

Advertisement

gizmag

MORE TOP STORIES >

New Horizons reveals new Pluto mysteries

Laser device may soon non-invasively monitor diabetics' glucose levels

Apple iPod touch (6th-generation): Early impressions

Review: Parrot Bebop drone and Skycontroller

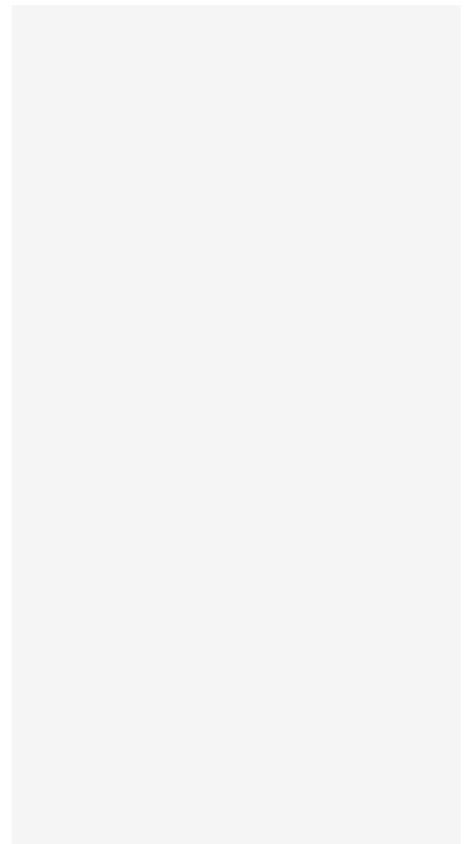
— COMPUTERS **FEATURE**

Creative AI: The robots that would be painters

By Richard Moss

040502 is a 2004 painting of pigment on paper by the robotic, artificially

Advertisement



intelligent painter AARON
Image Gallery (21 images)

Painting might be the last thing you'd expect computers to excel at. It's abstract, expressive, and tied to cultures, psychology, and subjectivity, whereas computers are objective, precise, and governed by the rules of mathematics. Painting, with its emotional reasoning and unclear meanings, appears to be the antithesis of a feeling, logical computer. But they aren't so far apart as they seem. Painting and other forms of visual art owe much to areas of mathematics such as geometry and perspective, and the algorithms that computers adhere to can in fact be made to generate images as varied and subtle as a human painter.



Much like its [musical counterpart](#), algorithmic art dates back to the time before computers were commonplace and in its purest sense requires no artificial intelligence whatsoever. You've probably seen examples of fractal art, which replicates patterns in a recursive, algorithmic way to often-stunning results that vary in appearance from geometric to organic to alien.

Traditionally, algorithmic art involves a human coming up with a concept that an algorithm then generates or visualizes either from scratch or based on existing material. An extreme example of this is Nagoya University researchers Yasuhiro Suzuki and Tomohiro Suzuki's [evolutionary painting algorithm](#), which takes example paintings of a given style and progressively mutates them – cutting and splicing and flipping elements, throwing out at each evolution any images that don't match the user's initial stylistic choices. But algorithmic art is more commonly used in the sense of images that are generated by computer code written by people like [Dextro](#), who is one of the leading practitioners of algorithmic/generative art.

As with music, game development, and writing, much of the attention from artists and scientists has been placed upon algorithms and intelligent tools that augment the artist's creativity. The [Processing](#) programming language was designed as an electronic sketchbook for artists and designers, while some of the better-known apps for algorithmic artists include [Ultra Fractal](#), [Scribble](#), and [Fragmentarium](#).

There are now over a dozen separate kinds of algorithmically-based art, including fractal art, genetic art, cellular automata, proceduralism, and transhumanist art. And there are multitudes of websites such as [The Algorists](#), [Algorithmic Worlds](#), and [The compArt database Digital Art](#) that celebrate the work of artists who use algorithms.



Advertisement

Advertisement

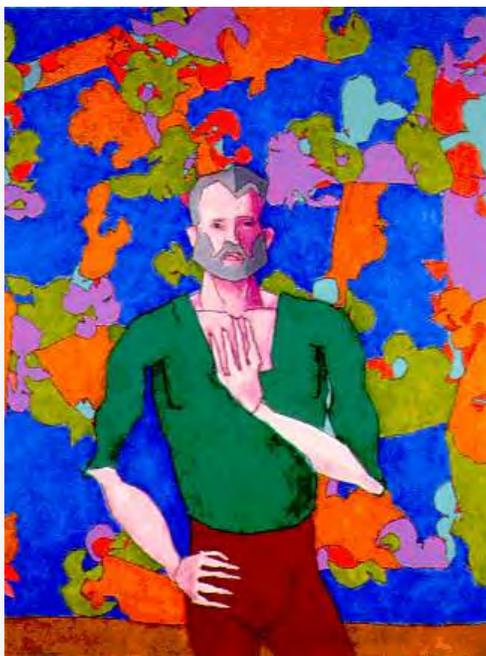


Harold Cohen watches AARON paint in 1995

But there are some who would teach computers to paint like humans, to push them beyond the point of being an extension of the artist and into the territory of artist themselves. The pioneer in this regard is a former artist and University of California San Diego professor called [Harold Cohen](#). He started working on an art-creating program called AARON in 1973, while a visiting scholar at Stanford University's Artificial Intelligence Lab.

AARON's capacity to paint improved year after year as its maker taught it more difficult or complex techniques. It learned to situate objects or people in 3D space in the 1980s, and could paint in color from 1990 onwards. In time its paintings found their way into many of the world's major art museums and onward into the hands of private collectors who paid hundreds or even thousands of dollars for AARON's art.

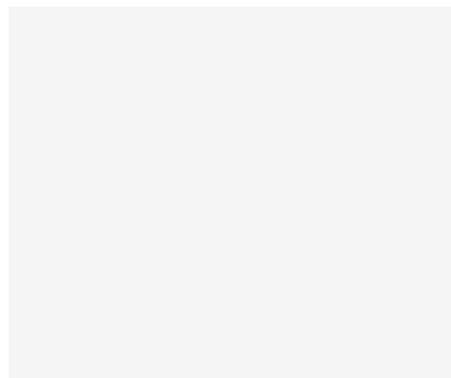
AARON paints not with pixels, we should note, but with real paint on an actual canvas. Cohen built a painting machine for his painting AI. He taught it to mix paint (fabric dyes, not oil), and even gave it an imagination of sorts. Enough of one, at least, that it can paint still life and portraits of human figures without photos or other human input as reference.



AARON learned to use color in a decorative motif in 1992 (Photo: Becky Cohen)

"I was always fascinated by the fact that we can look at a drawing and 'see' a face, when we are actually looking at a few dirty marks on a piece of paper," [he once told PBS](#) in an interview. It was through thinking about that mystery and how we humans develop the ability to generate those symbolic marks that AARON was born and further developed over a span of four decades.

Cohen never showed AARON any images, but rather taught his robot with lists of object/body elements and the relationships between them. Fundamental rules, essentially, that allow a robot that has never seen a human or a chair or flower to nevertheless paint something that looks like an abstract representation of those things.



Advertisement

THE TOUGHEST
CABLE ON EARTH



GET IT FOR \$24.99

Advertisement



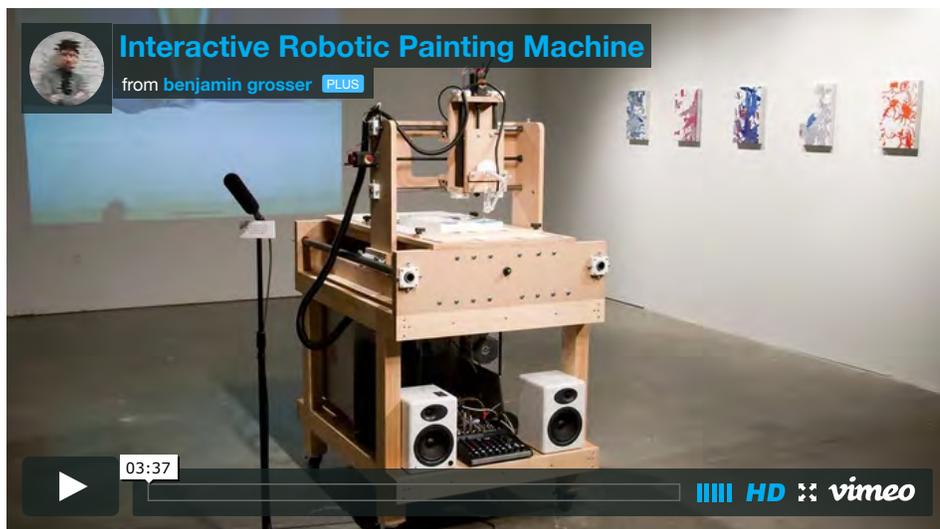
YOU'RE A VISIONARY.
SHOULDN'T YOUR
WEBSITE BE ONE, TOO?
PRE-ORDER NOW
LOCK IN YOUR EARLY ADOPTER DISCOUNT

AARON paints whatever it wants within the confines of its knowledge, driven only by its limited version of an imagination and not at any point instructed explicitly by Cohen or anybody else. And its end product, as you may have noticed from the images included in this article, has matured from childish lines and squiggles to portraits of people posing awkwardly near pot plants and most recently to abstract explorations of color and shape.

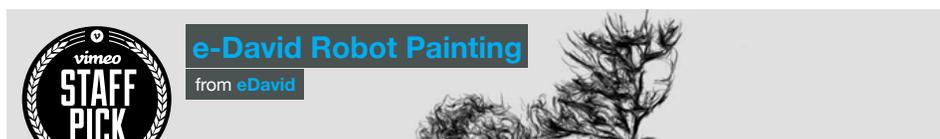


Final Approach, one of AARON's most recent paintings, from 2013

Others have tried to build on Cohen's pioneering work with their own AI painters. In 2011, for instance, Benjamin Grosser debuted his [Interactive Robotic Painting Machine](#), which paints abstract pictures with oil on canvas and, curiously, responds to the sounds in its environment, using those sounds of conversations, background audio, and its own mechanisms to influence the painting as it progresses. You can see a video of it in action below.



Another project, [e-David](#) by Oliver Deussen and Thomas Lindemeier from the University of Konstanz in Germany, uses a visual feedback loop to try to mimic how a human painter creates their work (see the video below).



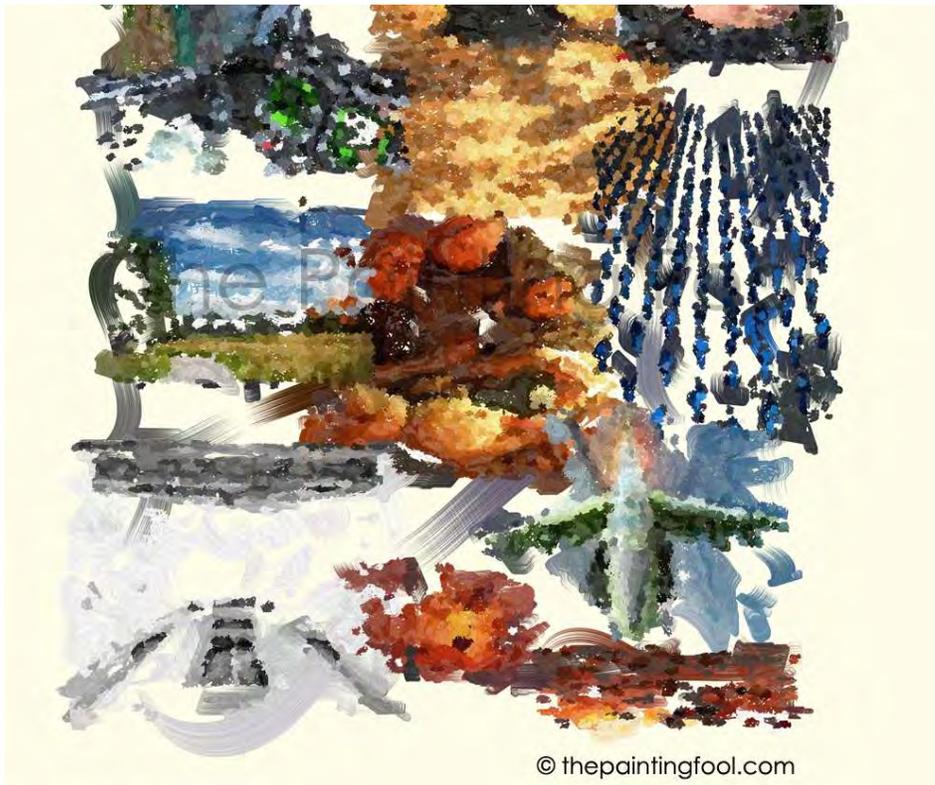


The idea, they explain on e-David's website, is that the technical aspects of painting are essentially "optimization processes in which color is manually distributed on a canvas until the painter is able to recognize the content, regardless if it is a representational painting or not." E-David does not concern itself with such lofty ambitions as creative intent, but another ambitious project has exactly that in mind.

Goldsmiths College, University of London professor [Simon Colton](#) brought [The Painting Fool](#) to life back in 2001, with his initial intention simply being to write software that could convert photos into art. The program gradually took on greater sophistication, moving from graphics filters and other elements you'd expect in programs like Photoshop into emotional modelling where photos could be altered according to keywords that were first supplied by an external agent and then self-determined by the AI.

After several years of development, The Painting Fool was granted an imagination – an ability to interpret ideas more loosely and to evolve its scene depictions without necessarily having any specific reference image – and then the ability to paint with intent. Colton describes one example where the software downloaded a news story on the war in Afghanistan and identified keywords that it used to find relevant images on Flickr. The Painting Fool then developed a painterly rendition of these images in a collage, juxtaposing a fighter plane with an explosion, a family, an Afghan girl, and a field of war graves.





Its intent here was flimsy in that it was more representing the journalist's work visually, but it was to Colton and his colleagues an indication that AI could indeed paint with intent if only you give it the ability to form opinions (especially through natural language processing and analysis of news media).

That's an important step because their long-term goal with The Painting Fool is to have it not be notable because it is a computer pretending to be an artist, as is the case now, but rather because it is a creative artist in its own right. They played with the concept of artists being moody creatures with an exhibition called "You Can't Know my Mind" in which The Painting Fool produced portraits of visitors that were affected by the previous 10 news articles it had read on The Guardian website. (It

was reading during the exhibit.)

If the software was reading a lot of articles with a very negative sentiment, it would refuse to paint, while the rest of the time its paintings were directed (by asking the sitter for a particular expression) and styled according to one of a selection of mood-appropriate adjectives. And at the end of each painting it self-assessed how well it went in achieving the goal, whether it was successful in putting intent into practice.

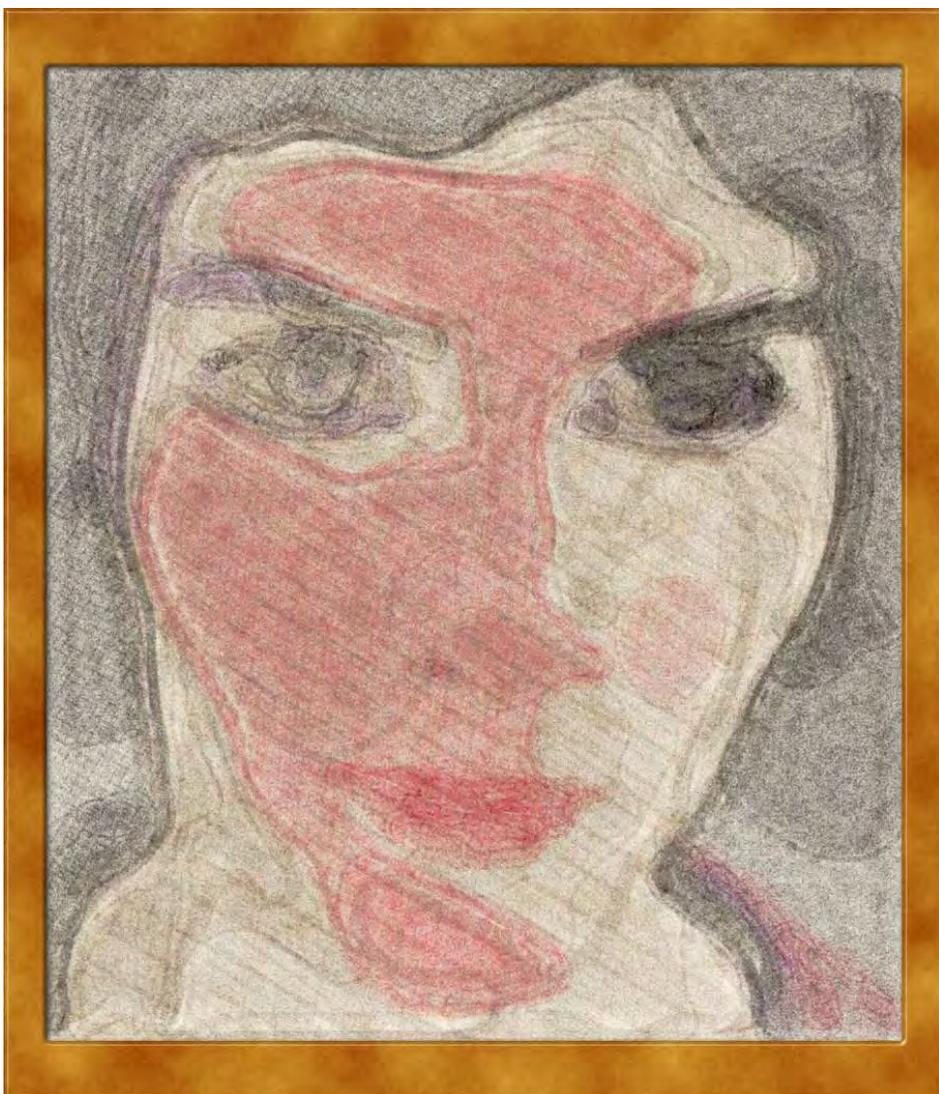
The exhibition got covered on [Wired UK](#) and [Vice](#), among other publications, with Colton noting to Vice reporter Nadja Sayej that events like this are important to illustrate that the software embodies "certain intelligent behaviors associated with creativity, such as skill, appreciation, imagination, learning, intention, and reflection."

The Painting Fool is not human and it never will be, but it can appreciate human emotions and develop crude mechanisms for guiding its own painting process by self-defined moods.

As with AI composers, computer programs that paint their own paintings have developed a nuance in recent years that in a child or apprentice we might call magnificent. But in computers there is still a reluctance to see the beauty.

When I spoke to [David Cope](#) and [Francisco Javier Vico](#) for our creative AI in music composition feature, both noted considerable opposition and suspicion that's been

composition feature, both noted considerable opposition and suspicion that's been mounted at their work, sometimes through anti-computer sentiment, but more often because of a tendency to downplay the achievement as a poor imitation of a professional human artist or composer or as a joke not to be taken seriously.



One of The Painting Fool's earliest projects involved producing [222 portraits](#) of the lead actress from the film Amelie, each with a different emotional value

In AI-authored art there is the same problem, with many dismissing the work of AARON, The Painting Fool, and others of the sort as uncreative or imitative because they believe that computers are incapable of creative intent. But beyond the layered,

unresolved arguments about the nature of creativity, which we will dig into at the end of this series of feature articles, there lies great potential for us to appreciate the artistic collaboration that is possible between humans and computers.

Computers can create subtle, nuanced, original artwork, be it from imaginative interpretations of human instructions or from emergent combinations of mathematical sequences. But the question is: Will we accept them as such? Or will the computer artist always remain an imitator and/or an enabler of the *real*, human artist?

Next week we tackle a different kind of artistry with a deep dive into the issues surrounding computational architecture and design.

206 100 24 0 24

Like Tweet +1 Pin it Share

Share

Advertisement



More from Gizmag



LiquidPiston unveils 40-bhp X2 rotary engine with 75 percent thermal efficiency



Creative AI: Computer composers are changing how music is made



New world record set for farthest hoverboard flight



LiquidPiston unveils X Mini engine

More from around the web




 209 Seconds That Will Make You Question Your Entire Existence
 BuzzFeed Video on The Scene


 Fotokite: A Drone Unlike Any Other
 Dell Power More


 Watch This Robot Build a Bridge
 U.S. Chamber of Commerce Foundation


 10 Favorite New Car Models Coming Out This Year
 Kelley Blue Book

Recommended by 

About the Author

Richard is a freelance writer and journalist based in Melbourne, Australia. He's contributed to Ars Technica, Edge Magazine, Polygon, and many other publications. When not writing or trying to read the entire internet, you'll likely find him dancing, playing games, dabbling in creative stuff, or learning about whatever catches his eye.



[All articles by Richard Moss](#)

Tags

» [Robots](#) » [Artificial Intelligence](#) » [Paint](#) » [Art](#) » [Creative](#) » [Algorithms](#)
 » [Creative AI](#)

Post a Comment

Login with your Gizmag account:

Email:

Password:

[Register here](#)

[Forgotten your password?](#)

Login with Facebook:



Related Articles



[Creative AI: Software writing software and the broader challenges of computational creativity](#)

March 2, 2015



[Creative AI: Computer composers are changing](#)



how music is made

January 26, 2015



Creative AI: Teaching computers to be reporters and storytellers

February 9, 2015



Creative AI: Algorithms and robot craftsmen open new possibilities in architecture

February 23, 2015



Creative Sound Blaster Roar portable Bluetooth speaker

February 27, 2014



Review: Creative Sound Blaster Roar portable Bluetooth speaker

June 9, 2014

Looking for something? Search our articles

Search

Subscribe to Gizmag's email newsletter

Subscribe