

WebAL Comes of Age: A review of the first 21 years of Artificial Life on the Web

Tim Taylor

tim@tim-taylor.com

Joshua E. Auerbach

joshua.auerbach@epfl.ch

Josh Bongard

jbongard@uvm.edu

Jeff Clune

jeffclune@uwyo.edu

Simon Hickinbotham

simon.hickinbotham@york.ac.uk

Charles Ofria

ofria@msu.edu

Mizuki Oka

mizuki@cs.tsukuba.ac.jp

Sebastian Risi

sebr@itu.dk

Kenneth O. Stanley

kstanley@cs.ucf.edu

Jason Yosinski

yosinski@cs.cornell.edu

Abstract

We present a survey of the first 21 years of web-based Artificial Life (“WebAL”) research and applications, broadly construed to include the many different ways in which Artificial Life and Web technologies might intersect. Our survey covers the period from 1994—when the first WebAL work appeared—up to the present day, together with a brief discussion of relevant precursors. We examine recent projects, from the period 2010–2015, in greater detail in order to highlight the current state of the art. We follow the survey with a discussion of common themes and methodologies that can be observed in recent work and identify a number of likely directions for future work in this exciting area.

Keywords: WebAL, web technologies, HTML5, open science, crowdsourcing, crowdfunding, public outreach

1 Introduction

In recent years there has been a growing body of work in Artificial Life (ALife) that makes use of web technology in one way or another.

Over the last five years or so, web technologies have shifted away from proprietary browser plugins and towards standardized, native application programming interfaces (APIs) for providing graphics, animation, multimedia and other advanced features. This progress is due to the development and adoption of the HTML5 language and associated API specifications introduced by the World Wide Web Consortium (W3C).¹

This movement has made it much easier to develop and deploy rich web-based applications that work reliably and consistently on any browser, across multiple platforms and devices. It is therefore unsurprising that a number of high-profile ALife projects have emerged over this period that utilize web technology in various different ways. We refer to such work as *WebAL*, and broadly construe the field to include the multitude of ways in which ALife and the Web might intersect. Examples include the creation of massively distributed user-guided evolutionary systems, the creation of open science platforms, the use of web-based applications for public outreach and engagement, and the use of crowd-funding platforms for supporting the development of ALife systems. As we demonstrate in this review, and summarize in Section 7, there are many other points of intersection in addition to these.

In light of this emerging trend, an inaugural Workshop on Artificial Life and the Web (WebAL-1)² was held at the 14th International Conference on the Synthesis and Simulation of Living Systems (ALIFE 14) in New York City on 31 July 2014 [128]. Inspired by the success of the workshop, a number of its participants decided to collaborate on writing a comprehensive review of the field—the result of which is the current paper.³

Although recent years have witnessed a rapid growth in this area, the first web-based ALife systems date back 21 years to the mid-1990s, with various antecedents employing alternative forms of network technology dating back to the 1970s. In Figure 1 we present a timeline showing some of the main WebAL projects discussed in this paper plotted against developments in the underlying web technology. We will refer to this figure throughout the review.

The organization of the rest of the paper is as follows. In Section 2 we define the scope of this review and describe the methodology adopted to collect the materials upon which the review is based. In Section 3 we review various non-web or non-ALife projects that have heavily influenced the state of the art of modern-day WebAL. In Section 4 we cover the first true WebAL systems that appeared in the 1990s soon after the birth of the Web

¹See <http://www.w3.org/TR/html5/> and <https://en.wikipedia.org/wiki/HTML5>.

²<http://alife.org/ws/webal1>

³This paper is a vastly expanded and more detailed revision of a previous review paper [127], with additional contributions from many of the participants and organizers of the WebAL-1 Workshop.

itself, and in Section 5 we explore developments in the following decade, the 2000s. In Section 6 we look in slightly more detail at WebAL systems that have appeared since 2010, in order to present the current state of the art. On the basis of the work reviewed, we identify some of the important emerging themes in Section 7, and comment on likely directions for future research as well as projects in active development. Finally, we present our conclusions in Section 8.

2 Review Scope and Methodology

As described in Section 1, we have approached this review with an open-minded perspective of what constitutes WebAL, in order to examine the many different meeting points of ALife and the Web. Although such an approach allows us to address a variety of topics that might not have been covered in a more narrowly-focused review, there is a danger that it might lead us into vast areas of research where there is no hope of providing a comprehensive, coherent review that would be of interest to readers of this journal. This danger is compounded by the fact that the boundaries of ALife research are indeterminate, and the distinction between Web applications and other forms of Internet application is also becoming increasingly blurred (particularly in the case of mobile apps and social media). To guard against these dangers, we have taken a number of measures to enhance the focus of the review and to provide a clear methodology.

To determine what counts as ALife in this context, we considered work that describes itself as ALife, or is published in the ALife literature, as a *de facto* definition of the field. Specifically, we searched for relevant work among all papers ever published in publications affiliated with the International Society for Artificial Life—that is, every issue of the *Artificial Life* journal, and every proceedings volume of the *International Conference on Artificial Life* and the *European Conference on Artificial Life*.⁴ We acknowledge that other ALife-related journals and conferences exist, but we needed to draw a line somewhere in order to keep the task manageable. The initial search yielded approximately 450 potentially relevant papers, although many of these were spurious results. These papers were then investigated in more detail in order to determine which were of genuine relevance. In addition to this systematic search, we also tapped into the collective knowledge of the co-authors of this paper (all of whom are actively working in WebAL), as well as asking a number of other colleagues associated with the field for feedback on an early draft.

Furthermore, we have attempted to draw a distinction between work relating to the Web and work related more generally to the Internet. We use the terms *WebAL* and *NetAL*, respectively, for the narrower and broader fields. The major focus of this review is on WebAL, although it has not always been sensible to draw a clear distinction here. In cases

⁴Where electronic copies of papers were available, searches were performed using the following keywords: (*web, internet, Java, applet, HTML, Flash, browser, Android, iOS, screensaver, JavaScript*). Where no electronic copies were available, abstracts were read to determine potential relevance.




Figure 1: A timeline of selected WebAL projects in relation to developments in web technology. The color of dot for each project indicates the main technology used (green: basic HTML and CGI scripts; blue: JavaScript; orange: Java; red: Flash; black: Other).

where we feel that specific NetAL work is especially relevant, we have included it in our review. In particular, we highlight some of the most relevant NetAL work in Box 1 and Box 2. For completeness, any other NetAL work that was uncovered in our systematic search of the *Artificial Life* journal and conference proceedings is listed in an annotated bibliography (Appendix A).

Having established our criteria for what work to review, there are particular aspects of this body of work that we have chosen to focus on. In terms of current work, we are primarily concerned with modern HTML5 and associated APIs as a *platform*, an *environment*, and an *enabler* for ALife. We are especially interested in how the modern Web enables *new* ways of working that were previously not possible or feasible—at the end of the review we highlight the most important of these in Section 7. In our review of earlier work, we are primarily interested in pioneering work that used ALife and the Web in novel ways, and in tracing the important antecedents to today’s WebAL projects. There are other areas of work that we have chosen not to focus on because they are large fields which would have taken us too far from our core topics, and because comprehensive reviews of those areas are available elsewhere. These include the general areas of web crawling and information retrieval, distributed multi-agent systems, and frameworks for distributed evolutionary algorithms. We do, however, discuss specific work in these fields that have a particular ALife-oriented focus, and we provide pointers to broader reviews of these areas where appropriate.

3 The Precursors to WebAL

Many of the WebAL projects discussed in later sections involve the use of network technology and distributed human interaction to produce some kind of emergent artifact. Seen in this general light, a wide variety of potentially relevant antecedents can be found. We cannot provide an exhaustive review of such a broad body of prior work, but instead highlight a few significant landmarks and waypoints.

The arts world provides many interesting studies of the combination of network technology, communication, emergent behavior, and inhabited virtual worlds. In 1977, artists Kit Galloway and Sherrie Rabinowitz exhibited the *Satellite Arts* project, with funding and technical support from NASA [26, 84]. The project was a live performance piece involving two sets of dancers, located on opposite coasts of the USA, connected by a satellite-linked video feed that blended the images of both sets of dancers to produce a live, shared virtual space in which the performance happened. The project was one of the first examples of the use of network (in this case, satellite) technology to explore distinctions such as real/virtual, subject/object and here/there. In the years since the *Satellite Arts* project, the arts world has continued to push the boundaries of technology to explore issues of networks, communication, virtuality and emergent behaviour.⁵

⁵To briefly mention just two more recent examples: Karlheinz Stockhausen’s *Helikopter-Streichquartett*

Besides the arts world, computer games have an equally long history of development of shared virtual worlds. Real-time, shared virtual worlds date back to the development of the original MUD (Multi User Dungeon) by Roy Trubshaw in 1978 [83]. MUD initially allowed multi-user play via ARPANet before being licensed to CompuServe,⁶ where it ran until 1999. The MUD source code has recently been acquired by Stanford University Library as part of an ongoing effort to preserve virtual worlds [58]. MUD spawned a diverse and distinguished lineage of massively multiplayer online games (MMOs) that still thrives today [83].

The 1970s witnessed the emergence of computer viruses on mainframes, and by the early 1980s viruses and worms were also appearing on microcomputers [125]. One of the first examples of a worm that spread via the Internet, causing widespread damage and attracting the attention of the mainstream media, was Robert Morris' *Internet Worm of November 1988* [23]. Referring to the growth of computer viruses, Christopher Langton warned in his introduction to the proceedings of the Second Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems (*Artificial Life II*, 1990):⁷

“There is a caution here that we all must attend to. Attempts to create Artificial Life may be pursued for the highest scientific and intellectual goals, but they may have devastating consequences in the real world, if researchers do not take care to insure that the products of their research cannot ‘escape’, either into computer networks or into the biosphere itself.” [50, p.18]

Although the general history of computer viruses is beyond the scope of this review, we will refer back to Langton's warning when we discuss recent developments with client-centric WebAL in Section 7.7. For further discussion of the broader topic of computer viruses and their relation to ALife, we refer the interested reader to Eugene Spafford's review paper [114].

Relevant antecedents to WebAL can also be found in the development of distributed evolutionary systems, with theoretical work on parallel genetic algorithms starting in the 1960s and implementations in the 1980s—see [14] for a good review. One of the most ambitious examples of a distributed evolutionary system is Karl Sims' *Evolved Virtual Creatures* [110], released in 1994, which ran on a Connection Machine CM-5 supercomputer across 1024 cores. In this simulation, Sims fully evolved the morphology and behavior of a population of virtual organisms in a 3D world, inspiring many subsequent systems. Furthermore, Sims' gallery installations of interactive evolutionary art, such

(Helicopter String Quartet), first performed in 1995, involved a coordinated live performance of four musicians each flying in separate helicopters [122]. Elsewhere, Matthew Fuller's group performance project *The Human Cellular Automaton*, first performed in 2000, implemented cellular automata rules (such as Conway's Game of Life) in a crowd of human participants [133, pp. 173–174], thereby creating a distributed human-powered computer.

⁶<http://british-legends.com/CMS/index.php/about-mud1-bl/history>

⁷Similar concerns were also raised by Harold Thimbleby, who presented at the Artificial Life II workshop but did not appear in the proceedings; his ideas were later published elsewhere [129].

as *Genetic Images* (1993)⁸ and *Galápagos* (1997),⁹ were important precursors for WebAL evolutionary art systems such as *Picbreeder* (discussed in Section 5.1.1) and later projects (Section 6.1).

Huge advances have been made in the more general field of distributed computing over the last couple of decades. The WebAL projects discussed in later sections have developed in the context of noteworthy successes in several major scientific projects using distributed computing and Internet crowdsourcing, including *SETI@Home* [5], *Foldit* [46], and *Galaxy Zoo* [56]. The continued advances in large-scale distributed systems over this period have also stimulated the growth of the new field of autonomic computing, which aims to use bio-inspired techniques to produce large-scale, self-managing distributed I.T. systems that are substantially more robust than traditional complex computing systems [45].

In this section we have only briefly touched upon some of the most relevant precursors to WebAL. By the early 1990s, the technological milieu was pregnant with many of the ideas outlined above. As the Internet matured and the World Wide Web was born, such ideas were invigorated by the development of easy-to-use, standardized network technology and the mushrooming uptake of Internet and Web technology not just by special interest groups but by the general public the world over. In the next section, we look at the earliest examples of work that truly deserves the label WebAL.

4 The 1990s: Early WebAL

One of the first examples of WebAL was Michael Witbrock and Scott Neil-Reilly's evolutionary art project *International Interactive Genetic Art 1 (IIGA1)*, developed in 1994, along with the subsequent IIGA2 system developed by John Mount [134].¹⁰ In this work, the fitness of images generated by genetic programming representations was determined by taking the average ratings for each image as scored by users accessing the system via the Web. Both IIGA1 and IIGA2 attracted tens of thousands of users, and the work was mentioned in *Wired* magazine [115]. The IIGA2 site, built upon CGI scripts and HTML forms,¹¹ racked up nearly 4000 generations (each of 10 images) on its original server, with over 115,000 page views of the images. The use of the Web as a means of collecting distributed aesthetic evaluation of 2D and 3D images in evolutionary art systems remains a strong component of WebAL work to this day, as will be highlighted in later sections.¹²

⁸<http://www.karlsims.com/genetic-images.html>. Like *Evolved Virtual Creatures*, this work also ran on a massively parallel Connection Machine supercomputer.

⁹<http://www.karlsims.com/galapagos/>

¹⁰See also <http://www.win-vector.com/blog/2009/06/what-is-genetic-art/>.

¹¹<http://www.mzlabs.com/MZLabsJM/page4/page4.html>

¹²The work on web-based evolutionary art systems reported here is part of a larger field of work on evolutionary art, much of which is not web based. A good review of the wider field can be found in [52].

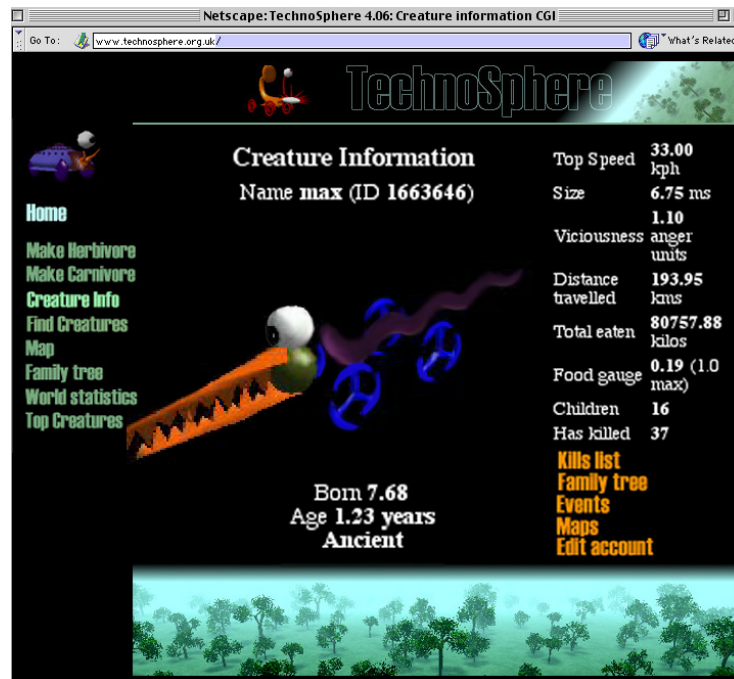


Figure 2: The *TechnoSphere* web interface.

Moving from 2D to 3D, the year 1995 saw the launch of the web-based artificial life virtual world *TechnoSphere*, created by Jane Prophet and Gordon Shelley [87]. The front-end of the system was a website where users could design their own creatures by selecting from a limited range of predesigned body parts (see Figure 2). Once created, the user submitted their creature to the web server, and it was tagged with the user's email address and a unique ID. Submitted creatures were released into a 3D virtual world (which was not rendered live on the website), featuring a fractally generated landscape and other creatures (many of which were designed by other users), subject to ecological rules that governed their interactions. At key events, the users would receive email updates. For example, when creatures interacted with each other, the email addresses of both authors were shared, to facilitate discussion. Users could even request “postcards” of their creatures, which were generated by rendering a scene showing the creature in its current location. In 1996 the *TechnoSphere* world reached a peak population of 90,000 creatures. In 1998, work started on a version with real-time 3D rendering [88], which was exhibited at a number of art galleries and museums over the period 1999–2001, although this version ran on a local network of PCs rather than on the Web.¹³

An early example of using the Web to mix real and virtual worlds—and thereby addressing some of the same topics investigated by the pre-Web *Satellite Arts* project described

¹³At the time of writing, a new version of *TechnoSphere*, in the form of an augmented reality mobile app, is currently under development (see Section 6.2.2).

Box 1: Highlights of other NetAL work from the 1990s

The *Network Tierra* project, initiated around 1995, was a networked version of Tom Ray's well known system *Tierra*, in which the digital organisms were native, self-reproducing code structures running in a virtual computer. The goal was to use the Internet to create a world-wide, distributed, complex environment in which the digital organisms could roam and freely evolve. Project participants ran specially designed servers which ran *Tierra* locally and managed network communications with other *Tierra* servers in a secure and contained manner.^a Over a period of 5 years, Ray and co-workers used *Network Tierra* to investigate the evolution of complexity in parallel programs (their analogy to multicellular organisms). Results were mixed, but the project ultimately failed to achieve an evolutionary increase in the number of differentiated parallel processes (cell types) [96].^b

A more applied variation of the ideas behind *Network Tierra* can be seen in Filippo Menczer and colleagues' 1995 paper on evolving agents for information retrieval on the Web [64]. This work proposed an algorithm where agents searched for documents driven by a user's query, gaining energy and reproducing dependent upon their success. The algorithm was tested on a small set of 116 connected web pages, although this early work did not employ actual on-line agents. A full review of the field of online information retrieval agents is beyond the scope of this paper, but a summary of work in this field in the 1990s can be found in [47].

Elsewhere, in 1996 David Ackley presented a description of his *ccr* system [1], which has been under development intermittently since 1989.^c *ccr* is an Internet-based software system that aimed to provide a large-scale, distributed common environment in which humans and artificial agents could communicate and interact. Some further analysis of the system was published in 2000 [2], and the system's key server, launched in 1996, is still running today.^d Although Ackley says that many of the lessons learned from the system have not yet been presented in the literature, the project has influenced his more recent research on indefinitely scalable computing [3, 4]

Another early example of the use of the Internet for distributed communication is provided by Luc Steels and colleagues' long-running *Talking Heads* project—initiated in 1999—which studied the development of language in communicating robots. In these experiments, physical robots were situated in various locations around the world, and virtual agents (i.e. the controllers for the robots) could “teleport” over the Internet to inhabit different physical bodies. The public could also interact with the agents via the project's web site [120, 121].

A final example from this period is the work of Johnson and colleagues in 1998 [42], which presented an early discussion of the potential of the Internet to facilitate large-scale collaborative problem solving and self-organization of knowledge. As discussed in Section 7, these concepts form a core thread of a significant portion of more recent WebAL work.

^a<http://life.ou.edu/tierra/netfaq.html>

^bSubsequent work with the *Avida* system was able to generate the spontaneous evolution of differentiated cell types, but these experiments were run on local clusters, not on the Web [33].

^c<http://keys.ccrcentral.net/ccr/history/>

^d<http://keys.ccrcentral.net/ccr/data/central-stats-v0.3+.html>



Figure 3: *Nerve Garden*. Composite image of the Nerve Garden Island with interface, butterfly in the sunshine and lightening storm and bee in flight over the island.

in Section 3—was *Telegarden*.¹⁴ Building upon their work on web-based teleoperated robotics [32], Ken Goldberg and colleagues created a system in which web users could tend to a remote real-world garden, using a teleoperated robotic arm to plant and water seedlings. *Telegarden* ran successfully for 10 years (1995–2004), amassing 9000 active users by August 1996¹⁵ and attracting coverage from many mainstream media channels. A sociological study of the community of users was reported in [62].

Bruce Damer and colleagues' *Nerve Garden*, launched in 1997, shared some of the ideas of *Telegarden*, but swapped the mixed-reality aspect for a shared 3D virtual environment [19]. The system allowed users to grow 3D plant models, generated by L-systems, on a client-side Java application. Users could then select their favorite plants, name them, and submit them to a central server where they would be planted in a shared VRML-based¹⁶ 3D environment that could be viewed by anyone with a browser suitably equipped with a VRML viewer plugin (see Figure 3).

By the mid-1990s, projects such as *Telegarden* were attracting growing attention in the popular media. In some cases, the rapidly expanding media interest in cyberculture even led to funding for new WebAL projects. Following the publication of Kevin Kelly's book *Out of Control* [44], Absolut Vodka (sponsors of Kelly's website) funded the development

¹⁴<http://goldberg.berkeley.edu/garden/>

¹⁵<http://web.archive.org/web/19980203215817/http://telegarden.aec.at/html/nyt.html>

¹⁶<http://www.w3.org/Markup/VRML/>

of a web-based genetic art system called *Absolut Kelly*, built upon the ideas of distributed intelligence and the “hive mind” discussed in the book. The system, developed by Jeffrey Ventrella, used an interactive genetic algorithm to allow users to evolve images featuring the distinctive outline of Absolut’s vodka bottles.¹⁷

Extending the hive mind approach to evolving art even further, and taking its name from the sci-fi novel *Do Androids Dream of Electric Sheep?* [24], Scott Draves’ *Electric Sheep* is a distributed, interactive artwork system launched in 1999 [25].¹⁸ The software exists in the form of a screensaver that can be downloaded by users and used to render frames of abstract artworks. Once complete, these frames are sent back to a central server, where they are used to generate animations (known as “sheep”). The animations are compressed by the server and sent back to the client-side screensavers for display (the system is therefore not web-based as such, but relies on Internet communications between the server and client-side screensavers). Users can vote for their favorite animations, with votes being used by a genetic algorithm to guide the future evolution of new sheep. The sheep are therefore the result of massively distributed computation and human participation from tens of thousands of users.

While genetic art projects such as *Electric Sheep* and the *Absolut Kelly* employed web users purely for the aesthetic selection step of the evolutionary process,¹⁹ the shared virtual environment projects such as *TechnoSphere* and *Nerve Garden* provided users with tools to directly manipulate the design of their contributed virtual inhabitants. In *TechnoSphere* this manipulation was accomplished by the composition of standard elementary parts, and in *Nerve Garden* by providing an interface to adjust various parameters of the L-system.

A somewhat more abstract approach, in which users entered written text in a web interface which was then treated as a genetic code and mapped to a 3D creature, was developed by Laurent Mignonneau and Christa Sommerer in a series of projects they developed in the late 1990s [112]. The first of these projects was *Life Spacies*, introduced in 1997 and followed by *Life Spacies II* in 1999 [69]. This project was an interaction environment installed in a museum in Tokyo and connected to a website through which users from all over the world could design virtual creatures by entering text messages that would then be mapped into 3D creatures and introduced into the environment displayed at the museum. Once uploaded to the shared environment, the creatures could interact with each other in the environment. Furthermore, visitors to the museum could interact with the creatures by touching them,²⁰ thereby blurring the boundary between

¹⁷Example images can be seen at <http://www.ventrella.com/Tweaks/Absolut/absolut.html>. The work is also described in [133, 45] and on Kevin Kelly’s website: <http://kk.org/ct2/2007/10/17/mutating-art-from-chaos/>.

¹⁸<http://electricsheep.org>. See also [133, pp. 154–157].

¹⁹Another early example of this kind was the *Artificial Painter* project, a genetic art system that included a limited web-based interface allowing user-guided aesthetic selection (archived at http://web.archive.org/web/19980503002715/http://gracco.irmkant.rm.cnr.it/luigi/lupa_ap.html) [81].

²⁰This interaction was achieved using a camera-tracking system that captured a visitor’s image and

real and virtual worlds (a recurring theme that has come up in various of the projects already discussed). A related web-based system, *Verbarium*, was also introduced in 1999, and allowed users to create shapes and forms in real-time using the same idea of a text-to-form encoding and an online interactive text editor [112, 113].

In addition to arts projects, the mid-1990s saw the birth of many WebAL-related projects in the fields of computer games and animation.

1996 saw the release of the ALife focused game *Creatures*, designed by Steve Grand.²¹ The characters in the game were digital life forms, called *Norns*, that were capable of lifetime learning, and possessed a physiology, drives, and communication abilities, all of which could evolve over generations. Although the first version of the game ran on standalone PCs, a growing online community of players soon started exchanging their *Norns* via enthusiast websites [16, 41]. By 2001, *Creatures Docking Station* was released, an Internet-based add-on to *Creatures 3* that allowed *Norns* to travel between different online worlds.²²

The growing popularity of the Java programming language in the mid-1990s (Figure 1), and the ease by which it allowed programs to be distributed via applets embedded in web pages, led to a flourishing of websites hosting ALife related applets. In 1996, Craig Reynolds ported his well-known *Boids* work, which had originally been published in 1987 [97], to a Java applet on his web site.²³ Various other *Boids* applets also appeared around the same time.²⁴ One particularly noteworthy example was *Floys* by Ariel Dolan,²⁵ which extended the basic *Boids* design by adding territorial behaviors to the creatures [80]. *Floys* (along with *Boids*) was described in a popular science article in *Scientific American* in 2000, which highlighted the potential for amateur scientists to modify the code and use it for their own investigations [15].

An example of a more extensive ALife-related Java application was developed by the British design group *Soda Creative* in 1998. Their system, *Soda Constructor*,²⁶ employed a 2D physics engine and presented users with an online editor with which they could construct creatures based upon mass-spring systems with oscillating muscles. By mid-2000, the popularity of the game had soared through “word of email”, and an online forum enabled users to share their creations.²⁷ *Soda Creative* won an Interactive Arts

projected it into the display, thereby integrating them into the virtual environment [69].

²¹[http://en.wikipedia.org/wiki/Creatures_\(artificial_life_series\)](http://en.wikipedia.org/wiki/Creatures_(artificial_life_series))

²²See http://creatures.wikia.com/wiki/Docking_Station. The information in this paragraph has been verified by personal communication with Steve Grand.

²³See <http://www.red3d.com/cwr/boids/applet/>. Date of original publication confirmed by personal communication with Craig Reynolds.

²⁴For example, Conrad Parker’s implementation <http://www.vergenet.net/~conrad/boids/>. Various other examples are listed at <http://www.aridolan.com/ad/Alife.html>.

²⁵<http://www.aridolan.com/ofiles/JavaFloys.html>

²⁶<http://soda.co.uk/work/sodaconstructor>

²⁷<http://www.acmi.net.au/soda.htm>

BAFTA Award in 2001 for their work.²⁸ In 2002, they teamed up with Queen Mary University London to develop *Sodarace*, a shared online environment where users from around the world could pit their creations against each other in competitions [63].²⁹ The development of *Sodarace* was supported by the UK's Engineering and Physical Sciences Research Council, and had a strong public outreach and educational flavor.³⁰

Elsewhere, an early example of work employing a distributed web-based genetic algorithm was reported in 1998 by Jens Astor and Christoph Adami [6] (with further experiments reported a couple of years later [7]). The work investigated the evolution of developmental artificial neural networks (ANNs), and utilized a server program written in Java that farmed out evaluations of individual ANNs to heterogeneous clients running a browser-based Java applet (or a local desktop Java program).³¹ The authors saw the potential of web-based distributed architectures for running computational tasks at a massive scale. However, their publications only reported results from experiments run over a local network; they cited Java's lack of speed as a problem, although they foresaw that improvements in Java compiler technology would make this kind of architecture more feasible in the future.³²

As demonstrated in the preceding discussion, by the late 1990s WebAL was already a fertile and vibrant field of research. A short review paper entitled “*ALife Meets Web: Lessons Learned*”, published in 1998, summarized some of the work discussed in this section and drew some conclusions about what might be the most productive areas for WebAL research going forward [80]. Areas highlighted included using the Web as a shared testing ground—a global laboratory—for ALife experiments, and using the close parallels between the Web and natural biological environments (e.g. both are large, dynamic, heterogeneous, noisy, and distributed) to inspire the design of complex WebAL worlds. As we will see later on (Section 7), both of these areas have become important components of current WebAL research.

5 The 2000s: WebAL Develops

We begin this section by describing a rather different kind of intersection between ALife and the Web that took place at the Artificial Life VII conference in 2000, held in Portland, Oregon. One aspect of the conference theme of “*Looking backwards, looking*

²⁸<http://awards.bafta.org/award/2001/interactive/interactive-arts>

²⁹See also <http://sodarace.net/>, <http://soda.co.uk/work/sodarace-online-olympics>

³⁰In 2013, Szerlip and Stanley developed an open-source browser-based version of Sodarace, called *IESoR* [123]. It features a developmental encoding of creatures suitable for evolutionary experiments, and is designed to be an accessible platform that other researchers can easily use.

³¹An archived copy of the project's website is available at <http://web.archive.org/web/20050126162950/http://norgev.alife.org/>.

³²They were absolutely right that this kind of architecture is now more feasible, although in modern systems this can be accomplished using native HTML5 technology rather than Java—see Section 7.2.

forwards” was a drive to better understand the interests and opinions of the community regarding the status and direction of ALife as a field of study. In addition, the organizers wished to canvas the community on the establishment of a professional society. In order to achieve these tasks, Steen Rasmussen and colleagues set up a web-based survey which allowed open-ended responses to a number of questions formulated by a group of community members [93]. The survey was implemented using Active Server Pages³³ technology, and stored responses in a database for further analysis. The survey authors produced histograms and mind-maps from the results in order to better understand the interests and concerns of the community, and also to investigate differences in opinions between respondents from computer science and biology. In addition, the positive responses to the questions regarding the establishment of a professional society led to the formation of the International Society for Artificial Life. In reporting the survey design and analysis, Rasmussen *et al.* explained that they wished to use the Web to harness the collective intelligence of the ALife community. In addition to reporting the specific results of their survey, they also described a more general methodology for performing this kind of web-based collective intelligence gathering and self-organization of knowledge from a community [93]. This is an early example of WebAL *crowd creativity*—we will come across other examples later in the paper, and briefly summarize these projects in Section 7.1.

Returning to more familiar flavors of WebAL, the general increase in the computing power of desktop PCs in the 2000s over the 1990s allowed the introduction of more sophisticated web browsers, and meant that it was feasible to run more powerful programs on the client side (within the browser). Hence, the 2000s witnessed an increase in complexity of WebAL projects, made possible by the parallel improvements in hardware and web technology in addition to the scientific progress of the field itself. In the remainder of this section we look at how the changing landscape of web technology affected the development of WebAL during the 2000s.

5.1 The changing technology landscape of WebAL

5.1.1 Java applications and applets of growing sophistication

Countless Java applets appeared in the late 1990s and early 2000s that simulated ALife related concepts, some of which were already described in Section 4. In addition, various cellular automata (CA) were popular.³⁴ While many of the earlier examples were fairly simple, some serious academic projects made their code available as applets as a means of distribution to encourage experimentation by other researchers (e.g. [68]).³⁵

³³<https://msdn.microsoft.com/en-us/library/aa286483.aspx>

³⁴For lists of many of these, see <http://cell-auto.com/links/> and <http://www.aridolan.com/ad/CA.html>.

³⁵As the decade developed, the growing popularity of agent-based modeling toolkits such as NetLogo [111] made it easier than ever to package simulation models as Java applets.

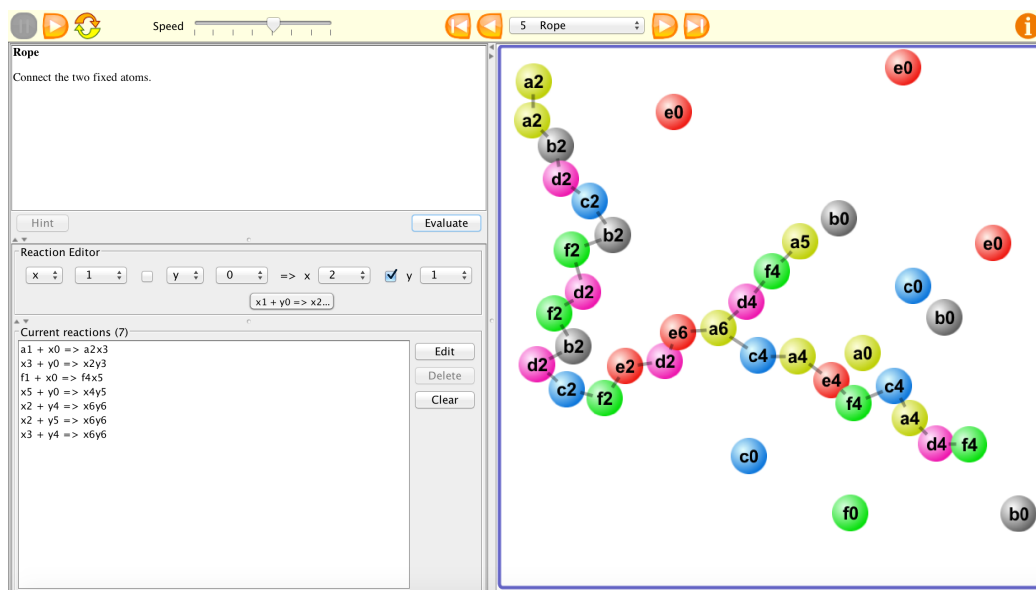


Figure 4: The *Organic Builder* user interface.

In addition, some more elaborate projects were developed as interactive educational tools, including *Soda Constructor* (described in Section 4 but further developed in the 2000s). Another example is *Organic Builder*,³⁶ an accessible tool for experimenting with artificial chemistries, developed by Tim Hutton and first launched in 2005. The system allowed users to edit the reaction rules of an artificial chemistry, and immediately observe the results in a browser-based graphical simulation of interacting particles (see Figure 4). A series of progressively harder challenges was presented, to test a user's skill in devising reaction rules to achieve particular behaviors. All 19 challenges set by Hutton were solved by users, sometimes in very unexpected and ingenious ways. As well as solving the challenges individually, some users also discussed and shared their results on an online discussion forum.³⁷ Reflecting on the experience of running the system for a number of years, Hutton commented “[t]hough it was initially intended as a set of challenges to be tackled as a game, the users experimented with the system far beyond this and discovered several novel forms of self-replicators. When searching for a system with certain properties such as self-replication, making the system accessible to the public through a Web site is an unusual but effective way of making scientific discoveries, credit for which must go to the users themselves for their tireless experimentation and innovation” [39]. Thus, in addition to the education and outreach goals, *Organic Builder* also provides an early example of web-mediated crowdsourced human computation for solving complex tasks.

Another important WebAL project in the late 2000s was the evolutionary art system

³⁶<http://organicbuilder.sourceforge.net/>

³⁷<http://groups.google.com/group/OrganicBuilder/>

Picbreeder [109, 108]. The system, launched in 2007 and still running today, allows users to evolve two-dimensional images using a Java applet which sends results back to a central server. The evolved images and their lineages can then be viewed via the project's website.³⁸ The images are encoded by a neural network variant called *compositional pattern-producing networks* (CPPNs) [116] and evolved through an implementation of the *neuroevolution of augmenting topologies* (NEAT) algorithm [118, 119].³⁹ A major innovation of *Picbreeder* was its support for a process called *branching*, which enables genuine collaborative interactive evolution; users are allowed not only to evolve their own images but also to select and continue evolving images produced by other users. As branch accumulates upon branch, the system facilitates the interactive evolution of deep lineages of evolved pictures encompassing the contributions of many users, and accordingly also enables the collective exploration of a vast search space of images. Results from *Picbreeder* demonstrate that it is possible to effectively harness the input of many unrelated users through branching. Some examples of images evolved with *Picbreeder* are shown in Figure 5.

Java-based projects such as *Organic Builder* and *Picbreeder* capitalized on the opportunities for distributed interaction and crowdsourced computation afforded by the Web. Elsewhere, other projects continued to utilize Java and applets purely as a means for easy distribution, dissemination and engagement. A prominent example is Hiroki Sayama's *Swarm Chemistry* project⁴⁰ [104], introduced in 2006 and still being actively developed.⁴¹

5.1.2 Adobe Flash

In addition to Java, Adobe Flash⁴² was another widely used technology in the 2000s for delivering animated multimedia web content (see Figure 1). There are various examples of Flash-based WebAL projects; one of the most notable being a web-based genetic algorithm for evolving cars to run over uneven terrain in a 2D simulated physics environment, which first appeared in 2008.⁴³ This simulation served as the inspiration for the more

³⁸<http://picbreeder.org>

³⁹A earlier, but much more limited project in web-based evolutionary art based upon the NEAT algorithm was the *Living Image Project* (<http://w-shadow.com/li/>), written by Jānis Elsts and operational over 2006–2007. The server-side application was partially based upon Mattias Fagerlund's Delphi NEAT package (<http://nn.cs.utexas.edu/?neatdelphi>), and generated PNG image files of evolved images to be displayed on the client-side web page.

⁴⁰<http://bingweb.binghamton.edu/~sayama/SwarmChemistry/>

⁴¹Swarm Chemistry has inspired various more recent independent projects, including an iOS app (<https://itunes.apple.com/us/app/emergent/id965513030>) and an Adobe Flash-based web version (<http://flexmonkey.blogspot.co.uk/2013/08/advection-swarm-chemistry-with.html>).

⁴²https://en.wikipedia.org/wiki/Adobe_Flash

⁴³The original site is now defunct, but it can still be seen on the Internet Archive at <https://web.archive.org/web/20100729074214/http://www.wreck.devisland.net/ga/>. We have been unable to trace the author of this work, beyond his anonymous announcement on Reddit (https://www.reddit.com/r/programming/comments/7i22c/genetic_programming_evolution_of_mona_lisa/c06pt65).




Figure 5: Examples of evolved images from *Picbreeder*.

recent and very popular reimplementations of the idea called *BoxCar2D*.⁴⁴ Developed in 2011 by UC Santa Cruz graduate Ryan Weber, *BoxCar2D* was also written in Flash, and used a Flash port⁴⁵ of the popular Box2D physics engine.⁴⁶ Also in 2011, shortly after the arrival of *BoxCar2D*, Rafael Matsunaga released a reimplementations of the system using native HTML5 technologies rather than Flash.⁴⁷

5.1.3 WebAL in online virtual worlds

Beyond the realms of academic research, Linden Lab's online virtual world *Second Life* was launched in 2003 and gained massive worldwide popularity over the following years.⁴⁸ A number of ALife related projects were developed within this platform, two of the most notable being *Svarga* and *Terminus*, which both came to prominence in 2006. *Svarga*, created by *Second Life* user Laukosargas Svarog, was an island with a

⁴⁴<http://boxcar2d.com/>

⁴⁵<http://sourceforge.net/projects/box2dflash/>

⁴⁶<http://box2d.org/>

⁴⁷See <http://rednuht.org/test/simulator/>. Matsunaga has more recently also written other web-based evolutionary systems using 2D physics, including one for biped walkers (http://rednuht.org/genetic_walkers/) and another for a quadruped (<http://rednuht.org/geneticat/>).

⁴⁸<http://secondlife.com/>

fully functioning ecosystem comprising a weather system and various types of plants and animals.⁴⁹ The island can still be visited in *Second Life* today,⁵⁰ although it was purchased from Svarog by Linden Lab in 2010⁵¹ and some of its original ecosystem features may no longer be present. Shortly after the release of *Svarga*, a separate effort was launched by the *Ecosystem Working Group* and associated with the in-game location *Terminus*.⁵² The group's aim was to develop an open-source programming language that would not only allow developers to freely create their own creatures, but would also allow the creatures in *Terminus* to interact and evolve using a shared language. However, the project apparently ran into funding and resource problems, and is no longer available.⁵³

5.1.4 Native HTML and HTTP technologies

The 2000s also saw the beginnings of a drive for native technologies and APIs in preference to Java applets and other techniques that required browser plugins. This drive accelerated towards the end of the decade, boosted by developments such as the beginnings of work on the HTML5 specification in the mid-2000s,⁵⁴ and the emergence of fast JavaScript engines offering substantial speed-up for client-side code, led by the introduction in 2008 of the V8 JavaScript engine by Google⁵⁵ (see Figure 1).

Even before these developments, there are examples of WebAL work that chose to focus on native technologies. An interesting early WebAL project that explored the potential of distributed computation and native client-side storage was William Langdon's *Pfeiffer* website, released in late 2001 and still running today⁵⁶ [49]. This browser-based system allowed users to evolve 2D patterns described by L-systems (see Figure 6). A user was presented with a variety of patterns on screen, and could select those they thought were good and bad, which directly influenced their evolutionary fitness. The patterns were presented as GIF files that had been generated by the server based upon the genetic description of an L-system and then sent to the user's browser. The user could also select patterns to be parents for a new offspring. Surviving patterns were made persistent on the client-side using HTTP cookies.⁵⁷ Users could name their favorite patterns, and save them, in which case they were not only stored locally but also uploaded to the system's global server where they would become available to be sent to other users. *Pfeiffer* therefore implemented distributed web-based evolution with aesthetic selection.

⁴⁹http://nwn.blogs.com/nwn/2006/05/god_game.html, <http://nwn.blogs.com/nwn/2010/03/svarga-returns.html>

⁵⁰<http://secondlife.com/destination/svarga>

⁵¹<http://nwn.blogs.com/nwn/2010/03/svarga-returns.html>

⁵²<http://news.nationalgeographic.com/news/2007/03/070308-second-life.html>

⁵³<http://forums-archive.secondlife.com/191/83/133314/1.html>

⁵⁴<http://www.w3.org/TR/html5/introduction.html>

⁵⁵https://en.wikipedia.org/wiki/JavaScript_engine

⁵⁶<http://www0.cs.ucl.ac.uk/staff/W.Langdon/pfeiffer.html>

⁵⁷https://en.wikipedia.org/wiki/HTTP_cookie

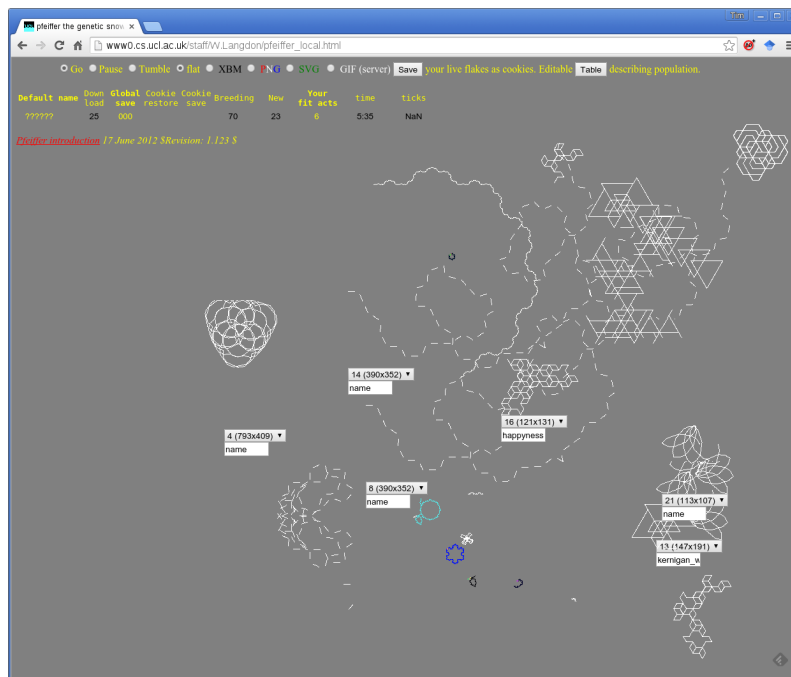


Figure 6: The *Pfeiffer* user interface.

Also noteworthy is Alexander Wait’s project *Quantum Coreworld*, first reported in 2004 [132]. The basic system was inspired by Rasmussen *et al.*’s early ALife system *Coreworld* [92], with the addition of “physics” inspired by quantum mechanics. The program was written in the C language and ran continually on a web server. The novel WebAL aspect of *Quantum Coreworld* was that it allowed users to inject “biotic” or “abiotic” changes into the system by uploading an instruction file via the system’s web interface.⁵⁸

An early example of a WebAL project using native HTML5 technologies is AlteredQualia’s *Image Evolution* website,⁵⁹ developed in 2008. This was inspired by Roger Johansson’s earlier (non web-based) work on evolving a polygon-based representation of the Mona Lisa.⁶⁰ At the time, the required HTML5 canvas element for drawing graphics was not widely supported by browsers, although the situation has now greatly improved, as we’ll see in the next section.

⁵⁸It appears that this system was not developed much further after its initial publication.

⁵⁹<http://alteredqualia.com/visualization/evolve/>

⁶⁰<http://rogeralsing.com/2008/12/07/genetic-programming-evolution-of-mona-lisa/>

Box 2: Highlights of other NetAL work from the 2000s

The 2000s also saw the continued development of various ALife related desktop (non-browser-based) systems that made use of Internet connections for collaborative functionality.

A prominent example of this type of project was *Golem@Home*, developed by Hod Lipson and Jordan B. Pollack at Brandeis University and launched in the summer of 2000.^a *Golem@Home* was an extension to the *Golem* project that sought to combine the evolution of simulated robotic systems with 3D printing technology to create a highly automated system for the design and manufacture of real-world robots [57]. *Golem@Home* allowed the general public to download a screensaver that harnessed idle CPU cycles to evolve simulated robots and animate some of them on the user's screen. If a network connection was available, evolved robots would occasionally migrate between different user's machines. Each screensaver kept a list of the IP addresses of other users' machines, allowing a highly distributed system without the need of a central server to manage the process. Over the course of approximately 12 months, *Golem@Home* attracted over 30,000 participants who jointly contributed several million CPU hours. From a technical perspective, the project was highly successful, although from a scientific perspective the authors did not observe the increase in evolved robot complexity they had hoped for.^b

Other examples of desktop-based systems with collaborative features implemented with Internet connections also continued to come from the area of ALife related computer games research. In 2003, Stanley and colleagues at the University of Texas at Austin started developing the computer game *NERO*, which allowed users to train a team of in-game agents using a real-time version of the NEAT architecture [117]. Once trained, the team could compete against an opposing team designed by another (often remote) user. Battle mode ran on a server such that both users could watch the battle while running the program on separate Internet-connected machines.^c

A few years later, Stanley and colleagues at the University of Central Florida produced *Galactic Arms Race (GAR)*, another desktop-based video game using Internet connectivity for collaborative game-play [34, 35].^d *GAR* includes a genetic algorithm to evolve new weapons (based upon particle systems) according to the weapons preferred by current users. In single-player mode, the weapons evolve according to the single user, but in full multiplayer Internet mode the weapons evolve in separate lineages based upon the aggregate usage of all players. The end result is the continual introduction of new in-game content based upon the players' tastes.

^a<http://www.demo.cs.brandeis.edu/golem/download.html>

^bResults as reported on the cited web page.

^c*NERO* was originally distributed as a binary file running on Mac or Windows. In 2009 work commenced an open-source version called *OpenNERO* (<http://nn.cs.utexas.edu/?opennero>).

^d<http://gar.eecs.ucf.edu>

6 The 2010s: Current WebAL

The movement away from Java applets and proprietary plugins towards native code and standardized, open APIs has continued to gain momentum during the 2010s. Along with the development of HTML5 (see Figure 1), the possibilities for creating native browser-based application have been greatly expanded by the appearance of many open JavaScript APIs. These include, among others, *Web Storage*⁶¹ for persistent data storage on client machines, *WebSocket*⁶² for two-way communication between server and client machines, *Web Workers*⁶³ for running background scripts on the client machine, *Web Audio*⁶⁴ for processing and synthesizing audio, and *WebGL*⁶⁵ for hardware accelerated graphics. This movement to open standards and technologies can be expected to continue now that the World Wide Web Consortium (W3C) has published the final, complete specification for HTML5.⁶⁶

In 2010, Swedish creative studio *B-Reel Creative*, in collaboration with Google Creative Labs and the American music video director Chris Milk, designed a ground-breaking web-based video called *The Wilderness Downtown*.⁶⁷ The project, a Grand Prix Winner at the 2011 Cannes Advertising Awards and recipient of a host of other awards,⁶⁸ was a trailblazer for many of the possibilities opened up by open web technologies including HTML5 video, audio and canvas. It included several ALife related techniques, such as an interactive bird flocking simulation that reacted both to the audio and to mouse interaction, as well as procedural drawing and generative typefaces.⁶⁹

The Wilderness Downtown graphically illustrated the potential of the modern Web as a platform for ALife applications. In the years since it was released, the reach, ambition and volume of WebAL research has greatly expanded, as other projects begin to realize the potential offered by modern, native web technology. In this section, we discuss a number of recent and current projects in a little more detail than in previous sections, in order to reflect the current state of the art. We will look at developments in the fields of Evolutionary Art and Design (Section 6.1), Games (Section 6.2), Science and Education (Section 6.3), Frameworks for WebAL (Section 6.4) and, finally, some new directions for WebAL (Section 6.5).

⁶¹<http://www.w3.org/TR/webstorage/>

⁶²<http://www.w3.org/TR/websockets/>

⁶³<http://www.w3.org/TR/workers/>

⁶⁴<http://www.w3.org/TR/webaudio/>

⁶⁵<https://www.khronos.org/webgl/>

⁶⁶The final specification was published in October 2014. See <http://www.w3.org/blog/news/archives/4167>

⁶⁷<http://thewildernessdowntown.com/>. The project was an interactive interpretation of the band *Arcade Fire*'s song *We Used to Wait*. For further information, see <http://b-reel.com/projects/digital/case/57/the-wilderness-downtown/>.

⁶⁸<http://docubase.mit.edu/project/the-wilderness-downtown/>

⁶⁹<https://www.chromeexperiments.com/arcadefire/>

6.1 Evolutionary Art and Design

6.1.1 EndlessForms

The *Picbreeder* system’s approach to collaborative interactive evolution through branching, described in Section 5.1.1, has inspired various subsequent projects. The foremost example is *EndlessForms*,⁷⁰ created in 2011 by Jeff Clune, Jason Yosinski, Eugene Doan, Hod Lipson and colleagues at Cornell University [17]. The design and goals of the project were influenced by *Picbreeder*, but *EndlessForms* focuses on the design of 3D shapes rather than 2D images. The practical purpose of the project is to allow people to create unique physical objects and see the power of evolution in action, and the scientific purpose is to explore what complex morphologies can be created with a computational implementation of developmental biology. A screenshot of the *EndlessForms* homepage is shown in Figure 7.

Besides demonstrating the scientifically interesting coupled power of evolution and human interaction, *EndlessForms* also addresses an important real world problem. While 3D printing technology is rapidly advancing, most people do not know how to design their own 3D objects even though they may have strong opinions and preferences. That is, many people are naturally skilled critics who will know what they like when they see it, but they are not skilled designers who can create what they desire from scratch. *EndlessForms* enables people to create and design interesting physical objects—such as jewelry, doorknobs, candlesticks, and sculptures—without any technical knowledge. This ability will be a key component to the shift from mass manufacturing to small-scale custom manufacturing fueled by the 3D printer revolution, as well as the sharing and easy modification of such information in the rapidly forming “internet of things”.

6.1.1.1 Mechanism and user interface

In a similar vein to *Picbreeder*, users are presented with a population of 15 shapes on a screen and are asked to select the one or several that they prefer (see Figure 8a). The shapes rotate slowly in the user’s browser window, allowing them to see all sides of the 3D object.

Once users have created a shape to their liking, they are able to save it to their personal account, tag it with a string for later reference, and download the shape in the STL (STereoLithography) file format to enable printing on 3D printers. For users who wish to have a printed copy of their shape but do not possess a 3D printer, they click through to have the shape printed and mailed to them by a 3D printing company, *Shapeways*.⁷¹

Critically, users are also afforded the option to publish the shapes they evolve. Upon

⁷⁰<http://endlessforms.com>

⁷¹<http://www.shapeways.com>

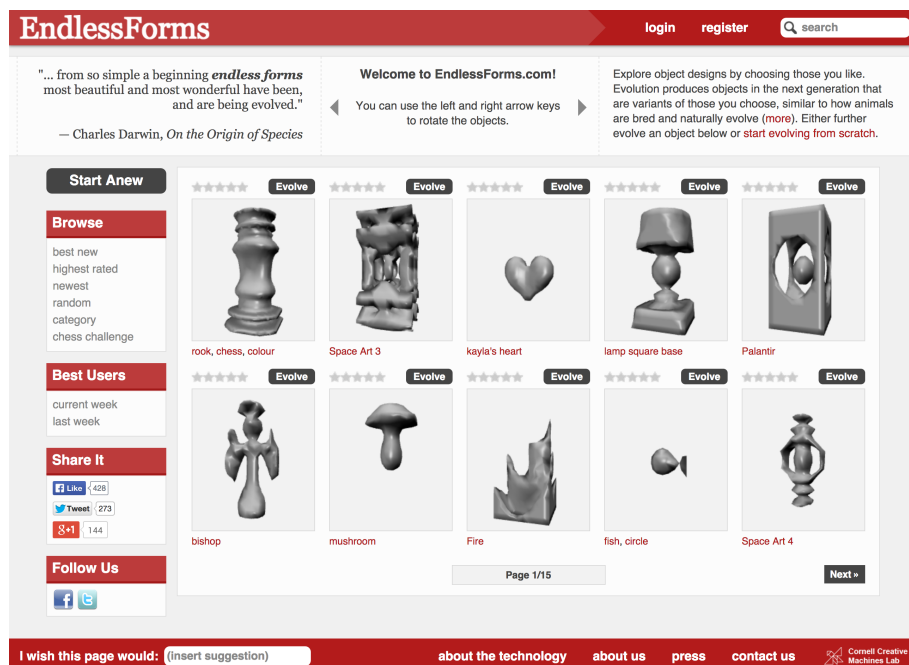


Figure 7: The *EndlessForms* homepage.

publication, the object appears on the homepage and becomes available for other users to download or further evolve. The authors have found that many of the most interesting shapes have been collaboratively evolved starting from previously published objects.

6.1.1.2 Technology

EndlessForms was created by combining several different web technologies (Figure 8b). On the server side, the main *EndlessForms* application is written in Python and uses the Django web framework.⁷² It handles all user interaction, stores important state—such as login information, user emails, publication records, organism fitness, etc.—in a database, and spawns and manages NEAT processes (as used in *Picbreeder*) to generate shapes. The client side code, written in JavaScript using the JQuery framework,⁷³ handles fetching each shape and drawing it using WebGL.⁷⁴ The site was initially constructed using shapes drawn in pure JavaScript, but this implementation was found to be slow and could not handle high resolution shapes. Thus, the site has become technologically possible only with the recent widespread adoption of WebGL in browsers to take advantage of the client machine’s graphics card for rendering.

⁷²<https://www.djangoproject.com/>

⁷³<https://jquery.com/>

⁷⁴<https://www.khronos.org/webgl/>

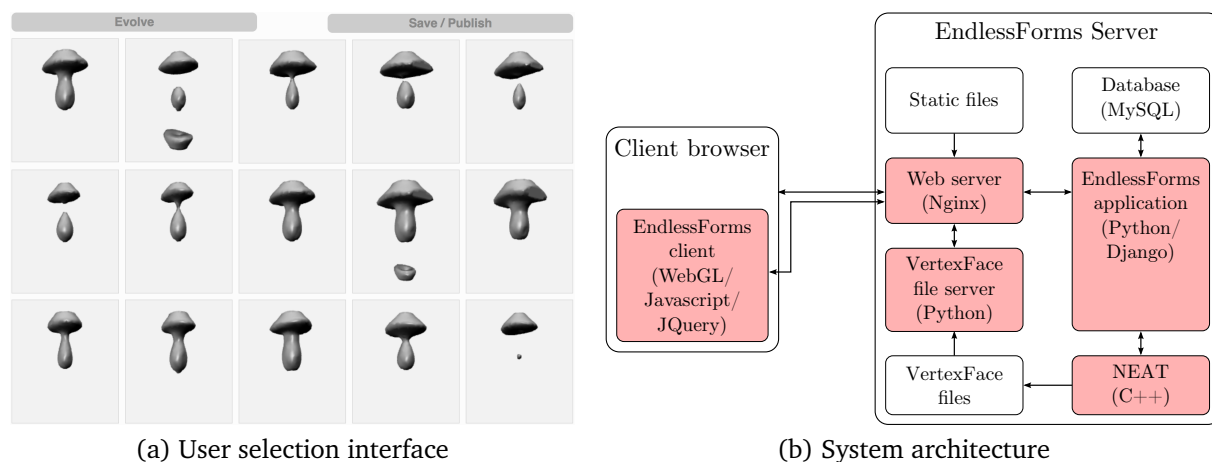


Figure 8: Example screenshot of the *EndlessForms* user selection interface (a), and architecture diagram showing the system’s main components and their interaction.

6.1.1.3 Results

As of 2015, users on *EndlessForms* have clicked through over 350,000 generations (screens of 15 individuals) to evolve over 5.3 million organisms. Many users have been driven to the site by coverage in the popular press, including *New Scientist*, *MSNBC.com*, *Slashdot*, *MIT Technology Review*, *KurzweilAI.net*, *Y-combinator Hacker News*, and the *Communications of the ACM*. It is only due to the large volume of users, eyes, and clicks that many interesting shapes have been able to be created. Some of these objects are shown in Figure 9, together with examples of physical objects created by 3D printing of some of the evolved forms.

6.1.2 Other recent work

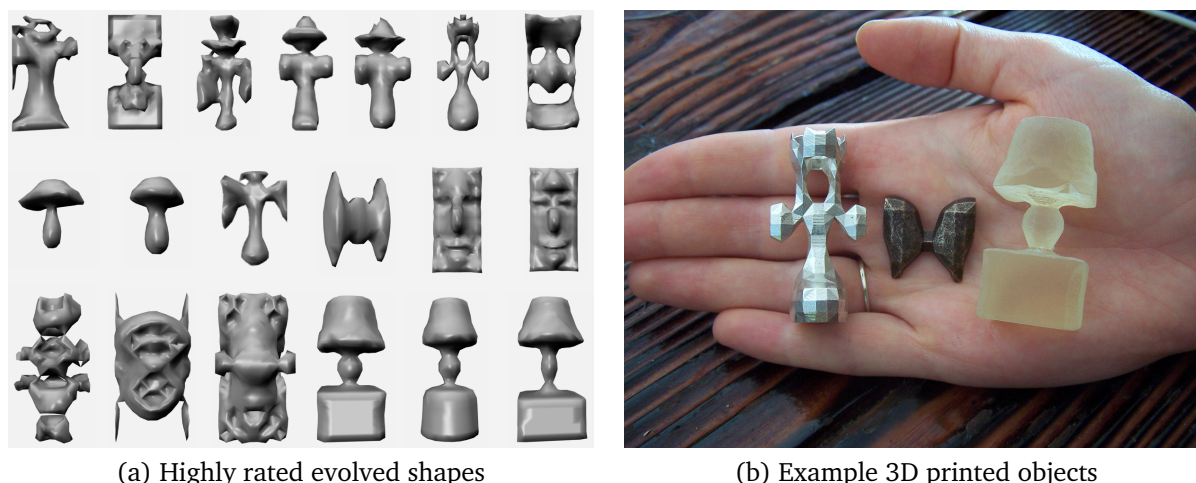
A recent project called *DrawCompileEvolve*,⁷⁵ created in 2013 by Rasmus Taarnby and Jinhong Zhang in Sebastian Risi’s lab at the IT University of Copenhagen, builds on some of *Picbreeder*’s web technology while adding the novel idea of a genotype-to-phenotype compiler [98]. This system allows users to draw a sketch of a particular image annotated with regularities (e.g. a butterfly with two symmetric wings) and then further evolve a CPPN representation of the image interactively [138].

Another new project with a strong *Picbreeder* flavour, but implemented using HTML5 SVG (Scalable Vector Graphics) support, is Craig Mandsager’s *Genolve* system (2014).⁷⁶

Elsewhere, a novel variety of WebAL was reported by Joshua Auerbach in 2012 while

⁷⁵<http://rasmustaarnby.dk/drawcompileevolve/>

⁷⁶<http://www.genolve.com/>. For further details, see <http://wildwebwidget.com/>.



(a) Highly rated evolved shapes

(b) Example 3D printed objects

Figure 9: Highly rated 3D shapes evolved in *EndlessForms* (a), and 3D-printed physical objects generated from evolved forms, in silver, bronze and plastic (b).

working at the University of Vermont [9]. This work evolved 2D images with a similar representation to that used in *Picbreeder*. However, the key difference was that the fitness of each image avoided the time consuming process of user selection. To accomplish this automation, the fitness function included a call to *Google Search by Image*⁷⁷—a Google service that performs an image search based upon a query image instead of (or in addition to) query text. Auerbach reasoned that, since images on the Web should be of interest to humans (otherwise they would not have been uploaded), interesting images should return many hits. Hence, the number of returned hits of a query image was used as a component of its fitness. Some example results are shown in Figure 10.

This work was novel in that it leveraged a third-party web service for fitness function evaluations, and presented one possibility for leveraging existing “big data” stores concerning human preferences. Additionally, the work presented several lessons for future research in this vein. Primary among these lessons were the constraints imposed by utilizing a service not meant for automated interaction. Calls to the service needed to be rate-limited in order to avoid being blocked by Google, which turned out to be a major setback to performing larger scale experimentation. This obstacle highlights the need to involve the owners of web services in future research, so that such issues can be avoided.

An alternative approach to automating the fitness evaluation of evolutionary design systems was described in 2015 by Anh Nguyen, Jason Yosinski and Jeff Clune at the University of Wyoming, in their work on *Innovation Engines* [73]. Although they employed a deep neural network (DNN) rather than a web service to implement the evaluation function, their system relied upon web services in an indirect but interesting way. The

⁷⁷<http://www.google.com/imghp?sbi=1>




Figure 10: Images from Auerbach’s work on automated image evolution using a fitness measure related to the number of hits returned by *Google Search by Image* [9].

DNN was trained on labeled images from the ImageNet Dataset⁷⁸. The evolutionary component of the system (based upon the CPPN/NEAT approach) was then challenged to generate images that the DNN would classify as representing a previously-trained label with high confidence. Although the system itself was not web-based, the ImageNet Dataset upon which the DNN was trained, which contained 1.3 million labeled images, was created with the help of Amazon’s *Mechanical Turk* service⁷⁹ [22]. *Mechanical Turk* is a web-based crowdsourcing marketplace that allows “requesters” (individuals or businesses) to farm out tasks that require human intelligence (in this case, labeling images according to various categories) to a crowdsourced group of workers who accept the work and get paid for completing it. Hence, crowdsourcing has played an important part in the *Innovation Engines* project.

Beyond two and three-dimensional forms, WebAL art applications have also been extended to sound generation. One such application is the *Breedsizer* system,⁸⁰ developed in 2015 by Björn Þór Jónsson, Amy Hoover and Sebastian Risi, which allows users to interactively evolve novel timbres in a manner similar to *Picbreeder* [43]. *Breedsizer* builds on the recently introduced *Web Audio API*⁸¹, along with the Worldwide Infrastructure for Neuroevolution framework (WIN) (which we discuss in more detail in Section 6.4.3).

⁷⁸<http://www.image-net.org>

⁷⁹<https://www.mturk.com/mturk/welcome>

⁸⁰<http://bthj.is/breedsizer/>

⁸¹<http://webaudioapi.com/>, <http://www.w3.org/TR/webaudio/>

6.2 Games

A variety of WebAL-related games have appeared in recent years, making use of several different technologies and platforms.

6.2.1 Desktop and Multiplatform WebAL Games

An example of using Facebook as a platform for WebAL is provided by *Petalz*,⁸² developed by Risi and colleagues first at the University of Central Florida and more recently at the IT University of Copenhagen. *Petalz* is a social game that allows users to breed and share evolved virtual flowers [99]. Released in 2012, the game is written using the Adobe Flash platform and delivered as a Facebook app. Facebook's Graph API⁸³ allows players to share their flower creations on other people's walls or sell them through an in-game marketplace. *Petalz* enables players to breed new flowers from purchased market seeds thereby facilitating meaningful collaborations between users. Like *EndlessForms*, *Petalz* also allows players to transfer their evolved flowers to the real world via 3D printing through the *Shapeways* API [101, 100].

Another platform that has been used for developing WebAL games is the *Unity Game Engine*⁸⁴, which facilitates (among other things) the use of Internet connections for collaborative gaming.⁸⁵ An example of such a game is *EvoCommander*⁸⁶ [40], which was inspired by ideas introduced in *GAR* and *NERO* (discussed in Box 2), and released in 2014. *EvoCommander* allows players to incrementally evolve arsenals of ANN-controlled behaviors such as ranged attack, flee, etc. Players can then battle other players' robots online, and through the novel game mechanic of "brain switching", selecting which evolved neural network is active at any point during a battle.

Unity has proven a powerful tool for a broad range of web game development. Another recent game that builds on its network capabilities is *FPSEvolver* [79], in which a group of players iteratively generate, play and improve multiplayer FPS levels to fit their particular preferences by voting on a selection of evolving levels. Other new games are in active development, such as *Evolve and Conquer*,⁸⁷ a StarCraft-style game focused on teaching evolutionary principles.

An interesting prospect for WebAL is to allow collaborative problem solving through a crowdsourced web-based approach. Online video games such as *Foldit*,⁸⁸ in which users

⁸²<https://apps.facebook.com/petalzgame>

⁸³<https://developers.facebook.com/docs/graph-api>

⁸⁴<https://unity3d.com/>

⁸⁵*Unity Personal edition* is free but proprietary software. It runs on many different platforms with the notable exception of Linux, although an experimental Linux build is currently under development (see <http://blogs.unity3d.com/2015/08/26/unity-comes-to-linux-experimental-build-now-available/>).

⁸⁶<http://jallof.com/thesis/>

⁸⁷<http://adamilab.msu.edu/evolve-and-conquer/>

⁸⁸<http://fold.it/portal/>

have to discover protein structures, hint at the power of crowdsourcing the brain's natural ability for certain tasks that involve pattern matching or spatial reasoning. A recently introduced WebAL system that similarly tries to harness human intuition is *Brain-Crafter*,⁸⁹ which allows users to collaboratively build artificial neural networks to control a simulated robot in an ALife setting [86]. The aim of this project is to ultimately facilitate a mixed-initiative process, in which a human and a computational creator take turns and propose changes to an evolving neural network [102].

6.2.2 Mobile and augmented reality WebAL

At the time of writing there has been little in the way of WebAL projects designed specifically for mobile platforms. However, we are aware of two significant projects of this type that are currently under development. *Wiggle Planet*,⁹⁰ a company founded by ALife veteran Jeffrey Ventrella (whose earlier work on the *Absolut Kelly* website was described in Section 4), is currently developing an augmented reality, mixed media game called *Polly Peck's Journey*. This is an interactive puzzle game, comprising a physical book and tablet application, that aims to encourage children to interact, learn and play in their natural environments. Children will be able to adopt the animated augmented reality characters in the game, called *wiglets*, and there are plans for the wiglets to be stored and shared in the Cloud. Another web-oriented aspect of *Polly Peck's Journey* is that its development was partially funded through a *Kickstarter* crowdfunding campaign that raised over US\$15,000 of development funding.⁹¹ Elsewhere, Jane Prophet and Mark Hurry are working on a new version of *TechnoSphere* (the original version of which, discussed in Section 4, was one of the earliest WebAL systems). *TechnoSphere 2.0*⁹² will be an augmented reality mobile application [89]. An Android app version of the program is currently undergoing beta testing.

A somewhat different example of the combination of ALife research and mobile apps is provided by *AppEco*,⁹³ an agent-based simulation model of mobile app ecosystems developed by Soo Ling Lim and Peter Bentley. The authors calibrated the model with real data about numbers of developers, apps and users collected from Apple's iOS ecosystem over three years. In a series of studies, they used it to investigate the effectiveness of various design strategies for developers [54], publicity strategies for new apps [53], and app store organization strategies [55].

⁸⁹<http://braincrafter.dk/>

⁹⁰<http://www.wiggleplanet.com/>

⁹¹<https://www.kickstarter.com/projects/1582488758/peck-pecks-journey-a-picture-book-that-spawns-virt>

⁹²<http://technosphe.re/>

⁹³<http://www.appeco.co.uk/>

6.3 Science and Education

Many of the projects discussed in Sections 6.1 and 6.2 implicitly introduce users to ALife concepts such as evolutionary algorithms. In this section, we review a number of recent projects that have more explicit pedagogical aims, and others developed as platforms for open science.

6.3.1 Evolutionary Robotics

The first three projects discussed below all focus on using the Web for teaching and outreach in the field of evolutionary robotics.

6.3.1.1 Ludobots

*Ludobots*⁹⁴ is an educational WebAL system developed by Josh Bongard and colleagues at the University of Vermont, and launched in 2012. The system serves as an infrastructure for those who wish to explore the design and evolution of robots. It is designed for students with a wide range of experience, from no programming expertise up to students who create and program new projects for other students to learn from. Students are gradually guided through a series of increasingly challenging projects, during which they learn about various concepts such as evolutionary algorithms, artificial neural networks, robotics, and embodied cognition.

The first project requires no programming: the student uses a web interface to create a robot by simply “connecting the dots” and observing the resulting robot behave in a web-embedded physics engine⁹⁵ (see Figure 11b) [131]. Later projects switch from the web-based simulation to a more powerful C++-based simulation running on a student’s computer (see Figure 11c). Importantly, *Ludobots* does not restrict students to use pre-specified robots or environments: students are free to modify their growing code base as they progress, and even modify the instructions of the assignments themselves. Figure 11a outlines the system’s curriculum flow.

Ludobots was inspired by open-ended web projects such as Reddit and Wikipedia: as the site grows, the growing user base continuously expands and improves the site’s content. Students may improve existing assignments, annotate those assignments with educational material, and create new projects of their own. The creators of *Ludobots* intend for this positive feedback of material to help overcome one of the major limitations of ALife projects: once an investigator publishes a paper describing their project, the project falls into disuse and is rarely replicated or extended by others.

⁹⁴<http://www.reddit.com/r/ludobots/>

⁹⁵Implemented using ammo.js (<https://github.com/kripken/ammo.js/>), a JavaScript port of the popular Bullet Physics engine (<http://bulletphysics.org>), and the Three.js WebGL programming library (<http://threejs.org/>).

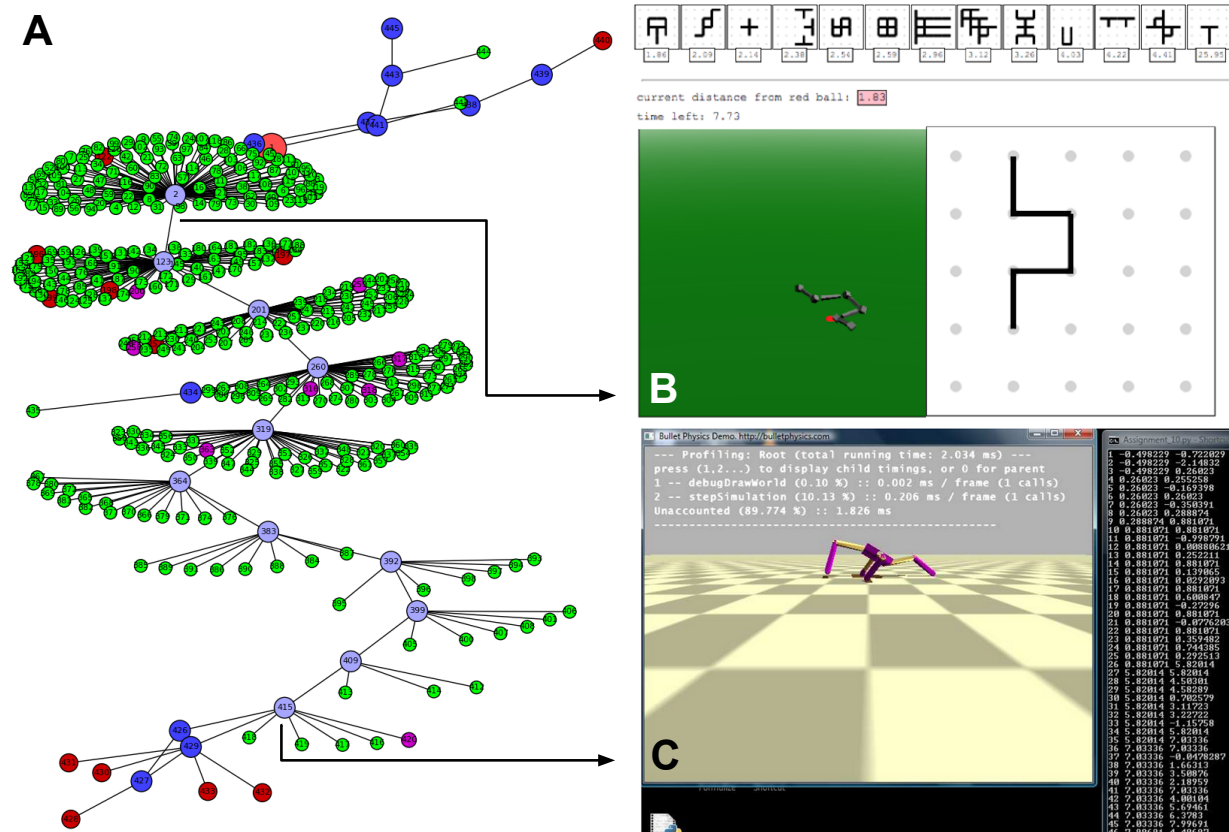


Figure 11: The *Ludobots* Educational Platform. The course tree (A) depicts all activity on the platform visually: light blue nodes represent programming assignments; green nodes represent student submissions; purple and red nodes represent resources and questions, respectively; dark blue nodes represent student-generated projects. The first project (B) requires no programming: the student simply creates a robot by “connecting the dots” and observing the resulting robot behave in a web-embedded physics engine. They can also see robots created by other students at the top of the page. Students progress through a series of programming assignments to the final assignment (C), which results in a basic evolutionary robotics test-bed. Students can then extend their system by following projects made by other students (the three dark blue dots at the bottom of (A)), or create new projects of their own and advertise them to their fellow students.

6.3.1.2 Evolve-A-Robot and WebGL Visualizer

In 2014, Jared Moore and colleagues at Michigan State University reported work on *Evolve-A-Robot*, a web-based evolutionary robotics simulation⁹⁶ [71]. Their work pushes the boundary of what can currently be achieved in web-based 3D rigid body simulation and evolutionary robotics: it embeds a whole evolutionary robotics system and 3D physics simulation in a JavaScript client-side web application.⁹⁷

In the same paper, the authors also report work on a WebGL-based visualizer⁹⁸ that provides a tool for the interactive 3D visualization of previously-recorded evolutionary robotics experiments. In contrast to pre-recorded results videos, the system allows the user to watch the system from different angles and focal points.

The authors state that their motivation for both of these projects is to “facilitate the exchange of ideas with other researchers as well as outreach to K-12 students and the general public.”

6.3.1.3 RoboGen

Another ongoing project that uses WebAL in an educational context is *RoboGen*TM, introduced in 2014 by Joshua Auerbach, Dario Floreano and colleagues at EPFL [8]. *RoboGen* is a platform for the co-evolution of robot morphologies and controllers, which focuses on the evolution of real (rather than virtual) robots. The goal of the project is to provide a simple and cost effective means of evolving robots in simulation and rapidly fabricating them in reality. This goal is accomplished through the use of inexpensive 3D-printers⁹⁹ and the use of simple, open-source, low-cost, off-the-shelf electronic components.

RoboGen has already been used by more than 100 students for course projects in EPFL’s Bio-Inspired Artificial Intelligence class, and there are plans for it to be part of a future Massive Open Online Course (MOOC) on this topic. Like *Ludobots*, *RoboGen* has been designed to be flexible in order to accommodate users with a diverse set of skill levels and backgrounds. Those with little familiarity with evolutionary computation or programming can familiarize themselves with new concepts in a controlled way without needing to write any code, whereas advanced users can dive more deeply by customizing the evolutionary algorithm or simulator or even by introducing new morphological building blocks. This open-endedness is a major advantage of the platform and addresses current

⁹⁶<http://evolve-a-robot.github.io/>

⁹⁷Like *Ludobots*, *Evolve-A-Robot* also uses a JavaScript port of the Bullet Physics engine: specifically, PhysiJS (<http://chandlerprall.github.io/Physijs/>), which is actually built on top of ammo.js but runs the physics simulation in a separate thread using Web Workers (<http://www.w3.org/TR/workers/>) to improve performance. The web-based simulation component of *Ludobots* does not (yet) incorporate an evolutionary algorithm, hence *Evolve-A-Robot* is novel in this respect.

⁹⁸<https://github.com/jaredmoore/WebGLVisualizer>

⁹⁹Such as the MakerBot Replicator 2x: <http://store.makerbot.com/replicator2x>




Figure 12: Screenshot of *RoboGen*'s WebGL-based visualization interface.

concerns regarding science laboratory education [61].

The project has developed rapidly since its introduction. Initially, its only web-oriented aspects were the *RoboGen* website and wiki¹⁰⁰ that provided access to the necessary utilities for evolving and manufacturing robots.¹⁰¹ However, the project's website now features a WebGL visualization engine—inspired by Moore's work described in Section 6.3.1.2—allowing for rich client-side 3D rendering of evolved robots within the browser (see Figure 12). Recently, a port has been released of the complete *RoboGen* software stack that runs natively in the browser. This was accomplished by compiling the original *RoboGen* C++ code¹⁰² with *Emscripten*,¹⁰³ a source-to-source compiler that generates highly efficient JavaScript from C/C++ source code. The ability to run the complete system directly in the browser without the need to download and install additional software drastically lowers the barrier to entry of using *RoboGen* for students, researchers and hobbyists. While similar to, and inspired by, the projects mentioned above, *RoboGen*-on-the-web is distinct in its concern for evolving and simulating robots that can be easily manufactured.

¹⁰⁰<http://www.robogen.org>

¹⁰¹Including an evolution engine complete with a physics simulator, as well as utilities both for generating design files of body components for 3D printing, and for compiling neural-network controllers to run on an Arduino microcontroller board (<http://www.arduino.cc>).

¹⁰²Including the *Open Dynamics Engine* physics simulator (<http://www.ode.org>).

¹⁰³<http://kripken.github.io/emscripten-site/>

There are plans to more fully capitalize on the potential offered by these web-based aspects of the system in future work. These include adding the facility for collaborative robot evolution, where users may upload new morphological building blocks and/or complete evolved robots for others to fabricate themselves, modify, or use as seeds for their own further evolutionary runs, and even to allow for distributing computational resources among users of the system. The creators of *RoboGen* believe that it will be through this collaborative exchange of ideas that the platform can truly blossom, a diversity of morphologies can be fabricated and tested, and users can benefit from each other's experience.

6.3.2 Other Science and Education Projects

The evolutionary robotics projects described in the previous section constitute a major strand of current WebAL research. But science and education-based WebAL is not limited to robotics; in this section, we look at projects in other subject areas.

6.3.2.1 Avida and Avida-ED

The *Avida* software¹⁰⁴ is a well-known research testbed that allows users to perform experiments on populations of digital organisms to investigate questions about evolutionary or ecological dynamics [74]. *Avida-ED*¹⁰⁵ is the educational version of this software, intended to allow students to get first-hand experience manipulating evolving populations, both to develop intuition about evolution and to generate a deeper understanding of the nature of science [85]. The *Avida* development team, based in the Digital Evolution Lab at Michigan State University, have recently created a fast JavaScript version of the code by compiling the original C++ version with *Emscripten*.¹⁰⁶ A new *Avida-ED* web interface¹⁰⁷ has been built for this web version of *Avida* and is in final testing before initial release. A full web reimplementations of the *Avida* research platform is also under development, and progressing rapidly.

6.3.2.2 The Ladybug Game

In 2012, Terence Soule and colleagues at the University of Idaho developed a web-based system called *The Ladybug Game*¹⁰⁸ as an interactive tool for teaching children about

¹⁰⁴<http://avida.devosoft.org/>

¹⁰⁵<http://avida-ed.msu.edu/>

¹⁰⁶<http://kripken.github.io/emscripten-site/>

¹⁰⁷<https://github.com/devosoft/avida/tree/AvidaEd-Web>

¹⁰⁸<http://www2.cs.uidaho.edu/~tsoule/ladybug-lesson/>

evolution.¹⁰⁹ The system was originally written in Processing¹¹⁰ (from which a Java applet was generated), and later rewritten in JavaScript. *The Ladybug Game* presents the student with a simulation of a ladybug and a population of aphids on a colored background, where each aphid's color is genetically determined and the closeness of match between an aphid's color and the background affects its chance of being eaten by the ladybug. Over time the aphids' colors evolve in response to predation by the ladybug. The simulation was used as a tool in a series of five interactive lessons to demonstrate the roles of *selection*, *inheritance*, and *variation* in evolution. The user could interactively change the color of the leaf as the simulation proceeded, as well as select which of the three evolutionary features were active. *The Ladybug Game* was exhibited at the USA Science and Engineering Festival in 2012.¹¹¹

6.3.2.3 Swarm Grammars GD

Another recent example of an educational WebAL system is *Swarm Grammars GD*,¹¹² a WebGL-driven website by Sebastian von Mammen and Sarah Edenhofer at the University of Augsburg, Germany. The authors developed the tool as a means of introducing STEM research, and specifically ALife ideas, to high-school students. The system allows students to interactively configure and manipulate a *swarm grammar* system, which combines concepts from L-systems and flocking algorithms (see Figure 13). Initial studies were conducted as part of a Girls-in-STEM program at the University of Augsburg, Germany. In their 2014 paper, the authors discuss their experience of introducing the system to the students (aged between 12 and 15) and allowing them to explore its generative capabilities [130]. This preliminary study generated mixed results, but the authors use the experience to identify several ways in which the system should be developed in future work.

6.3.2.4 StringmolWeb

Tim Hutton's work on *Organic Builder*, a web-based platform for allowing users to learn about and experiment with artificial chemistries, was discussed in Section 5.1.1. A more recent system with similar goals is *StringmolWeb*,¹¹³ developed in 2010 by Adam Nellis, Ed Clark and Simon Hickenbotham at the University of York, UK. The system provides a

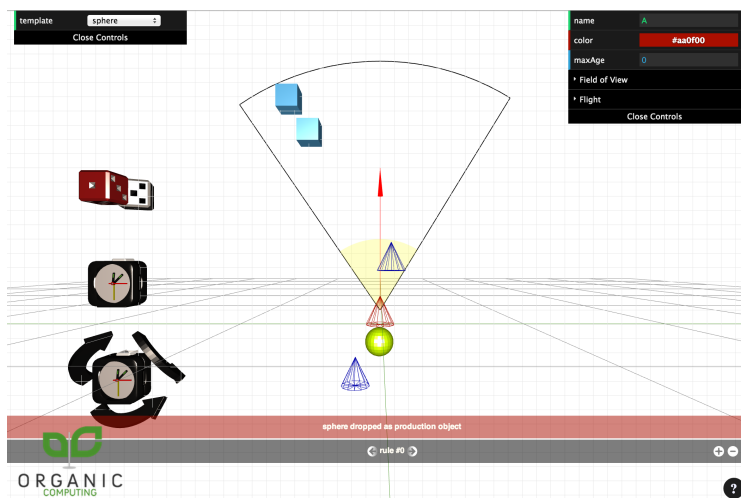
¹⁰⁹For further details, see also <http://beacon-center.org/blog/2012/12/03/beacon-researchers-at-work-teaching-evolution-the-ladybug-game/>

¹¹⁰<https://processing.org/>

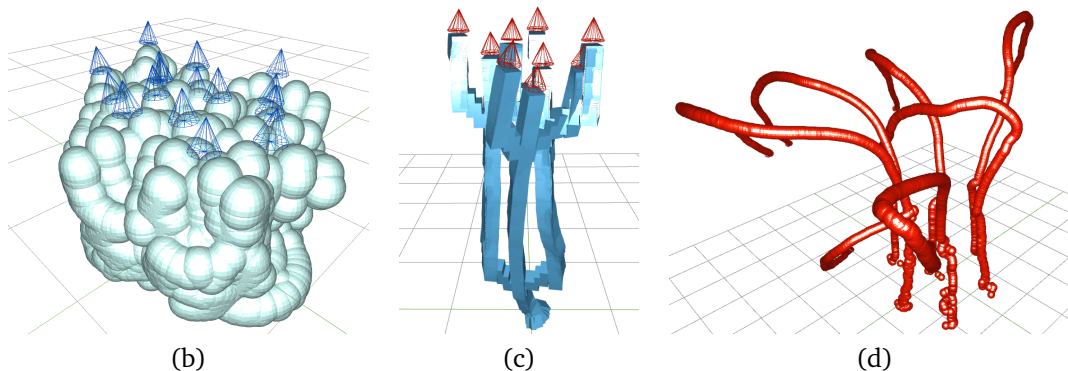
¹¹¹<http://www.usasciencefestival.org/component/content/article/66-2012-exhibits/213-2012-exhibits?&offset=12>

¹¹²<http://www.vonmammen.org/SG-GD/>. "GD" stands for generative design.

¹¹³<https://www.cs.york.ac.uk/nature/plazzmid/external/RUTSAC11/StringmolWeb/>. Source code available at <https://github.com/adamnellis/StringmolWeb>, and for the underlying *Stringmol* simulator at <https://github.com/franticspider/stringmol>.



(a) Rule editor interface



(b)

(c)

(d)

Figure 13: The *Swarm Grammars GD* user interface (a), and examples of agent behavior with traces enabled (b-d).

web-based interface to the *Stringmol* artificial chemistry system [36]. The client-side application is written in native HTML and JavaScript, and communicates via CGI¹¹⁴ scripts with a server running an instance of the *Stringmol* system (which is written in C++). It was used as an interactive educational tool at a tutorial workshop on artificial chemistries at the ECAL 2011 conference.¹¹⁵

6.3.2.5 OpenWorm

A prominent example of an ALife-related *open science* project is *OpenWorm*, which aims to create a whole system simulation of a living organism. Specifically, the goal is to develop a detailed 3D dynamic simulation of the nematode *C. elegans* [82]. Although the simu-

¹¹⁴https://en.wikipedia.org/wiki/Common_Gateway_Interface

¹¹⁵<https://www.cs.york.ac.uk/nature/plazzmid/external/RUTSAC11/>

lation itself is not web-based, the core team are distributed across the world and have regular team meetings using web-based collaboration tools. The project website actively seeks to recruit new members to the team, including scientists, programmers, artists and writers.¹¹⁶ All code, data and models produced by the project are open-source under the MIT license. The project also pursues a crowdfunding approach, seeking donations via the website and via a successful *Kickstarter* campaign that raised over US\$120,000 in 2014.¹¹⁷ The *OpenWorm* project therefore demonstrates multiple ways in which the Web can successfully facilitate large, open, collaborative projects of this kind.

6.3.2.6 The Broader Landscape

The WebAL-related science and education projects described above are part of the broader landscape of developments in web-based education. These projects were highlighted because of the connection of their subject matter or their investigators (or both) to the Artificial Life community. A review of the broader field is beyond the scope of the current paper, but we here provide some pointers to other work to give a flavor of more general developments.

In the area of web-based tools for primary and secondary education in biology, the *e-Bug* project¹¹⁸ is a more elaborate example than *The Ladybug Game* described in Section 6.3.2.2. *e-Bug* is a large European Union-funded project to develop educational material to teach microbiology. A number of different web-based games were developed in the project, aimed at different age groups, and the final versions were translated into 11 different European languages [27].

A somewhat different example of the intersection of biological science education and Web technology is provided by Kayhan Moharreri and colleagues' work on the *EvoGrader*¹¹⁹ system [70]. While the projects described above have focused on web-based delivery of learning and educational *content*, the *EvoGrader* system provides a free web-based tool for the automated *assessment* of undergraduate biology students' understanding of concepts relating to evolution and natural selection.

Of course, the Web also serves as the underlying delivery platform for many Massive Open Online Courses (MOOCs) beyond the more ALife-specific projects discussed in the preceding sections. The major MOOC providers¹²⁰, and other more focused platforms¹²¹ provide courses covering many topics in biology, evolution, complex systems and many other topics relevant to ALife.¹²²

¹¹⁶<http://www.openworm.org/>

¹¹⁷<https://www.kickstarter.com/projects/openworm/openworm-a-digital-organism-in-your-browser>

¹¹⁸<http://www.e-bug.eu/>

¹¹⁹<http://www.evograder.org>

¹²⁰E.g. <https://www.coursera.org/>, <https://www.edx.org> and <http://online.stanford.edu/>.

¹²¹E.g. <http://www.complexityexplorer.org/>.

¹²²The implementation of scalable assessment mechanisms is a challenge for MOOCs. Current approaches

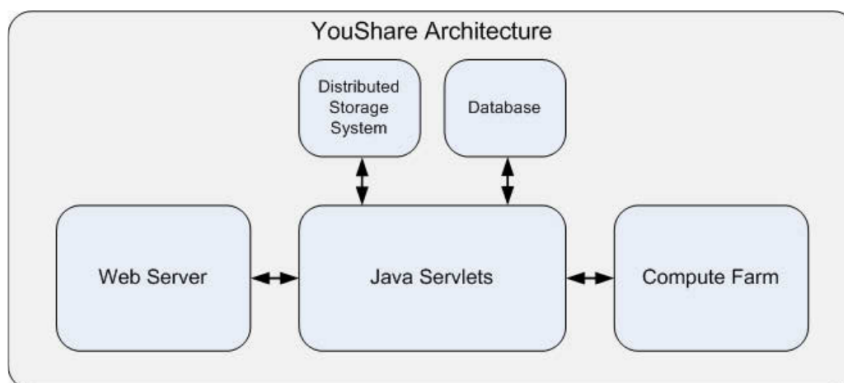


Figure 14: Overview of the *YouShare* architecture.

6.4 Frameworks for WebAL

Recent years have seen the emergence of a number of more general frameworks and platforms for WebAL research. These systems provide users with the tools and infrastructure to conduct particular kinds of WebAL experiments without having to write everything from scratch; within a given domain, they provide general-purpose facilities that can be reused by multiple research projects. Four such frameworks are highlighted in this section, followed by a brief discussion of other relevant work.

6.4.1 YouShare and the ALife Zoo

In 2013, Hickinbotham and colleagues reported WebAL-related work that made use of *YouShare*, a system that has the capability to run software designed for any platform in a web-based environment [37]. *YouShare* adopts a software-as-a-service model that exploits Virtual Machine (VM) technology and links it to a web browser. This design removes the need to port software to each target operating system (OS), since the software can be run inside a VM that runs the OS that the software was originally developed on. The only requirement is that the OS must also be able to run Java. The system architecture is illustrated in Figure 14. A web server, built using the Google Web Toolkit (GWT), interacts with a suite of Java Servlets. These servlets, in turn, process the request for a service: a database maintains the service, its VM requirements, and the data to be analyzed by the service. The requisite components of the service are then fetched from a (distributed) storage system to a compute farm, which deploys the VM and the service inside it. The service is run, and the user is notified when the job has finished, again via the web portal. By this process, it is possible to hook up to a browser any ALife

typically involve either the use of simple forms of assessment (e.g. multiple choice questions or formulaic questions with well defined answers), or peer grading [10]. Systems such as *EvoGrader*, which allow automated assessment of free-form written responses, could provide an important additional tool for realizing the full potential of MOOCs and other web-based educational tools.

system that can run in a Java-enabled VM. It is also possible to chain services together with analysis tools to build *workflows* of analysis.

To explore the way that ALife systems can exploit this technology, Hickinbotham et al. wrapped three well-known ALife software technologies for deployment on *YouShare*, to create the *ALife Zoo* [37]. The systems were *Stringmol* [38], *Avida* [74] and *Tierra* [94]. The authors wrapped each software package as a service, and the ‘test’ implementations that were provided with the source code were used as examples of how to use the system.

The *YouShare* framework has the potential to run or access a cloud- or grid-based HPC system, allowing large-scale ALife experiments to be run relatively easily. For example, it could form a suitable platform on which to implement a modern version of the *Network Tierra* project discussed in Box 1, with considerably less coding effort than was required for the original system.

Priorities identified by the authors for future development of *YouShare* and the *ALife Zoo* include facilities for visualization of a simulation’s progress, and the provision of software hooks to more easily allow new systems to utilize the services provided in a modular fashion.

6.4.2 COEL

Another recent project aimed at providing a framework for ALife related work is COEL, an open-source web-based chemistry simulation framework introduced in 2014 by Peter Banda and colleagues at Portland State University [11].¹²³ In contrast to the approach taken with *YouShare* of providing infrastructure and virtual machine technology to run existing heterogeneous simulation software, COEL provides a single, purpose built web-based simulation framework.

The COEL designers aimed not only to introduce a well-designed system that avoids researchers from having to reinvent the wheel by coding their own simulations, but also to create a simulation that: (1) runs on COEL’s own grid rather than on the client’s machine (as with *YouShare*); (2) allows geographically distributed teams to work together on a single platform; (3) provides remote database storage and backup facilities for experimental data; and (4) provides integrated visualization tools (a feature identified by the *YouShare* developers as an important next step).

COEL offers a unified, web-based environment for the definition, simulation and analysis of chemical reaction networks. It can be run from any browser without installation—embedded visualization is implemented using Google’s Chart API¹²⁴, and the server side is implemented using a variety of modern Java Virtual Machine (JVM)-based technologies. Without the need for installation, the authors argue that COEL has a larger potential audience than existing desktop-based systems. They also suggest that the cloud-based

¹²³<http://coel-sim.org>

¹²⁴<https://developers.google.com/chart>

data-storage facilities promote collaboration, sharing of results, and building upon past work by others.

The authors' vision for COEL extends beyond chemical reaction network simulation: their hope is for it to "become a common platform for diverse unconventional computing models" [11, p. 20].

6.4.3 WIN (Worldwide Infrastructure for Neuroevolution)

Systems such as *Picbreeder* (Section 5.1.1) and related projects use the Web as a platform for long-running, open, collaborative experiments in artificial evolution, but each required considerable effort to develop. Paul Szerlip and Ken Stanley are currently developing the *Worldwide Infrastructure for Neuroevolution (WIN)* project, which aims to provide a general-purpose solution that will greatly reduce the effort required to build such systems [124]. Also under development is an associated public web-based front-end for ongoing experiments built with the WIN platform, called WIN Online.¹²⁵ WIN has been designed as a lightweight and expandable collection of event-driven *Node.js*¹²⁶ modules, allowing developers to use just those packages that they require.

The designers of WIN aim "to make it trivial to connect any individual or lab platform to the world, providing both a stream of online users, and archives of data and discoveries for later continuation" [124, p. 901]. Although currently still at prototype stage, the authors have demonstrated its use in rapidly reimplementing both *Picbreeder* and *IESoR* (the recent browser-based version of *Sodarace*, described in Section 4). The WIN-based reimplementation of the latter, *win-IESoR*, additionally featured interactive evolution, which was not present in the original *IESoR* implementation. These reimplementations show that WIN can be successfully employed to dramatically reduce coding effort by reusing standard libraries. Furthermore, the data-archiving facilities of the system provide the possibility of reuse of *results* for further evolution in future experiments.

6.4.4 Empirical

Several of the projects mentioned in previous sections have used the *Emscripten*¹²⁷ source-to-source compiler that converts C or C++ code into high-performance JavaScript¹²⁸. The *Empirical* library¹²⁹, currently being developed by Charles Ofria and others at Michigan State University, works with *Emscripten* to simplify the construction of ALife experimental testbeds. *Empirical* provides tools to: (1) build data-driven web pages as part of

¹²⁵<http://winark.org>

¹²⁶<http://nodejs.org>

¹²⁷<http://kripken.github.io/emscripten-site/>

¹²⁸See <http://devosoft.org/an-introduction-to-web-development-with-emscripten/> for a quickstart guide to working with Emscripten.

¹²⁹<https://github.com/mercere99/Empirical>

Emscripten-compiled C/C++ software, (2) build dynamic configuration systems, (3) serialize results for easy analyses or continuation of runs, and (4) hook into a variety of JavaScript packages including *D3.js*¹³⁰ for visualization. Many of the tools are built to work appropriately for either web-distribution or native compilation using any standard C++ compiler, and various high-performance core C++ tools are provided that are generally useful in scientific software. *Empirical* is currently being used to develop the full reimplementation of the *Avida* platform, described in Section 6.3.2.1.

6.4.5 Other work

We conclude this section with a brief look at some other projects of relevance to WebAL frameworks.

The *SimWorld Agent-based Grid Experimentation System (SWAGES)*, initially conceived over 15 years ago and more recently developed as a component of the *Agent Development Environment (ADE)* project¹³¹, is an extendable distributed experimentation platform for large-scale agent-based simulations [107, 106]. At a high level, the architecture is somewhat similar to that of COEL. Implemented in Java, the system allows for the automatic parallelization and distribution of simulations, as well as providing analysis and visualization facilities, all controlled via a web-based interface. A full review of the field of distributed multi-agent systems is beyond the scope of this paper, but such a review can be found in [103].¹³²

Elsewhere, and also bearing some similarities to COEL (Section 6.4.2), Brucer Damer and colleagues' *EvoGrid* project sought to employ large-scale, distributed artificial chemistry simulations to study the origin of life.¹³³ The project was based upon a distributed architecture linking compute clusters via a web-based simulation management system. Some initial proof-of-concept results were reported in 2010–2012 [20, 21], although Damer and colleagues are currently concentrating on more traditional origins-of-life research (e.g. [18]).

In contrast to the kinds of projects mentioned above which involve the development of large, comprehensive frameworks for distributed WebAL, an alternative approach is to use more general-purpose existing tools—separately or in combination—to create cheap (in terms of development effort), lightweight WebAL frameworks.

One example of such an approach is provided by Atanas Radenski, who showed how the MapReduce model¹³⁴ could be used to distribute large-scale lattice-based ALife sim-

¹³⁰<http://d3js.org/>

¹³¹<http://ade.sourceforge.net/>

¹³²We also give an honorable mention to *breve.js* (<http://artificial.com/breve.js/>), a slimmed-down, browser-based version of Jon Klein's popular *breve* 3D simulation platform for multi-agent systems <http://www.spiderland.org/>. At the time of writing, *breve.js* is described as “work in progress”.

¹³³<http://www.evogrid.org/>

¹³⁴<https://en.wikipedia.org/wiki/MapReduce>

ulations in the cloud [91]. Specifically, he demonstrated how discrete and continuous versions of Conway’s Game of Life could be implemented on Amazon Elastic MapReduce.¹³⁵ Radenski investigated various optimization techniques for his approach, and discussed how his design could be used as a prototype for other such work.

Some interesting examples of lightweight approaches to web-based evolutionary algorithms have been published by Juan-Julián Merelo and colleagues in recent years. These include studies employing popular cloud-based storage services (specifically, Dropbox¹³⁶ and SugarSync¹³⁷) [67], and others implementing JavaScript-based evolutionary algorithms [66, 65]. Other work has studied how the performance of different architectures of web-based evolutionary algorithms is affected when client nodes join and leave the cluster during an evolutionary run [51]. These studies present some interesting ideas that could be applied to WebAL projects. A full review of the general field of web-based distributed evolutionary algorithms is beyond the scope of this paper, but recent reviews can be found in [30] and [31].

6.5 New Directions for WebAL

The work discussed up to this point in Section 6 has fairly naturally fallen into a handful of distinct categories. In this final subsection, we discuss preliminary work in a couple of projects that look at somewhat different aspects of WebAL.

6.5.1 EvoPopcorn: Distributed client-centric WebAL

Luis Zaman has recently described exploratory work with a distributed evolutionary system called *EvoPopcorn*, implemented with native web technology [137]. The main component of the system is a client-side application, written in JavaScript and HTML5, that presents the viewer with a simulation of a 2D arena in which simple organisms compete for food. Organisms die at random, and are replaced by mutated offspring of parents chosen according to their fitness (amount of food consumed). The resulting evolution of morphology and behavior of the organisms depends upon the distribution of food sources in the environment, which the user can manipulate.

Perhaps the most interesting aspect of this work in the current context is that Zaman extended the system so that evolved organisms could hop from one browser to another, using the *WebSocket* API.¹³⁸ Although the migration of organisms from one browser to another was physically routed through the system’s server, this work stands in contrast to most of projects discussed above in that computation is happening on the client-side (in the browser) rather than on a central server. The server is only acting as a relay

¹³⁵<https://aws.amazon.com/elasticmapreduce/>

¹³⁶<https://www.dropbox.com/>

¹³⁷<https://www.sugarsync.com/>

¹³⁸<http://www.w3.org/TR/websockets/>

station to route organisms from one browser to another. The conceptual architecture of *EvoPopcorn* is therefore similar to that of the *Golem@Home* project described in Box 2.¹³⁹ The main novelty is that *EvoPopcorn* is implemented using native web technology that works by a user simply visiting the website, with no plugins or software installation required. This work hints at the growing possibility of a much more distributed kind of WebAL, in which ALife organisms live “in the wild” (on browsers and local storage on client machines around the world) rather than being penned in on a central server and only released to clients on demand.

6.5.2 The Internet as a Living System

All of the work discussed in preceding sections has used the Web as a platform upon which to implement WebAL of one sort or another. Recent work by Mizuki Oka and colleagues takes a different perspective by asking the question: by measuring the characteristic dynamics and activity of the Web, might we consider the Web itself to be a living system?¹⁴⁰ In order to address this question, the dynamics of the Web are examined in terms of four characteristics that are essentially associated with living systems: excitability, autonomy, homeostasis and capacity to evolve.

A series of studies argue that aspects of the Internet can be regarded as: *an excitable medium*—specifically with reference to the dynamics of activity on social media sites [76]; *an autonomous system*—by applying the concepts of reactive and default modes of activity from brain sciences to social media and web search behavior [78]; *a homeostatic system*—by studying the adaptation and robustness of packet switching networks under varying data input loads [75]; and, *an evolving system*—in relation to vocabularies used in tags on a social media system [77].

These studies regard the Web as a complex chemical soup analogous to the primitive state of the earth. The work is guided by the possibility that the Web’s massive scale, complex dynamism, open richness, and social character could potentially be developing living systems spontaneously. Thus, the authors argue that it may be more profitable to study it using tools and concepts appropriate for understanding nervous systems, organisms, ecosystems and society, rather than using more traditional analytical tools employed in engineering and technological systems.

Elsewhere, Yuki Takeichi and colleagues have recently reported a study of activity on

¹³⁹In a similar vein to *Golem@Home*, the long running *DarwinBots* desktop ALife simulation (<http://www.darwinbots.com>) added an Internet Mode in 2011 (http://wiki.darwinbots.com/w/Internet_Mode). If enabled, this mode allowed bots evolved on one user’s machine to be teleported to other users’ machines via an FTP server.

¹⁴⁰While the analogies between networked computer systems and living systems have been highlighted before (e.g. [28]), Oka *et al.*’s work takes the idea of the Web as a living system more literally. Bedau *et al.* have previously suggested that social communities on the Web might display some aspects of genuinely living systems (in their terminology, they are borderline between “primary” and “secondary” living technology) [12], but Oka *et al.* develop this idea much further.

Twitter during major sporting events [126]. On the basis of their analysis, the authors are led to regard Twitter as an emergent “social sensor”—a distributed bionic extension to human sensory capability that exists and evolves in cyberspace but senses events in the physical world.

7 Emerging Themes and Future Directions

In the previous sections we surveyed the current landscape of WebAL and have seen how it has developed and expanded over its first two decades. In this final section we highlight some emerging themes that are apparent in current WebAL, and consider how these topics might influence future work.

7.1 Human/hybrid computation and crowd creativity

A great deal of the work surveyed here, especially the evolutionary art and design systems covered in Section 6.1 and some of the earlier systems discussed in Sections 4 and 5, involves the idea of *human or hybrid computation*, where some part of the computation is performed by human users of the system. Many of these systems use humans for the selection stage, but some allow users to exert more direct control over the design of the evolving artifact—for example, *TechnoSphere* and *Nerve Garden* (Section 4), *DrawCompileEvolve* (Section 6.1.2) and *BrainCrafter* (Section 6.2.1). The *Organic Builder* project (Section 5.1.1) is another example of a human computation system, this time based upon artificial chemistries rather than evolutionary systems. Elsewhere, the web-based survey of the ALife community reported by Rasmussen and colleagues (discussed at the start of Section 5) provides a somewhat different example of using the Web to harness the collective intelligence of a distributed group of users.

There is an increasingly large and active research community investigating methods for combining human and computational intelligence in appropriate ways so as to leverage and complement the strengths of both. Kosorukoff presented an interesting early exploration of different ways in which humans and computers can be deployed to create hybrid evolutionary algorithms [48]. There is a large literature on the more general areas of *human computation* and *crowd creativity*: for good reviews, see [60], [59], [90] and [136].¹⁴¹ An ALife-oriented review of *crowdsourcing* and discussion of important factors determining the success of crowdsourcing platforms, together with a report of simulation studies investigating such parameters, can be found in [13].

As web-based distributed computation systems become more powerful (see Section 7.2) and infrastructure frameworks such a WIN (Section 6.4.3) emerge to simplify the devel-

¹⁴¹An interesting recent study that conceptualizes human decision making and creativity as evolutionary computation is described by Sayama and Dionne [105]. Elsewhere, Woolley and Stanley have also recently explored novel architectures for human–computer collaboration within an ALife context [135].

opment of such systems, human–computer hybrid architectures will undoubtedly remain a prominent feature of much WebAL research in future years.

7.2 Distributed computation

The kind of web-based human computation described above is a special case of the more general concept of *distributed computation*. As mentioned in Section 3, theoretical and practical work on distributed artificial evolution goes back many decades. It is now becoming increasingly possible to use the Web as a distributed computation platform: HTML5 and related APIs such as *WebSocket*,¹⁴² *Web Workers*¹⁴³ and *Web Storage*¹⁴⁴ allow these kinds of distributed computation systems to be implemented using native technology.

At the same time, a number of technologies are currently being developed to allow fast client-side processing at speeds approaching those of local compiled code: Mozilla’s *asm.js*,¹⁴⁵ and Google’s *Native Client*,¹⁴⁶ are the most prominent of these. The *Emscripten* compiler,¹⁴⁷ used in several of the projects discussed in Section 6, converts C and C++ code to highly efficient JavaScript (using *asm.js*). Furthermore, the W3C has recently launched a *WebAssembly Community Group*¹⁴⁸ that aims to develop an efficient low-level programming language for fast client-side applications.¹⁴⁹ An alternative form of web-based distributed computation is provided by cloud computing platforms such as the established *Amazon Web Services*¹⁵⁰ and the more recent *Google Cloud Platform*.¹⁵¹

Elsewhere, David and Elena Ackley have advocated a more radical approach to designing distributed, scalable computational architectures that embrace stochastic events rather than attempt to prevent them at all costs [4]. As such, this hardware-based approach has a strong ALife flavor that could be relevant to future WebAL work and ideally bring WebAL techniques into more mainstream applications.

All of these developments contribute to easier implementation and faster execution of client-side processing and web-based distributed computation systems. We can therefore expect to see more WebAL projects along these lines in the near future.

¹⁴²<http://www.w3.org/TR/websockets/>

¹⁴³<http://www.w3.org/TR/workers/>

¹⁴⁴<http://www.w3.org/TR/webstorage/>

¹⁴⁵<http://en.wikipedia.org/wiki/Asm.js>

¹⁴⁶http://en.wikipedia.org/wiki/Google_Native_Client

¹⁴⁷<https://github.com/kripken/emscripten>

¹⁴⁸<https://www.w3.org/community/webassembly/>

¹⁴⁹<https://en.wikipedia.org/wiki/WebAssembly>

¹⁵⁰<http://aws.amazon.com/> (as used, for example, in Radenski’s work [91] discussed in Section 6.4.5)

¹⁵¹<https://cloud.google.com/>

7.3 Persistent systems

Traditional ALife experiments typically run for a few hours, days, or maybe weeks on a local machine or compute cluster: data are collected, results are written up, and no further experimentation is done. A feature of many of the web-based ALife systems reviewed here is that they are designed to run indefinitely, for as long as there are users who are interested in interacting with them. This requirement represents a profound change in the way that experiments are designed, showing some parallels with long-running evolution studies of real biological systems such as Richard Lenski's *E. coli* long-term evolution experiment [29]. Such a shift in methodology also introduces challenges in developing appropriate methods for data capture and analysis—we expect improvements in these areas to be a feature of future work in the area. Furthermore, using the HTML5 APIs mentioned above for client-side processing and data storage, or cloud-based processing and data-storage services, these systems can potentially be massively distributed and extended across space as well as time. Systems such as *Pfeiffer* (Section 5.1.4) and *Picbreeder* (Section 5.1.1) give some indication of the potential benefits of web-based experiments, and many other types of long-term experiment can be imagined for future projects.

7.4 Cumulative progress: Building upon past results

An important aspect of some of the projects reviewed, in addition to employing human computation and distributed, persistent systems, is that they allow users to *build upon the results* of previous users' work. Early examples include the *International Interactive Genetic Art* series and the *Electric Sheep* project in the 1990s (see Section 4). The enthusiast forums that developed around the *Creatures* game, allowing users to exchange their *Norns* with other users, can also be seen in this light (Section 4). This approach was introduced more explicitly in *Picbreeder*, with its *branching* method that allowed users to select and continue evolving images previously evolved by other users (Section 5.1.1). The branching process is also central to the more recent *EndlessForms* project (Section 6.1.1). From a more general perspective, open, web-based science frameworks such as COEL (Section 6.4.2) and WIN (Section 6.4.3), which provide data-archiving facilities for previous experiments, facilitate the process of using and building upon past results by others. Furthermore, platforms such as *YouShare* (Section 6.4.1) facilitate archiving and reuse of whole software systems from previous research projects. Thus, these kinds of web-based platforms can potentially provide tremendous advantages for pursuing collaborative WebAL research (and similarly for many other branches of science as well).

7.5 Open science, open education and public outreach

Many of the systems reviewed here, particularly in Section 6.3 (Science and Education) and Section 6.4 (Frameworks for WebAL), are designed to be platforms for open science or open education. Others are focused on public outreach and communication of ALife concepts. These systems are generally available for free, and, being web-based, require no installation—most of this work now uses native technology rather than relying on browser plugins, which might have been the case in the past. This work hints at the huge potential offered by web-based systems to open up science and education in ALife (and, of course, other topics too) to a much wider audience. Over the coming years we will surely continue to witness great innovations in the way that research, education and outreach is conducted on the Web, in ALife as in many other academic disciplines.¹⁵²

7.6 The Web as an arena for multi-user competitions

In addition to providing mechanisms for collaboration, Web platforms can also provide arenas in which competitions can be held between agents submitted by multiple users. Examples of work employing this kind of approach include *Sodarace* (Section 4), *NERO* (Box 2) and *EvoCommander* (Section 6.2.1). Taking a somewhat different approach, *Galactic Arms Race (GAR)* (Box 2) investigated how in-game content can be evolved based upon the behavior of multiple concurrent online players. The user challenges set in *Organic Builder* (Section 5.1.1) can also be seen as a form of multi-user competition, although the discussion of results on the online forum also gave the project a more cooperative flavor. It is likely that many additional forms of multi-user games and competitions will be explored in future WebAL projects.

7.7 Client-centric ALife: WebAL “in the wild”

Some of the early publications on WebAL, such as Ray’s paper on *Network Tierra* [95] (Box 1) and Langdon’s paper on *Pfeiffer* [49] (Section 5.1.4), discuss the possibility of ALife agents roaming the Internet and evolving in the complex environment that it provides. In the architectures employed by most of the work described here, the system’s server plays a central role in performing computations, farming data to client browsers and receiving the results. However, one can imagine much more client-centric architectures whereby artificial lifeforms exist primarily on users’ machines (made persistent between browser sessions by client-side storage mechanisms such as cookies or the new *Web Storage API*), and the server is used primarily as a routing system to allow organisms to migrate from one client to another (which can now be implemented natively

¹⁵²See [72] for a good review and outlook on the topic of *citizen science*, which intersects open science, education and outreach and is therefore relevant to the future direction of the WebAL projects discussed here.

using the *WebSocket* API).¹⁵³ There are glimmers of this approach in the *Pfeiffer* system (Section 5.1.4), and it is more strongly emphasized in the recent exploratory work on *EvoPopcorn* (Section 6.5.1) and in the forthcoming augmented reality mobile games *Polly Peck's Journey* and *TechnoSphere 2.0* (Section 6.2.2). It seems likely that this kind of ability for free-roaming agency—for WebAL “in the wild”—could be explored and exploited much more thoroughly, and we expect many more developments in this area in future work. However, for this kind of work to truly flourish, important concerns about safety, security, and preventing free-roaming ALife agents from evolving out of control—concerns raised by, among others, Chris Langton in the early days of ALife (see Section 3)—must be fully addressed.

7.8 Cloud APIs

The work by Auerbach [9] (Section 6.1.2) illustrates an approach to utilizing cloud interfaces and APIs (in his case, *Google Search by Image*) as a component of a computational intelligence system. Of the work reviewed here, this is the only example of a WebAL system that directly utilizes a cloud API to provide an intelligent component to its architecture. Nguyen and colleague's work [73], also described in Section 6.1.2, makes indirect use of Amazon's *Mechanical Turk* web-based crowdsourcing platform. It is not hard to imagine ways in which future WebAL systems could make more direct use of *Mechanical Turk*. More broadly, we can imagine many other ways in which cloud APIs could be employed to provide enhanced capabilities to WebAL systems. As the number of cloud services and open data sources continue to expand and Web API ecosystem architectures mature, we can expect to see many more WebAL systems using this kind of approach in future.

7.9 Crowdfunding

While not related to WebAL technology as such, another important way in which the Web can enhance ALife is through *crowdfunding* of research and applications.

In 2011, Steve Grand (author of the *Creatures* game discussed in Section 4) successfully secured *Kickstarter* funding of nearly US\$57,000 to develop a new ALife powered game called *Grandroids*, currently still under development.¹⁵⁴ More recently, the *OpenWorm* project (see Section 6.3.2.5) has raised substantial funds through crowdfunding efforts, including over US\$120,000 through a *Kickstarter* campaign in 2014. Another example is the company *Wiggle Planet*, discussed in Section 6.2.2, who raised over US\$15,000 in 2014 through a *Kickstarter* campaign to develop their augmented reality ALife game.

¹⁵³In other words, effectively a persistent, native, worldwide, peer-to-peer (P2P) architecture.

¹⁵⁴The initial target platforms for the game are PCs and Macs. For more details, see <https://www.kickstarter.com/projects/1508284443/grandroids-real-artificial-life-on-your-pc>

Between them, these three projects have raised nearly US\$200,000 of funding through *Kickstarter*. These examples demonstrate that it is possible (although still far from easy) to obtain substantial funding for ALife projects via crowdfunding.

8 Conclusion

We have explored many different ways in which web technologies and ALife techniques have intersected in past and current work. The projects we have reviewed demonstrate the diverse domains in which WebAL has been applied, including: collaborative design; human computation; education; outreach; persistent and long-running experiments; the archiving, sharing, reproduction, and reuse of scientific experiments and platforms; collaborative open science; art; computer games; crowdfunding; and more besides.

As web technology continues to develop, and particularly with the move towards native APIs in place of proprietary plugins, the potential for developing complex web-based ALife research and applications grows greater each year.

Whether or not a WebAL project is primarily focused on education or public outreach, the very nature of the Web means that WebAL research is inherently open and can reach a wide audience (unless steps are taken to actively prevent this accessibility). As funding councils around the world place increasing emphasis on the public understanding of science, WebAL is well placed to play an important role in the communication of ALife research to a wide and diverse audience. Furthermore, WebAL not only enables wide dissemination of results, but it also promotes public *engagement* and *participation* with ALife research.

Looking back over the research reviewed here, it is clear that great strides have been made over the last 21 years. However, as web technology and APIs develop, it is likely that current work represents only the tip of the iceberg of what could be possible. The work surveyed here represents a great showcase of some of the possibilities of WebAL, and yet we suspect there are many other possibilities, some as yet unimagined. Advances will doubtless be made in all of the areas outlined in our discussion of emerging themes and future directions (Section 7), and likely in completely different areas as well.

It is a truly exciting time to be involved in WebAL research. We expect the current rapid pace of development to continue, and indeed to accelerate, over the next few years. We look forward to witnessing the advances and achievements, both expected and unexpected, that will emerge from these efforts.

Acknowledgements

We would like to thank two anonymous reviewers for their insightful comments and constructive criticism on an earlier draft of this paper.

Thanks to Alan Dorin for providing references to relevant early work on ALife art. Thanks also to Bruce Damer, Steve Grand, Tim Hutton, Jane Prophet, Craig Reynolds, Hiroki Sayama, Terence Soule and Sebastian von Mammen for verifying details relating to their work.

We gratefully acknowledge use of Overleaf, the free, online collaborative \LaTeX authoring tool (<https://www.overleaf.com/>), in preparing the draft of this paper.

Tim Taylor and Simon Hickinbotham acknowledge funding from the EU FP7 project EvoEvo, grant number 610427. Joshua Auerbach acknowledges funding from the EU FP7 project Insight, grant agreement number 308943. Jeff Clune was supported by an NSF CAREER award, grant number 1453549. Charles Ofria acknowledges funding for the BEACON Center for the Study of Evolution in Action from the U.S. National Science Foundation, under Cooperative Agreement DBI-0939454. Jason Yosinski was supported by a NASA Space Technology Research Fellowship.

Image credits: Figure 2 courtesy of Jane Prophet, used with permission; Figure 3 courtesy of <http://www.biota.org/nervegarden/publish.html>, used according to the licensing information on that page; Figure 6 is a screenshot taken by one of the authors (TT), used with permission of Bill Langdon; Figure 4 courtesy of Tim Hutton, used with permission; Figure 13 courtesy of Sebastian von Mammem, used with permission. All other figures were provided by the authors.

A Annotated Bibliography of Other NetAL Work

Details of papers discovered in the systematic search of the Artificial Life journal and conference proceedings that have not been covered elsewhere in this review paper. For the purposes of this review, they were deemed to be either insufficiently focused on the Web (i.e. NetAL rather than WebAL), or insufficiently focused on ALife. They are listed here for completeness. For details of the search methodology adopted, see Section 2.

- [A1] M. Aguilera, I. Morer, X. E. Barandiaran, and M. G. Bedia. Quantifying political self-organization in social media. Fractal patterns in the spanish 15m movement on Twitter. In P. Liò, O. Miglino, G. Nicosia, S. Nolfi, and M. Pavone, editors, *Advances in Artificial Life, ECAL 2013*, pages 395–402. MIT Press, 2013.

An analysis of fractal scaling, self-organizing processes and different types of noise in real-world Twitter interaction networks.

- [A2] M. L. Best. An ecology of text: using text retrieval to study alife on the net. *Artificial life*, 3(4):261–287, 1997.

A study of ecological and evolutionary dynamics of posts to Usenet groups, using analysis tools including Latent Semantic Indexing.

- [A3] M. L. Best. Models for interacting populations of memes: Competition and niche behavior. In P. Husbands and I. Harvey, editors, *Fourth European Conference on Artificial Life*, pages 154–163. MIT Press, 1997.

Extension of work described in A[2], examining competition and niche behaviour.

- [A4] J. Bollen, B. Gonçalves, G. Ruan, and H. Mao. Happiness is assortative in online social networks. *Artificial Life*, 17(3):237–251, 2011.

An empirical study of the subjective well-being (SWB) of Twitter users finds that SWB is assortative across the Twitter social network.

- [A5] M. Humphrys. Distributing a mind on the internet: The world-wide-mind. In J. Kelemen and P. Sosík, editors, *Advances in Artificial Life: 6th European Conference, ECAL 2001, Prague, Czech Republic, September 10-14, 2001. Proceedings*, pages 669–680. Springer, 2001.

Proposes an architecture for collaborative, distributed AI research on the Internet whereby systems can be packaged as "mind servers" and made available for others to utilize as components in their own work.

- [A6] J. O. Kephart, J. E. Hanson, and J. Sairamesh. Price and niche wars in a free-market economy of software agents. *Artificial Life*, 4(1):1–23, 1998.

An investigation into mechanisms for information brokerage in the context of free-market economies of software agents on the Internet.

- [A7] K.-J. Kim and S.-B. Cho. A comprehensive overview of the applications of artificial life. *Artificial Life*, 12(1):153–182, 2006.

Section 3.8 of paper covers Internet applications, including some web-based systems that have been covered in the current paper, but other non-web-based systems that have not been covered here.

- [A8] H. Noser, I. Pandzic, T. Capin, N. Thalmann, and D. Thalmann. Playing games through the virtual life network. In C. G. Langton and K. Shimohara, editors, *Artificial Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*, pages 135–142. MIT Press, 1996.

Describes a shared, networked 3D virtual environment in which human actors can interact with autonomous agents and other humans. The system was run over an ATM (asynchronous transfer mode) pilot network, provided by Swiss Telecom, and experiments conducted with participants in Singapore and Geneva.

- [A9] C. O’Leary and M. Humphrys. Building a hybrid society of mind using components from ten different authors. In *Advances in Artificial Life: 7th European Conference, ECAL 2003, Dortmund, Germany, September 14-17, 2003, Proceedings*, pages 839–846. Springer, 2003.

This work builds upon A[5] by describing experiments in which components developed by ten different authors and made available as web services were used to solve a simple problem.

- [A10] F. Saffre and M. Shackleton. "Embryo": an autonomic co-operative service management framework. In S. Bullock, J. Noble, R. A. Watson, and M. A. Bedau, editors, *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 513–520. MIT Press, 2008.

Describes the design and simulation results of a decentralized service management framework aimed at providing self-managing (autonomic) architectures for application areas including web services and pervasive computing.

- [A11] D. Stow and C. Roadknight. Antigens, antibodies, and the world wide web. In J. Kelemen and P. Sosík, editors, *Advances in Artificial Life: 6th European Conference, ECAL 2001, Prague, Czech Republic, September 10-14, 2001. Proceedings*, pages 161–165. Springer, 2001.

Reports preliminary investigations in using an Artificial Immune System to recognize patterns of URLs taken from a cache log file.

References

- [1] D. H. Ackley. ccr: A network of worlds for research. In C. G. Langton and K. Shimohara, editors, *Artificial Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*, pages 116–123. MIT Press, 1996.
- [2] D. H. Ackley. Real artificial life: Where we may be. In M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen, editors, *Artificial Life VII: Proceedings of the Seventh International Conference on Artificial Life*, pages 487–496. MIT Press, 2000.
- [3] D. H. Ackley. Bespoke physics for living technology. *Artificial Life*, 19(3-4):347–364, 2013.
- [4] D. H. Ackley and E. S. Ackley. Artificial life programming in the robust-first attractor. In P. Andrews, L. Caves, R. Doursat, S. Hickinbotham, F. Polack, S. Stepney, T. Taylor, and J. Timmis, editors, *Proceedings of the European Conference on Artificial Life 2015*, pages 554–561. MIT Press, 2015.
- [5] D. P. Anderson, J. Cobb, E. Korpela, M. Lebofsky, and D. Werthimer. Seti@ home: an experiment in public-resource computing. *Communications of the ACM*, 45(11):56–61, 2002.
- [6] J. C. Astor and C. Adami. Development and evolution of neural networks in an artificial chemistry. In *Proceedings of the Third German Workshop on Artificial Life*, Bochum, 1998. Verlag Harri Deutsch. arxiv preprint: <http://arxiv.org/abs/adap-org/9807003>.
- [7] J. C. Astor and C. Adami. A developmental model for the evolution of artificial neural networks. *Artificial Life*, 6(3):189–218, 2000.
- [8] J. Auerbach, D. Aydin, A. Maesani, P. Kornatowski, T. Cieslewski, G. Heitz, P. Fernando, I. Loshchilov, L. Daler, and D. Floreano. RoboGen: Robot Generation through Artificial Evolution. In *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 136–137. The MIT Press, 2014.
- [9] J. E. Auerbach. Automated evolution of interesting images. In C. Adami, D. M. Bryson, C. Ofria, and R. T. Pennock, editors, *Artificial Life 13: Proceedings of the Thirteenth International Conference on the Simulation and Synthesis of Living Systems*, pages 629–630. MIT Press, 2012.
- [10] S. P. Balfour. Assessing writing in MOOCs: Automated essay scoring and calibrated peer review. *Research and Practice in Assessment*, 8:40–48, 2013.
- [11] P. Banda, D. Blount, and C. Teuscher. COEL: A web-based chemistry simulation framework, 2014. <http://arxiv.org/abs/1407.4027>.
- [12] M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen. Living technology: Exploiting life’s principles in technology. *Artificial Life*, 16(1):89–97, 2010.
- [13] T. Buecheler, R. Lonigro, R. M. Fuchsli, and R. Pfeifer. Modeling and simulating crowdsourcing as a complex biological system: Human crowds manifesting collective intelligence on the internet. In T. Lenaerts, M. Giacobini, H. Bersini, P. Bourguine, M. Dorigo, and R. Doursat, editors, *Advances in Artificial Life, ECAL 2011*, pages 109–116. MIT Press, 2011.
- [14] E. Cantú-Paz. A survey of parallel genetic algorithms. *Calculateurs parallèles, réseaux et systèmes répartis*, 10(2):141–171, 1998.
- [15] S. Carlson. Boids of a feather flock together. *Scientific American*, 283(5):112–114, 2000.

-
- [16] D. Cliff and S. Grand. The *Creatures* global digital ecosystem. *Artificial Life*, 5(1):77–94, 1999.
- [17] J. Clune and H. Lipson. Evolving three-dimensional objects with a generative encoding inspired by developmental biology. In T. Lenaerts, M. Giacobini, H. Bersini, P. Bourguine, M. Dorigo, and R. Doursat, editors, *Advances in Artificial Life, ECAL 2011*, pages 141–148. MIT Press, 2011.
- [18] B. Damer and D. Deamer. Coupled phases and combinatorial selection in fluctuating hydrothermal pools: a scenario to guide experimental approaches to the origin of cellular life. *Life*, 5(1):872–887, 2015.
- [19] B. Damer, K. Marcelo, F. Revi, T. Furmanski, and C. Laurel. Nerve garden: Germinating biological metaphors in net-based virtual worlds. In A. Adamatzky and M. Komosinski, editors, *Artificial Life Models in Software*, chapter 3, pages 67–80. Springer, 2005.
- [20] B. Damer, P. Newman, R. Gordon, T. Barbalet, D. Deamer, and R. Norkus. The EvoGrid: A framework for distributed artificial chemistry cameo simulations supporting computational origins of life endeavors. In H. Fellersmann, M. Dörr, M. M. Hanczy, L. L. Laursen, S. Maurer, D. Merkle, P.-A. Monnard, K. Støy, and S. Rasmussen, editors, *Artificial Life XII: Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*, pages 73–79. MIT Press, 2010.
- [21] B. Damer, P. Newman, R. Norkus, J. Graham, R. Gordon, and T. Barbalet. Cyberbiogenesis and the EvoGrid: a twenty-first century grand challenge. In J. Seckbach, editor, *Genesis – In The Beginning*, pages 267–288. Springer, 2012.
- [22] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009.
- [23] P. J. Denning. The science of computing: The internet worm. *American Scientist*, 77(2):126–128, 1989.
- [24] P. K. Dick. *Do Androids Dream of Electric Sheep?* Doubleday, 1968.
- [25] S. Draves. Evolution and collective intelligence of the electric sheep. In J. Romero and P. Machado, editors, *The Art of Artificial Evolution*, Natural Computing Series, pages 63–78. Springer Berlin Heidelberg, 2008. http://dx.doi.org/10.1007/978-3-540-72877-1_3.
- [26] S. Durland. Defining the image as place: A conversation with kit galloway, sherrie rabinowitz and gene youngblood. *High Performance*, 37:52–59, 1987. Text available at http://www.ecafe.com/museum/hp_gy_1987/hp_gy_1987.html.
- [27] D. Farrell, P. Kostkova, L. Lazareck, D. Weerasinghe, J. Weinberg, D. M. Lecky, N. Adriaenssens, T. K. Herotová, J. Holt, P. Touboul, et al. Developing e-Bug web games to teach microbiology. *Journal of Antimicrobial Chemotherapy*, 66(suppl 5):v33–v38, 2011.
- [28] S. Forrest, J. Balthrop, M. Glickman, and D. Ackley. Computation in the wild. In E. Jen, editor, *Robust Design: A Repertoire of Biological, Ecological, and Engineering Case Studies*, pages 207–230. Oxford University Press, 2005.
- [29] J. W. Fox and R. E. Lenski. From here to eternity—the theory and practice of a really long experiment. *PLOS Biology*, 13(6):e1002185, 2015.
- [30] P. García-Sánchez, J. González, P. A. Castillo, M. G. Arenas, and J. Merelo-Guervós. Service oriented evolutionary algorithms. *Soft Computing*, 17(6):1059–1075, 2013.

- [31] M. García-Valdez, L. Trujillo, F. Fernández de Vega, J. J. Merelo Guervós, and G. Olague. EvoSpace: A Distributed Evolutionary Platform Based on the Tuple Space Model. In A. I. Esparcia-Alcázar, editor, *Applications of Evolutionary Computation*, volume 7835 of *Lecture Notes in Computer Science*, pages 499–508. Springer, 2013.
- [32] K. Goldberg, M. Mascha, S. Gentner, N. Rothenberg, C. Sutter, and J. Wiegley. Desktop teleoperation via the world wide web. In *Proceedings of the 1995 IEEE International Conference on Robotics and Automation*, pages 654–659, 1995.
- [33] H. Goldsby, A. Dornhaus, B. Kerr, and C. Ofria. Task switching costs promote the evolution of division of labor and shifts in individuality. *Proceedings of the National Academy of Sciences, USA*, 34(109):13686–13691, 2012.
- [34] E. J. Hastings, R. K. Guha, and K. O. Stanley. Evolving content in the galactic arms race video game. In *Proceedings of the IEEE Symposium on Computational Intelligence and Games (CIG09)*, pages 241–248, Piscataway, NJ, 2009. IEEE.
- [35] E. J. Hastings, R. K. Guha, and K. O. Stanley. Interactive evolution of particle systems for computer graphics and animation. *IEEE Transactions on Evolutionary Computation*, 13(2):418–432, 2009.
- [36] S. Hickinbotham, E. Clark, S. Stepney, T. Clarke, A. Nellis, M. Pay, and P. Young. Specification of the stringmol chemical programming language version 0.2. Technical Report YCS-2010-458, University of York, Apr 2012.
- [37] S. Hickinbotham, M. Weeks, and J. Austin. The ALife Zoo: cross-browser, platform-agnostic hosting of artificial life simulations. In P. Liò, O. Miglino, G. Nicosia, S. Nolfi, and M. Pavone, editors, *Advances in Artificial Life, ECAL 2013*, pages 71–78. MIT Press, 2013.
- [38] S. J. Hickinbotham, E. Clark, S. Stepney, T. Clarke, A. Nellis, M. Pay, and P. Young. Diversity from a monoculture-effects of mutation-on-copy in a string-based artificial chemistry. In H. Feller-mann, M. Dörr, M. M. Hanczy, L. L. Laursen, S. Maurer, D. Merkle, P.-A. Monnard, K. Støy, and S. Rasmussen, editors, *Artificial Life XII: Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*, pages 24–31. MIT Press, 2010.
- [39] T. Hutton. The organic builder: A public experiment in artificial chemistries and self-replication. *Artificial Life*, 15(1):21–28, 2009.
- [40] D. Jallof, S. Risi, and J. Togelius. Evocommander: A novel game based on evolving and switching between artificial brains. *IEEE Transactions on Computational Intelligence and AI in Games*, PP(99):1–1, 2016.
- [41] D. Jepsen. They’re dead, Jim: Natural selection goes awry in this depressing update. In *Computer Gaming World*, volume 174, pages 364–366. Ziff Davis, January 1999.
- [42] N. Johnson, S. Rasmussen, C. Joslyn, L. Rocha, S. Smith, and M. Kantor. Symbiotic intelligence: self-organizing knowledge on distributed networks driven by human interaction. In C. Adami, R. Belew, H. Kitano, and C. Taylor, editors, *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*, pages 403–407. MIT press, 1998.
- [43] B. Jónsson, A. Hoover, and S. Risi. Interactively evolving compositional sound synthesis networks. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*. ACM, 2015.
- [44] K. Kelly. *Out of Control: The New Biology of Machines, Social Systems and the Economic World*. Addison Wesley, 1994.

-
- [45] J. O. Kephart and D. M. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.
- [46] F. Khatib, F. DiMaio, S. Cooper, M. Kazmierczyk, M. Gilski, S. Krzywda, H. Zabranska, I. Pichova, J. Thompson, Z. Popović, et al. Crystal structure of a monomeric retroviral protease solved by protein folding game players. *Nature structural & molecular biology*, 18(10):1175–1177, 2011.
- [47] M. Klusch, editor. *Intelligent Information Agents: Agent-Based Information Discovery and Management on the Internet*. Springer, 1999.
- [48] A. Kosorukoff. Human based genetic algorithm. In *IEEE International Conference on Systems, Man, and Cybernetics*, volume 5, pages 3464–3469. IEEE, 2001.
- [49] W. B. Langdon. Pfeiffer: A distributed open-ended evolutionary system. In B. Edmonds, N. Gilbert, S. Gustafson, D. Hales, and N. Krasnogor, editors, *AISB'05: Proceedings of the Joint Symposium on Socially Inspired Computing (METAS 2005)*, pages 7–13, 2005. ISBN 1 902956 48 4.
- [50] C. G. Langton. Introduction. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 3–24. Addison-Wesley, 1991.
- [51] J. L. J. Laredo, P. Bouvry, D. L. González, F. F. de Vega, M. G. Arenas, J. Merelo, and C. M. Fernandes. Designing robust volunteer-based evolutionary algorithms. *Genetic Programming and Evolvable Machines*, 15(3):221–244, 2014.
- [52] M. Lewis. Evolutionary visual art and design. In J. Romero and P. Machado, editors, *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, pages 3–37. Springer, 2008.
- [53] S. L. Lim and P. J. Bentley. App epidemics: Modelling the effects of publicity in a mobile app ecosystem. In C. Adami, D. M. Bryson, C. Ofria, and R. T. Pennock, editors, *Artificial Life 13: Proceedings of the Thirteenth International Conference on the Simulation and Synthesis of Living Systems*, pages 202–209. MIT Press, 2012.
- [54] S. L. Lim and P. J. Bentley. How to be a successful app developer: Lessons from the simulation of an app ecosystem. In *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation*, GECCO '12, pages 129–136, New York, NY, USA, 2012. ACM.
- [55] S. L. Lim and P. J. Bentley. Investigating app store ranking algorithms using a simulation of mobile app ecosystems. In *Evolutionary Computation (CEC), 2013 IEEE Congress on*, pages 2672–2679, June 2013.
- [56] C. J. Lintott, K. Schawinski, A. Slosar, K. Land, S. Bamford, D. Thomas, M. J. Raddick, R. C. Nichol, A. Szalay, D. Andreescu, et al. Galaxy zoo: morphologies derived from visual inspection of galaxies from the sloan digital sky survey. *Monthly Notices of the Royal Astronomical Society*, 389(3):1179–1189, 2008.
- [57] H. Lipson and J. B. Pollack. Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974–978, 2000.
- [58] H. Lowood. Stanford university libraries acquire mud1 source code, April 2014. <http://web.stanford.edu/group/htgg/cgi-bin/drupal/?q=node/1191>.
- [59] M. L. Maher. Design creativity research: From the individual to the crowd. In T. Taura and Y. Nagai, editors, *Design Creativity 2010 (Proceedings of the International Conference on Design Creativity)*, pages 41–47. Springer, 2010.

- [60] T. W. Malone, R. Laubacher, and C. Dellarocas. Harnessing crowds: Mapping the genome of collective intelligence. Technical Report 2009-001, MIT Center for Collective Intelligence, February 2009.
- [61] M. J. Mataric. Robotics education for all ages. In *Proceedings, AAAI Spring Symposium on Accessible, Hands-on AI and Robotics Education*, 2004.
- [62] M. L. McLaughlin, K. K. Osborne, and N. B. Ellison. Virtual community in a telepresence environment. In S. Jones, editor, *Virtual Culture: Identity and Communication in Cybersociety*, pages 146–168. Sage Publications, 1995.
- [63] P. W. McOwan and E. J. Burton. Sodarace: Adventures in artificial life. In A. Adamatzky and M. Komosinski, editors, *Artificial Life Models in Software*, chapter 5, pages 97–111. Springer, 2005.
- [64] F. Menczer, R. K. Belew, and W. Willuhn. Artificial life applied to adaptive information agents. In C. Knoblock and A. Levy, editors, *Information Gathering From Heterogeneous, Distributed Environments: Papers from the AAAI Spring Symposium*, pages 128–132. AAAI, 1995. Technical Report SS-95-08.
- [65] J.-J. Merelo, P. Castillo, A. Mora, A. Esparcia-Alcázar, and V. Rivas-Santos. NodeO, a multi-paradigm distributed evolutionary algorithm platform in JavaScript. In *GECCO Comp '14 Proceedings of the Companion Publication of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 1155–1162. ACM, 2014.
- [66] J. J. Merelo-Guervós and P. García-Sánchez. Modeling browser-based distributed evolutionary computation systems. *arXiv preprint arXiv:1503.06424*, 2015.
- [67] K. Meri, M. G. Arenas, A. M. Mora, J. Merelo, P. A. Castillo, P. García-Sánchez, and J. L. J. Laredo. Cloud-based evolutionary algorithms: An algorithmic study. *Natural Computing*, 12(2):135–147, 2013.
- [68] M. Métivier, C. Lattaud, and J.-C. Heudin. A stress-based speciation model in LifeDrop. In R. Standish, M. A. Bedau, and H. A. Abbass, editors, *Artificial Life VIII: Proceedings of the Eighth International Conference on Artificial Life*, pages 121–126. MIT Press, 2002.
- [69] L. Mignonneau and C. Sommerer. Creating artificial life for interactive art and entertainment. *Leonardo*, 34(4):303–307, 2001.
- [70] K. Moharreri, M. Ha, and R. H. Nehm. EvoGrader: an online formative assessment tool for automatically evaluating written evolutionary explanations. *Evolution: Education and Outreach*, 7(15), 2014.
- [71] J. M. Moore, A. J. Clark, and P. K. McKinley. Evolutionary robotics on the web with WebGL and Javascript. In *Proceedings of WebAL-1: Workshop on Artificial Life and the Web 2014, held at the 14th International Conference on the Synthesis and Simulation of Living Systems (ALIFE 14)*, Aug. 2014. <http://arxiv.org/abs/1406.3337>.
- [72] G. Newman, A. Wiggins, A. Crall, E. Graham, S. Newman, and K. Crowston. The future of citizen science: emerging technologies and shifting paradigms. *Frontiers in Ecology and the Environment*, 10(6):298–304, 2012.
- [73] A. Nguyen, J. Yosinski, and J. Clune. Innovation engines: Automated creativity and improved stochastic optimization via deep learning. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2015.

-
- [74] C. Ofria, D. M. Bryson, and C. O. Wilke. *Avida*. In *Artificial Life Models in Software, Second Edition*, pages 3–35. Springer, 2009.
- [75] M. Oka, H. Abe, and T. Ikegami. Dynamic homeostasis in packet switching networks. *Adaptive Behavior*, 23(1):50–63, 2015.
- [76] M. Oka, Y. Hashimoto, and T. Ikegami. Self-organization on social media: endo-exo bursts and baseline fluctuations. *PLOS ONE*, 9(10):e109293, 2014.
- [77] M. Oka, Y. Hashimoto, and T. Ikegami. Open-ended evolution in a web system. Late Breaking Papers at the European Conference on Artificial Life (ECAL 2015), 2015. <https://www.cs.york.ac.uk/nature/ecal2015/late-breaking/159.pdf>.
- [78] M. Oka and T. Ikegami. Exploring default mode and information flow on the web. *PLOS ONE*, 8(4):e60398, 2013.
- [79] P. T. Ølsted, B. Ma, and S. Risi. Interactive evolution of levels for a competitive multiplayer FPS. In *Proceedings of the 2015 IEEE Congress on Evolutionary Computation*. IEEE, 2015.
- [80] L. Pagliarini, A. Dolan, F. Menczer, and H. H. Lund. ALife Meets Web: Lessons Learned. In J.-C. Heudin, editor, *Virtual Worlds: First International Conference, VW'98 Paris, France, July 1–3, 1998 Proceedings*, pages 156–167. Springer, 1998.
- [81] L. Pagliarini, H. H. Lund, O. Miglino, and D. Parisi. Artificial life: A new way to build educational and therapeutic games. In C. G. Langton and K. Shimohara, editors, *Artificial Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*, pages 152–156. MIT Press, 1996.
- [82] A. Palyanov, S. Khayrulin, S. D. Larson, and A. Dibert. Towards a virtual *C. elegans*: A framework for simulation and visualization of the neuromuscular system in a 3D physical environment. *In Silico Biology*, 11(3):137–147, 2012.
- [83] S. Parkin. The man who made a game to change the world. Eurogamer (website), October 2014. <http://www.eurogamer.net/articles/2014-10-30-the-utopia-that-never-died>.
- [84] K. Paulsen. Image as Place: The Phenomenal Screen in Kit Galloway and Sherrie Rabinowitz’s *Satellite Arts 1977*. *Leonardo Electronical Almanac*, 19(2):98–111, 2013.
- [85] R. T. Pennock. Learning evolution and the nature of science using evolutionary computing and artificial life. *McGill Journal of Education*, 42:211–224, 2007.
- [86] J. Piskur, P. Greve, J. Togelius, and S. Risi. Braincrafter: An investigation into human-based neural network engineering. In *Proceedings of the IEEE Congress on Evolutionary Computation (IEEE CEC 2015)*, 2015.
- [87] J. Prophet. Sublime ecologies and artistic endeavors: Artificial life and interactivity in the online project TechnoSphere. *Leonardo*, 29(5):339–344, 1996.
- [88] J. Prophet. TechnoSphere: “real” time, “artificial” life. *Leonardo*, 34(4):309–312, 2001.
- [89] J. Prophet and H. Pritchard. SE Asian Ubicomp and ALife: Roaming and homing with TechnoSphere 2.0 computational companions. In *CHI 2015 Extended Abstracts (Workshop on Between the Lines: Reevaluating the Online/Offline Binary)*. ACM, 2015.

- [90] A. J. Quinn and B. B. Bederson. Human computation: A survey and taxonomy of a growing field. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1403–1412. ACM, 2011.
- [91] A. Radenski. Using MapReduce streaming for distributed life simulation on the cloud. In P. Liò, O. Miglino, G. Nicosia, S. Nolfi, and M. Pavone, editors, *Advances in Artificial Life, ECAL 2013*, pages 284–291. MIT Press, 2013.
- [92] S. Rasmussen, C. Knudsen, R. Feldberg, and M. Hindsholm. The Coreworld: Emergence and evolution of cooperative structures in a computational chemistry. *Physica D: Nonlinear Phenomena*, 42(1):111–134, 1990.
- [93] S. Rasmussen, M. J. Raven, G. N. Keating, and M. A. Bedau. Collective intelligence of the artificial life community on its own successes, failures, and future. *Artificial Life*, 9(2):207–235, 2003.
- [94] T. S. Ray. An approach to the synthesis of life. In *Artificial life II*, pages 371–408. Addison-Wesley, 1991.
- [95] T. S. Ray. A proposal to create a network-wide biodiversity reserve for digital organisms. Technical report, ATR, Japan, 1995. TR-H-133.
- [96] T. S. Ray. Selecting naturally for differentiation: Preliminary evolutionary results. *Complexity*, 3(5):25–33, 1998.
- [97] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *Computer Graphics (SIGGRAPH '87 Conference Proceedings)*, 21(4):25–34, 1987.
- [98] S. Risi. A compiler for CPPNs: Transforming phenotypic descriptions into genotypic representations. In *Proceedings of the 2013 AAAI Fall Symposium on How Should Intelligence be Abstracted in AI Research*, 2013.
- [99] S. Risi, J. Lehman, D. B. D’Ambrosio, R. Hall, and K. O. Stanley. Combining search-based procedural content generation and social gaming in the petalz video game. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE 2012)*, Menlo Park, CA, 2012. AAAI Press.
- [100] S. Risi, J. Lehman, D. B. D’Ambrosio, R. Hall, and K. O. Stanley. Petalz: Search-based procedural content generation for the casual gamer. *Computational Intelligence and AI in Games, IEEE Transactions on*, PP(99):1–1, 2015.
- [101] S. Risi, J. Lehman, D. B. D’Ambrosio, and K. O. Stanley. Automatically categorizing procedurally generated content for collecting games. In *Proceedings of the Workshop on Procedural Content Generation in Games (PCG) at the 9th International Conference on the Foundations of Digital Games (FDG-2014)*, New York, NY, USA, 2014. ACM.
- [102] S. Risi, J. Zhang, R. Taarnby, P. Greve, J. Piskur, A. Liapis, and J. Togelius. The case for a mixed-initiative collaborative neuroevolution approach. In *Proceedings of WebAL-1: Workshop on Artificial Life and the Web 2014, held at the 14th International Conference on the Synthesis and Simulation of Living Systems (ALIFE 14)*, Aug. 2014. <http://adsabs.harvard.edu/abs/2014arXiv1408.0998R>.
- [103] A. Rousset, B. Herrmann, C. Lang, and L. Philippe. A survey on parallel and distributed multi-agent systems. In *Euro-Par 2014: Parallel Processing Workshops*, pages 371–382. Springer, 2014.
- [104] H. Sayama. Swarm chemistry. *Artificial Life*, 15(1):105–114, 2009.

-
- [105] H. Sayama and S. D. Dionne. Studying collective human decision making and creativity with evolutionary computation, 2014. <http://arxiv.org/abs/1406.6291>.
- [106] M. Scheutz and J. Harris. An overview of the SimWorld agent-based grid experimentation system. In W. Dubitzky, K. Kurowski, and B. Schott, editors, *Large-Scale Computing Techniques for Complex System Simulations*, pages 59–80. Wiley, 2012.
- [107] M. Scheutz, P. Schermerhorn, R. Connaughton, and A. Dingler. SWAGES—an extendable distributed experimentation system for large-scale agent-based alife simulations. *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pages 412–419, 2006.
- [108] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, J. T. Folsom-Kovarik, and K. O. Stanley. Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary Computation*, 19(3):345–371, 2011.
- [109] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, and K. O. Stanley. Picbreeder: evolving pictures collaboratively online. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1759–1768. ACM, 2008.
- [110] K. Sims. Evolving 3d morphology and behavior by competition. In *Proceedings of Artificial Life IV*, pages 353–372, MA, 1994. MIT Press.
- [111] E. Sklar. Software review: NetLogo, a multi-agent simulation environment. *Artificial Life*, 13(3):303–311, 2007.
- [112] C. Sommerer and L. Mignonneau. VERBARIUM and LIFE SPACIES: Creating a visual language by transcoding text into form on the internet. In *IEEE Symposium on Visual Languages (VL'99) Conference Proceedings*, pages 90–95, Tokyo, 1999.
- [113] C. Sommerer and L. Mignonneau. Modeling emergence of complexity: the application of complex system and origin of life theory to interactive art on the internet. In M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen, editors, *Artificial Life VII: Proceedings of the Seventh International Conference on Artificial Life*, pages 547–554. MIT Press, 2000.
- [114] E. H. Spafford. Computer viruses as artificial life. *Artificial Life*, 1(3):249–265, 1994.
- [115] K. Spence. Genetic exhibitionism. *Wired*, 3.01:148, January 1995.
- [116] K. O. Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines Special Issue on Developmental Systems*, 8(2):131–162, 2007.
- [117] K. O. Stanley, B. D. Bryant, and R. Miikkulainen. Evolving neural network agents in the nero video game. In *Proceedings of the IEEE 2005 Symposium on Computational Intelligence and Games (CIG-05)*, pages 182–189, Piscataway, NJ, 2005. IEEE.
- [118] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127, 2002.
- [119] K. O. Stanley and R. Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal Artificial Intelligence Research (JAIR)*, 21:63–100, 2004.

- [120] L. Steels and F. Kaplan. Collective learning and semiotic dynamics. In D. Floreano, J.-D. Nicoud, and F. Mondada, editors, *Advances in Artificial Life: 5th European Conference, ECAL'99, Lausanne, Switzerland, September 13-17, 1999 Proceedings*, pages 679–688. Springer, 1999.
- [121] L. Steels, F. Kaplan, A. McIntyre, and J. Van Looveren. Crucial factors in the origins of word-meaning. In A. Wray, editor, *The Transition to Language*, pages 252–271. Oxford University Press, 2002.
- [122] K. Stockhausen. Helikopter-streichquartett. *Grand Street*, 56:213–225, 1996. Published by Jean Stein.
- [123] P. Szerlip and K. O. Stanley. Indirectly encoded Sodarace for artificial life. In P. Liò, O. Miglino, G. Nicosia, S. Nolfi, and M. Pavone, editors, *Advances in Artificial Life, ECAL 2013*, pages 218–225. MIT Press, 2013.
- [124] P. Szerlip and K. O. Stanley. Steps toward a modular library for turning any evolutionary domain into an online interactive platform. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, and H. Lipson, editors, *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 900–907. MIT Press, 2014.
- [125] P. Szor. *The Art of Computer Virus Research and Defense*. Addison-Wesley, 2005.
- [126] Y. Takeichi, K. Sasahara, R. Suzuki, and T. Arita. Twitter as social sensor: Dynamics and structure in major sporting events. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, and H. Lipson, editors, *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 778–784. MIT Press, 2014.
- [127] T. Taylor. Artificial life and the web: WebAL comes of age. In *Proceedings of WebAL-1: Workshop on Artificial Life and the Web 2014, held at the 14th International Conference on the Synthesis and Simulation of Living Systems (ALIFE 14)*, Aug. 2014. <http://arxiv.org/abs/1407.5719>.
- [128] T. Taylor, J. Auerbach, J. Bongard, J. Clune, S. Hickinbotham, and G. Hornby, editors. *Proceedings of WebAL-1: Workshop on Artificial Life and the Web 2014*, Aug. 2014. <http://arxiv.org/abs/1406.2507>.
- [129] H. W. Thimbleby, I. H. Whitten, and D. J. Pullinger. Concepts of cooperation in artificial life. *IEEE Transactions on Systems, Man, and Cybernetics*, 25(7):1166–1171, 1995.
- [130] S. von Mammen and S. Edenhofer. Swarm grammars GD: Interactive exploration of swarm dynamics and structural development. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, and H. Lipson, editors, *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 312–319. MIT Press, 2014.
- [131] M. Wagy and J. Bongard. Crowdfunding: A novel approach for designing bioinspired machines. In S. P. Wilson, P. F. Verschure, A. Mura, and T. J. Prescott, editors, *Biomimetic and Biohybrid Systems*, volume 9222 of *Lecture Notes in Computer Science*, pages 293–303. Springer, 2015.
- [132] A. Wait. The quantum coreworld: Competition and cooperation in an artificial ecology. In J. Pollack, M. A. Bedau, P. Husbands, R. A. Watson, and T. Ikegami, editors, *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, pages 280–285. MIT Press, 2004.
- [133] M. Whitelaw. *Metacreation: Art and Artificial Life*. MIT Press, 2004.

- [134] M. Witbrock and S. Neil-Reilly. Evolving genetic art. In P. Bentley, editor, *Evolutionary Design by Computers*, chapter 10, pages 251–259. Morgan Kaufmann, 1999.
- [135] B. G. Woolley and K. O. Stanley. A novel human-computer collaboration: combining novelty search with interactive evolution. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*, pages 233–240. ACM, 2014.
- [136] L. L. Yu, J. V. Nickerson, and Y. Sakamoto. Collective creativity: Where we are and where we might go. In T. W. Malone and L. von Ahn, editors, *Proceedings of Collective Intelligence*, 2012. <http://arxiv.org/abs/1204.2991>.
- [137] L. Zaman. EvoPopcorn: A web-native distributed artificial life simulation. In *Proceedings of WebAL-1: Workshop on Artificial Life and the Web 2014, held at the 14th International Conference on the Synthesis and Simulation of Living Systems (ALIFE 14)*, Aug. 2014. <http://arxiv.org/abs/1406.2507>.
- [138] J. Zhang, R. Taarnby, A. Liapis, and S. Risi. DrawCompileEvolve: Sparking interactive evolutionary art with human creations. In *Evolutionary and Biologically Inspired Music, Sound, Art and Design*, pages 261–273. Springer, 2015.