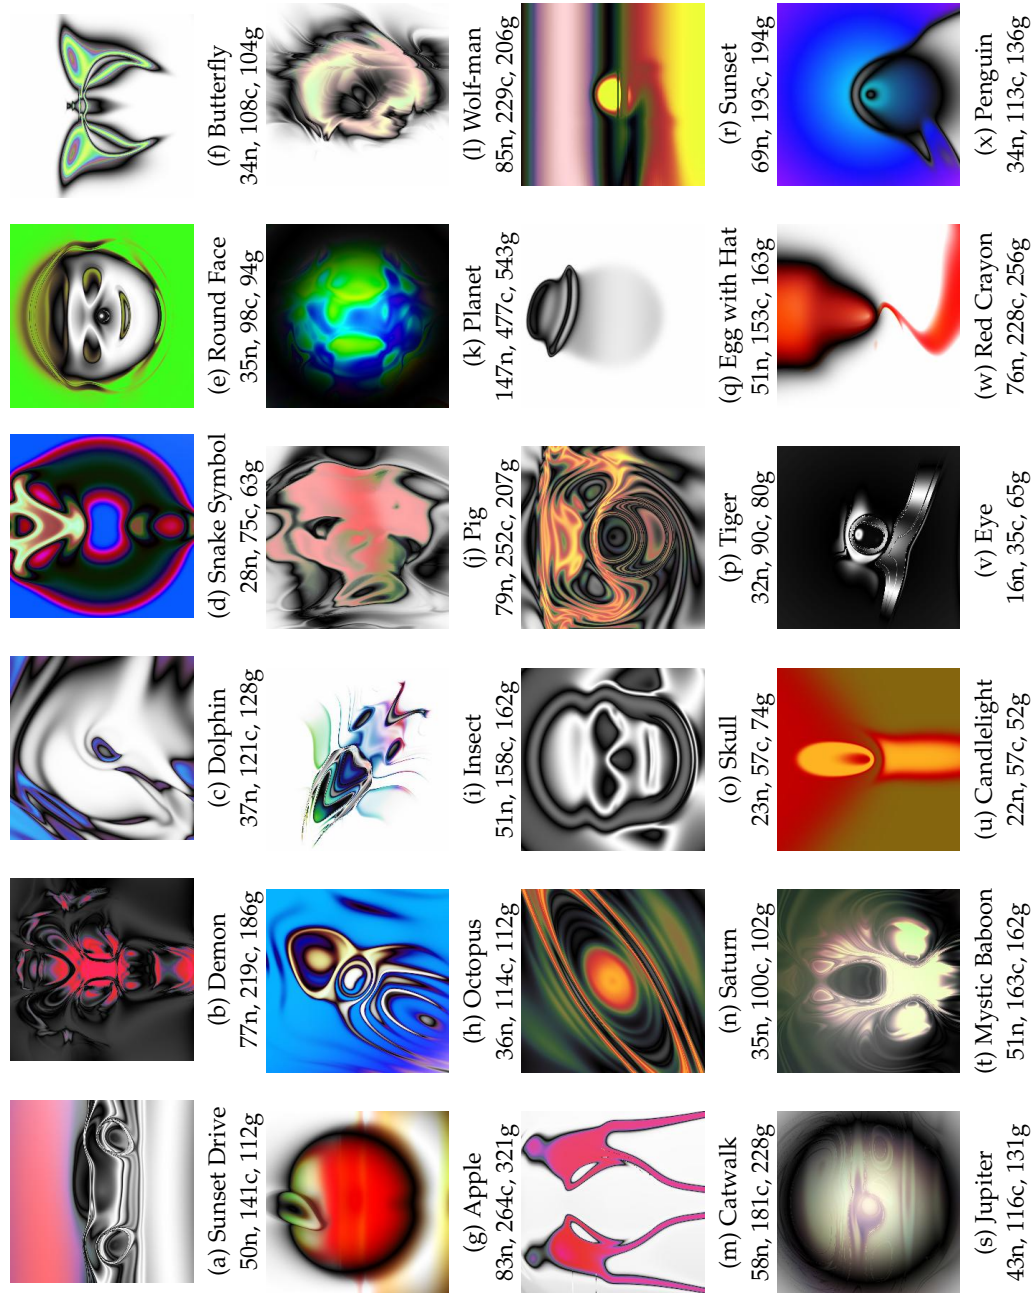


Figure 13: **Example evolved designs.** Example images created completely by Picbreeder's collaborative evolutionary process are shown with their associated number of nodes n , connections c and cumulative generations g . These images embody a variety of themes including mechanical (a), animal (c,d,f,h,i,j,p,x), humanoid (b,e,l,m,o), landscape (r), astronomical (k,n,s), and surreal (q,v).

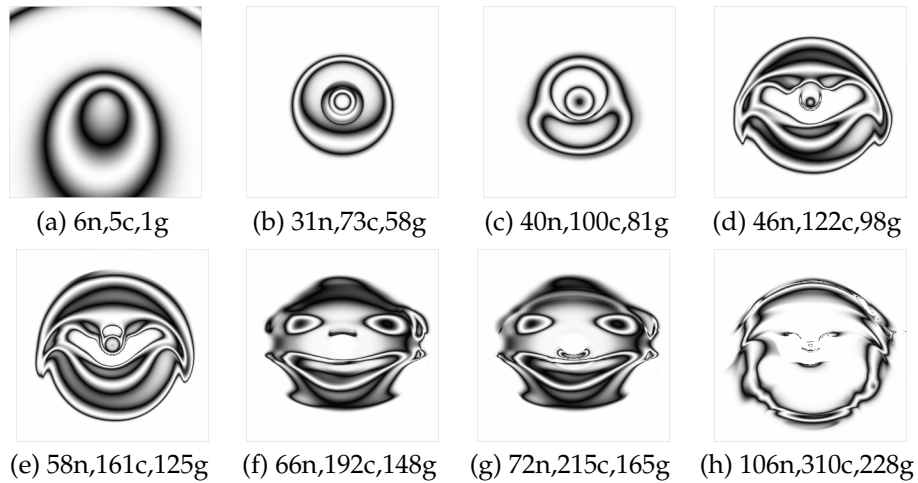


Figure 14: **Sequence of Descendant Face Patterns** The chronological sequence (a)-(h) displays successive progeny evolved in the Picbreeder client. The number of functions and connections in the generating CPPN is shown below each pattern, followed by the number of generations it took to evolve. The sequence exhibits a continual elaboration of regular form.

Picbreeder's CPPN representation and IEC client allow users to evolve a succession of elaborations on the same theme, as shown in the sequence of faces (each one the representative of one series in the chain) in figure 14. The images gradually become more elaborate and recognizable as a face as evolution progresses. This sequence demonstrates elaboration through complexification, acquiring new features while preserving the properties of previous generations.

Interestingly, despite no explicit design communication among users, many synergistic progressions among users emerged. Figure 15 depicts a phylogenetic tree created through the collaboration of 30 Picbreeder users. Each image represents a unique published series, which is separated from its parent by one or more generations of evolution. The tree reveals how the vast search space of images can be navigated with the help of human guidance, without requiring human expertise. 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, facilitating collaboration through the design itself.

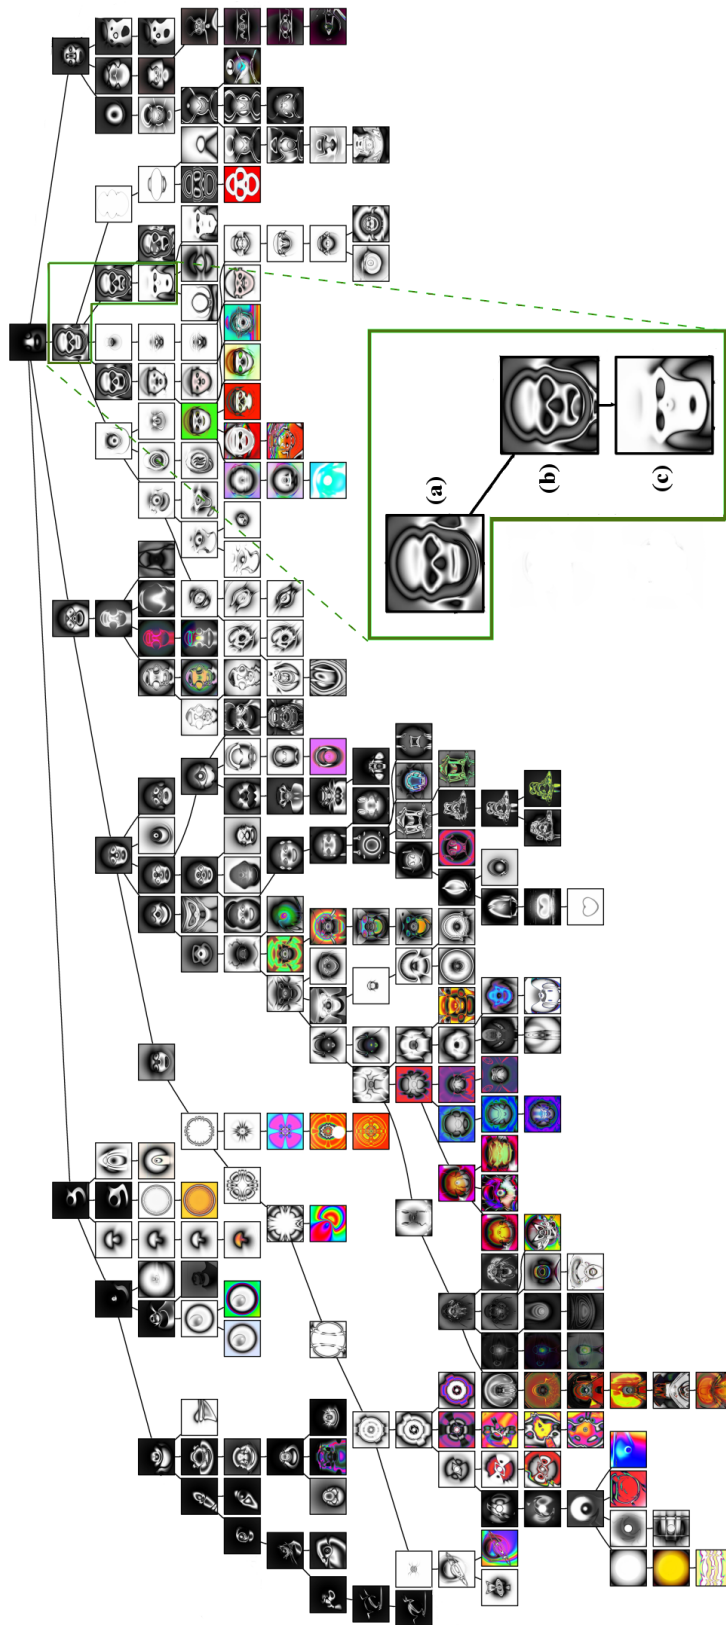


Figure 15: **A phylogeny of design.** This tree depicts the efforts of 30 different users. Each image represents a branching action from its parent. Note that the parent and the child are typically separated by many generations. The tree reveals how Picbreeder serves as an engine to proliferate new designs by building and expanding on the structure of other designs. The inset shows a minor change, from (a) to (b), and a larger one, from (b) to (c).

6 Discussion

With over 7,500 published image series contributed to the site, Picbreeder is among the largest online collections of evolved content. The collected statistics show that a Picbreeder-style approach to interactive evolution can increase the number of generations that content is evolved. It also produces images that users perceive as high in quality as shown through ratings. The images evolved through Picbreeder embody a variety of styles and motifs, reaffirming the CPPN as a flexible representation and NEAT as a capable evolutionary algorithm. Several users have evolved extensive collections of images, and have thereby enjoyed a novel method of creative collaboration. Finally, Picbreeder's techniques for highlighting noteworthy images and users have encouraged participation, while simultaneously promoting fruitful evolutionary lineages. However, even with these successes, there are several remaining challenges that the Picbreeder experiment uncovers. An analysis of its successes and failures reveals important lessons for the next generation of IEC systems.

To begin, simply comparing the art from Picbreeder with well-known classical art suggests that the evolved art is still categorically different. It is an interesting question whether such a difference signifies an artistic gap, or whether it introduces valuable novelty in its own right. This question will likely be the subject of significant discussion in the years to come. In addition, several existing interactive evolutionary systems, most notably GP-based ones (Sims, 1993, 1997; Rooke, 2002), have already generated images that may be considered as attractive or more attractive than the ones generated in Picbreeder. By employing complex aesthetic functions (e.g. gradients, fractals, and blurs) into their basic function sets, GP-based methods can accelerate the search for appealing images by biasing the search space. Future Picbreeder-like systems may employ GP (e.g. Teller 1999), expression trees and complex base functions as well. However, while some IEC systems may achieve abstract beauty, many Picbreeder pictures do still achieve the unusual distinction in IEC of resembling something recognizable despite little bias in the CPPN representation. One possible reason that Picbreeder images have yet to express the complexity and variety of traditional art is that most of the CPPNs evolved in Picbreeder still have relatively simple structure; the average published image is composed of 50.26 nodes and 146.70 connections. In principle, the CPPNs in Picbreeder may need orders of magnitude more complexity to express the nuance of traditional art.

Yet many generations and high CPPN complexity are not enough to ensure interesting images. There is surprisingly little relationship between the total number of generations evolved and how much the community appreciates the image: The cumulative number of generations evolved and the average rating given to an image (only for images which have been rated at least once) are correlated with a 0.088 coefficient. Nor is the image structure a significant predictor of ratings: Nodes and connections are correlated with ratings with coefficients of 0.12 and 0.11 respectively. Focusing instead on the number of unique users branching an image does not tell a significantly different story. Taking a sample of images published between July 1, 2009 and August 1, 2009 (to control for older images having more time to be branched) the number of unique users who branched the image by August 1, 2010 is correlated with cumulative generations with a coefficient of -0.02 (and with nodes and connections both with coefficients of -0.03). If one assumes that more generations should yield higher quality images, this fact appears counterintuitive. However, analogies from the traditional art world may aid in understanding the lack of correlation. Though the traditional arts can be thought to "evolve" over time, numerous classics are still appreciated (e.g. the

Mona Lisa) as much or more than further “evolved” contemporary art (Mohen et al., 2006). Audiences thus may appreciate the classics while at the same time look for new and challenging artistic directions.

A significant challenge that Picbreeder faced early is that visitors spent too little time on the site. Potential users were visiting the site in sufficient number, but were not staying long enough to evolve anything (i.e. 30 seconds or less). One possible explanation is that visitors to the site could not understand what could be done on the site that was interesting or valuable. To provide more immediate gratification for visitors, a “Quick Start” function was enabled to allow the user to start evolving right away and *anonymously* publish an image, thereby avoiding the sign-up process.

Of all of the published images, 84.5% were published by registered users and 15.5% were contributed anonymously. It is better to receive anonymous contributions than none at all. However, if users are hesitant to sign up, and choose to publish images infrequently and anonymously to the site, it can undermine the community-based mechanisms designed to encourage collaboration and participation. Users who are not registered and familiar with the community may not derive the same inspiration from interest generated by their evolved images. Long-term user participation and a sense of community remains a challenge: To date, slightly over 500 users have registered. Yet online communities with orders of magnitude more users are not uncommon. Unfortunately, most visitors to the site are merely spectators. It may be that a system like Picbreeder would need a larger core of users to form a critical mass of collaborative creativity.

For some users, Picbreeder is an enjoyable pastime. As an example, one user has published over 2,500 images. By contrast, the average user published 26.52 images (17.47 images when excluding those of the most prolific user from the average). In the absence of compensation, entertainment is clearly a vital aspect of the system. Exploring a design space can be fun and rewarding. Picbreeder enables users who would traditionally be left out of a creative process to become an active part of a design community. However, to ensure continued and significant contributions, Picbreeder and similar systems must prove their worth as a viable pastime to a larger audience. This challenge is essential for CIE to grow in relevance.

7 Future Directions

In the short term, future work will focus on the challenges that remain in the field of IEC. The mechanisms for community interaction should continue to improve to attract larger audiences. For example, future such systems may allow users to save individuals to their own computers and pass them through e-mail. Additional analysis can also help to clarify further users’ preferences for complexity versus simplicity.

The over-7,500 image repository that has accumulated on Picbreeder presents interesting future educational possibilities. In combination with the Image DNA viewer (figure 5), the repository can become a unique teaching tool for students studying how information is encoded in networks such as CPPNs or ANNs, and how such structures evolve. It may even provide intuition to early students of biology by explicitly depicting such concepts as gene conservation in some images and diversification in others.

More broadly, an important implication of the ability of Picbreeder users to collaboratively discover needles in the design space’s haystack is that other types of artifacts may also be possible to evolve collaboratively through a system like Picbreeder. Any digital artifact that can potentially be evolved through an IEC process, including music, product designs, architectures, synthesized voices, and intelligent agents can be

adapted within the Picbreeder framework. Furthermore, it is possible that designs created by experts can be converted to an evolvable form and injected into the system as a starting platform for evolution by an unskilled online community. Picbreeder-like systems can also potentially augment creativity in a traditional design process by providing a starting point from which an expert designer can develop ideas. Picbreeder provides the necessary precedent to show that such applications can actually work.

This technology also has commercial potential. In the future, a car, clothing, or furniture company might deploy a Picbreeder-like system to evolve new designs, potentially manufacturing the most popular ones. In the near term, Picbreeder could realistically evolve patterns that could be painted onto any of the aforementioned objects. As the popularity of personal rapid prototyping machines increases (Lipson and Pollack, 2000; Malone and Lipson, 2007), Picbreeder-like systems may evolve three-dimensional objects that can be downloaded and automatically built. The potential for such commercial applications is an important real-world incentive for the field of CIE as a whole. Yet for such applications to be realized, large numbers of users must be attracted and rewarded for their contributions. It is possible that combining CIE with online catalogs (e.g. clothing or furniture) will excite end-users, but this effect remains to be demonstrated commercially.

8 Conclusions

Picbreeder is an online system for collaborative interactive evolution. By allowing users to save their evolved images back to the server for other users to branch and continue, Picbreeder facilitated the creation of a large and diverse online repository of evolved content. As a result of branching, images can evolve for several times more generations than is possible through single-user IEC. Picbreeder also employs a number of mechanisms to encourage users to participate in the evolutionary process and contribute to the creative community. Though there is a long way to go before IEC systems can rival traditional design, Picbreeder shows that groups of users can evolve significant content, and can creatively interact in a new way.

References

- (2006–2007). Living Image Project. <http://www.w-shadow.com/li/>.
- (2007). Apophysis. <http://www.apophysis.org>.
- (2007). Wikipedia. <http://www.wikipedia.org/>.
- (2008). Sourceforge. <http://sourceforge.net/>.
- Aaltonen, T. et al. (2009). Measurement of the top quark mass with dilepton events selected using neuroevolution at CDF. *Physical Review Letters*. To appear.
- Amores, A., Force, A., Yan, Y.-L., Joly, L., Amemiya, C., Fritz, A., Ho, R. K., Langeland, J., Prince, V., Wang, Y.-L., Westerfield, M., Ekker, M., and Postlethwait, J. H. (1998). Zebrafish HOX Clusters and Vertebrate Genome Evolution. *Science*, 282:1711–1784.
- Angeline, P. J. (1995). Morphogenic Evolutionary Computations: Introduction, Issues and Examples. In McDonnell, J. R., Reynolds, R. G., and Fogel, D. B., editors, *Evolutionary Programming IV: The Fourth Annual Conference on Evolutionary Programming*, pages 387–401, San Diego, CA. MIT Press.
- Balogh, J., Dubbin, G., Do, M., and 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.

- Baluja, S., Pomerleau, D., and Jochem, T. (1994). Towards Automated Artificial Evolution for Computer-Generated Images. *Connection Science*, 6(2 and 3):325–354.
- Bentley, P. J. and Kumar, S. (1999). The Ways to Grow Designs: A Comparison of Embryogenies for an Evolutionary Design Problem. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-1999)*, pages 35–43, Orlando, FL. Kaufmann.
- Biles, A. (2010). Genjam. <http://www.ist.rit.edu/~jab/GenJam.html>.
- Biles, J. (1994). Genjam: A Genetic Algorithm for Generating Jazz Solos. In *International Computer Music Conference*, pages 131–137, Aarhus, Denmark.
- Bongard, J. C. and Pfeifer, R. (2001). Repeated Structure and Dissociation of Genotypic and Phenotypic Complexity in Artificial Ontogeny. In Spector, L., Goodman, E. D., Wu, A., Langdon, W. B., Voigt, H.-M., Gen, M., Sen, S., Dorigo, M., Pezeshk, S., Garzon, M. H., and Burke, E., editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 829–836. San Francisco: Kaufmann.
- Brown, D. C. and Birmingham, W. P. (1997). Understanding the Nature of Design. *IEEE Expert*.
- Caldwell, C. and Johnston, V. S. (1991). Tracking a criminal suspect through "Face-Space" with a genetic algorithm. In *Proceedings of Fourth International Conference on Genetic Algorithms*, pages 416–421, San Diego, CA.
- Carroll, S. B. (1995). Homeotic Genes and the Evolution of Arthropods and Chordates. *Nature*, 376:479–485.
- Carstensen, P. and Schmidt, K. (1999). Computer supported cooperative work: New challenges to systems design. In Itoh, K., editor, *Handbook of human factors*.
- Collins, N. (2002). Interactive Evolution of Breakbeat Cut Sequences. In *Proceedings of Cyberonica, Institute of Contemporary Arts*, London, England.
- 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*. Longman, Essex, U.K.
- Dawkins, R. (1989). The Evolution of Evolvability. In Langton, C. G., editor, *Artificial Life*, pages 201–220. Addison-Wesley, Reading, MA.
- Draves, S. (2005). The Electric Sheep Screen-Saver: A Case Study in Aesthetic Evolution. In *Proceedings of EvoMUSART*, Lausanne, Switzerland.
- Fagerlund, M. (2003–2006). DelphiNEAT homepage. <http://www.cambrianlabs.com/mattias/DelphiNEAT/>.
- Fagerlund, M. (2005). DelphiNEAT-based genetic art homepage. <http://www.cambrianlabs.com/mattias/GeneticArt/>.
- Ferstl, H. (2006). SharpNEAT-based genetic art homepage. <http://www.cs.ucf.edu/~kstanley/GenArt.zip>.
- Goldschmidt, G. and Porter, W. L., editors (2004). *Design Representation*. Springer.
- Graham, P. (2004). *Hackers & Painters: Big ideas from the computer age*. O'Reilly, Sebastopol, California.
- Greenfield, G. R. (2000). Evolving Expressions and Art by Choice. *Leonardo*, 33(2):93–99.
- Gruau, F., Whitley, D., and Pyeatt, L. (1996). A Comparison Between Cellular Encoding and Direct Encoding for Genetic Neural Networks. In Koza, J. R., Goldberg, D. E., Fogel, D. B., and Riolo, R. L., editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 81–89, Stanford University. MIT Press.

- Grudin, J. (1994a). Computer-Supported Cooperative Work: Its History and Participation. *Computer*, 27(4):19–26.
- Grudin, J. (1994b). Groupware and Social Dynamics: Eight Challenges for Developers. *Communications of the ACM*, 37(1):93–104.
- Hammond, S. (2007). Imagebreeder. <http://www.imagebreeder.com/>.
- 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*, LNCS 4448, pages 527–536, Valencia, Spain. Springer-Verlag.
- Hart, W. E., Kammeyer, T. E., and Belew, R. K. (1994). The Role of Development in Genetic Algorithms. Technical Report CS94-394, University of California San Diego, San Diego.
- Hastings, E., Guha, R., and 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.
- Hoover, A. and 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.
- Hornby, G. S. and Pollack, J. B. (2002). Creating High-Level Components with a Generative Representation for Body-Brain Evolution. *Artificial Life*, 8(3):223–246.
- Husbands, P., Gerny, G., McIlhagga, M., and Ives, R. (1996). Two applications of genetic algorithms to component design. *Evolutionary Computing. LNCS 1143*, pages 50–61.
- Johanson, B. and Poli, R. (1998). GP-music: An interactive genetic programming system for music generation with automated fitness raters. *Proceedings of the Third Annual Conference: Genetic Programming*, pages 181–186.
- Jorda, S. (1999). Faust music on line: An approach to real-time collective composition. *Leonardo Music Journal*, 9(5):5–12.
- Kohl, N., Stanley, K., Miikkulainen, R., Samples, M., and Sherony, R. (2006). Evolving a Real-World Vehicle Warning System. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2006)*, pages 1681–1688, Seattle, WA.
- Krause, K. (1996). Kai’s Power Tools 3. [Software].
- Langdon, W. B. (2005). Pfeiffer – A Distributed Open-ended Evolutionary System. In Edmonds, B., Gilbert, N., Gustafson, S., Hales, D., and Krasnogor, N., editors, *AISB’05: Proceedings of the Joint Symposium on Socially Inspired Computing (METAS 2005)*, pages 7–13, Hatfield, UK.
- Lee, J.-Y. and Cho, S.-B. (1999). Incorporating human preference into content-based image retrieval using interactive genetic algorithm. In *GECCO*, page 1788, Orlando, FL.
- Lipson, H. and Pollack, J. (2000). Automatic Design and Manufacture of Artificial Lifeforms. *Nature*, (406):974–978.
- Lund, H. H., Pagliarini, L., and Miglino, P. (1995). Artistic Design with GA and NN. In *Proceedings of the 1st Nordic Workshop on Genetic Algorithms and Their Applications (1NWGA)*, pages 97–105, Vaasa, Finland.
- Machado, P. and Cardoso, A. (2002). All the truth about NEvAr. *Applied Intelligence Special Issue on Creative Systems*, 16(2):101–118.
- Malone, E. and Lipson, H. (2007). Fab@Home: The Personal Desktop Fabricator Kit. *Rapid Prototyping Journal*, 13(4):245–255.
- Martin, A. P. (1999). Increasing Genomic Complexity by Gene Duplication and the Origin of Vertebrates. *The American Naturalist*, 154(2):111–128.

- McCormack, J. (2008). *Facing the Future: Evolutionary Possibilities for Human-Machine Creativity*, pages 417–451. Springer.
- McCormack, J. P. (1993). Interactive Evolution of L-system Grammars for Computer Graphics Modelling. In Green, D. G. and Bossomaier, T., editors, *Complex Systems: From Biology to Computation*, pages 118–130. IOS Press, Amsterdam.
- Mohen, J.-P., Menu, M., and Mottin, B. (2006). *Mona Lisa: Inside the Painting*. Harry N. Abrams, Inc.
- Nelson, G. L. (1993). Sonomorphs: An application of genetic algorithms to growth and development of musical organisms. In *4th Biennial Art and Technology Symp.*, pages 155–169, New London, CT.
- Nishino, H., Takagi, H., Cho, S., and Utsumiya, K. (2001). A 3D Modeling System for Creative Design. In *Proceedings of the 15th International Conference on Information Networking*, pages 479–487, Beppu City, Oita, Japan. IEEE Press.
- Risi, S., Vanderbleek, S. D., Hughes, C. E., and Stanley, K. O. (2009). How Novelty Search Escapes the Deceptive Trap of Learning to Learn. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2009)*, Montreal, Canada. ACM Press.
- Romero, J. and Machado, P., editors (2007). *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*. Springer.
- Rooke, S. (2002). *Eons of Genetically Evolved Algorithmic Images*, chapter 13, pages 339–365. Morgan Kaufmann.
- Secretan, J., Beato, N., D’Ambrosio, D. B., Rodriguez, A., Campbell, A., and Stanley, K. O. (2008a). Picbreeder: Collaborative Interactive Evolution of Images. *Leonardo*, 41(1):98–99.
- Secretan, J., Beato, N., D’Ambrosio, D. B., Rodriguez, A., Campbell, A., and Stanley, K. O. (2008b). Picbreeder: Evolving Pictures Collaboratively Online. In *CHI ’08: Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1759–1768, Florence, Italy. ACM.
- Shneiderman, B. (2007). Creativity Support Tools: Accelerating Discovery and Innovation. *Communications of the ACM*, 50(12).
- Sims, K. (1991). Artificial Evolution for Computer Graphics. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH ’91)*, pages 319–328, Las Vegas, NV. ACM Press.
- 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>.
- Smith, J. R. (1991). Designing Biomorphs with an Interactive Genetic Algorithm. In Belew, R. K. and Booker, L. B., editors, *Proceedings of the 4th International Conference on Genetic Algorithms (ICGA-91)*, pages 535–538, San Diego, CA. Morgan Kaufmann.
- Stanley, K. O. (2006). Exploiting Regularity Without Development. In *Proceedings of the AAAI Fall Symposium on Developmental Systems*, Menlo Park, CA. AAAI Press.
- 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., Bryant, B. D., and Miikkulainen, R. (2005a). Evolving Neural Network Agents in the NERO Video Game. In *Proceedings of the IEEE 2005 Symposium on Computational Intelligence and Games*, Essex University, Colchester, Essex, UK.

- Stanley, K. O., Kohl, N., and Miikkulainen, R. (2005b). Neuroevolution of an Automobile Crash Warning System. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1977–1984, Washington, DC.
- Stanley, K. O. and Miikkulainen, R. (2002a). Efficient Reinforcement Learning through Evolving Neural Network Topologies. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2002)*, pages 569–577, New York, NY.
- Stanley, K. O. and Miikkulainen, R. (2002b). Evolving Neural Networks Through Augmenting Topologies. *Evolutionary Computation*, 10:99–127.
- Stanley, K. O. and Miikkulainen, R. (2003). A Taxonomy for Artificial Embryogeny. *Artificial Life*, 9(2):93–130.
- Stanley, K. O. and Miikkulainen, R. (2004a). Competitive coevolution through evolutionary complexification. *JAIR*, 21:63–100.
- Stanley, K. O. and Miikkulainen, R. (2004b). Evolving a Roving Eye for Go. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2004)*, pages 1226–1238, Seattle, WA. Springer Verlag.
- Żumlanski, S. R., Wu, A. S., and Hughes, C. E. (2006). Conflict Resolution and a Framework for Collaborative Interactive Evolution. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 512–517. AAAI Press.
- Takagi, H. (2001). Interactive Evolutionary Computation: Fusion of the Capacities of EC Optimization and Human Evaluation. *Proceedings of the IEEE*, 89(9):1275–1296.
- Taylor, M. E., Whiteson, S., and Stone, P. (2006). Comparing evolutionary and temporal difference methods in a reinforcement learning domain. In *GECCO 2006: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1321–1328, Seattle, WA.
- Teller, A. (1999). Neural Programming and Internal Reinforcement. In Angeline, P., O'Reilly, U., Langdon, W., and Spector, L., editors, *Advances in Genetic Programming III*. MIT Press.
- Todd, S. and Latham, W. (1992). *Evolutionary Art and Computers*. Academic Press, London.
- Tokui, N. and Iba, H. (2000). Music composition with interactive evolutionary computation. In *Proceedings of 3rd International Conference on Generative Art*, Milan, Italy.
- Trujillo, L., Olague, G., Lutton, E., and de Vega, F. (2008). Discovering several robot behaviors through speciation. *Lecture Notes in Computer Science*, 4974:164.
- Unemi, T. (1994). Genetic algorithms and computer graphic arts. *Journal of Japan Society for Artificial Intelligence*, 9(4):518–523.
- 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*, pages 288–291, Adelaide, SA.
- Viégas, F. B., Wattenberg, M., van Ham, F., Kriss, J., and McKeon, M. (2007). Many Eyes: A Site for Visualization at Internet Scale. In *Proceedings of Infovis*, Sacramento, CA.
- Watanabe, T. and Takagi, H. (1995). Recovering system of the distorted speech using interactive genetic algorithms. In *Systems, Man and Cybernetics, 1995. Intelligent Systems for the 21st Century., IEEE International Conference on*, volume 1, pages 684–689, Vancouver, BC.
- World, L. (1996). Aesthetic Selection: The Evolutionary Art of Steven Rooke. *IEEE Computer Graphics and Applications*, 16(1):4–5.
- Zigmond, M. J., Bloom, F. E., Landis, S. C., Roberts, J. L., and Squire, L. R., editors (1999). *Fundamental Neuroscience*. Academic Press, London.