

DELAY

DELA

DEL

DE

D

Art in the Age of CAPTCHA

Jimena Canales

'I am not a robot,' wrote Philippe Decrauzat, in notes for a series of paintings displayed on the white walls of an art gallery. The artist concluded his notes by referring to the CAPTCHA, a widely used test to tell computers and humans apart. With clockwork regularity, philosophers, scholars and artists raise their voices to argue that computers cannot make art. Yet every day more and more art involves computing and robotics in a creative and original way. At stake in these debates are dystopian and utopian concerns about mechanisation, artificial intelligence, and the fragility of our basic humanity. Can art save us? Decrauzat's work invites us to look at art in the age of CAPTCHA. Instead of tackling the question of art's exceptionalism theoretically and philosophically, he investigates it practically, visually ... viscerally.

Thou doth protest too much, we might say after hearing the old arguments against computer art and intelligence. For more than half a century already, computers have played key roles in the art world. In 1960, in his article on 'Inventing the Future', the scientist and engineer Dennis Gabor asked, 'Will the machine go a step farther and cut out also the creative artist?' While Gabor hoped that 'machines *will* never replace the creative artist,' he 'in good conscience [...] cannot say that they never *could*.'¹ The editors of *Computers and Automation* magazine had different hopes. To see just how far computers could enter into the art world, they launched one of the first annual computer art contests. The winning image usually graced the cover of their August issue.

Abstract art and Op art had the most to lose — or gain, depending on the perspective — from the new


challenge. The artist Piet Mondrian was among the first whose work was impersonated by a machine programmed specifically to produce work that was practically indistinguishable from the artist's, and possibly better. While working at Bell Telephone Labs, A. Michael Noll, the computer art pioneer, programmed a computer to create *Computer Composition With Lines* (1964), closely mimicking Mondrian's *Composition with Lines* (1917). Bridget Riley's Op art creations soon suffered a similar fate with *Ninety Parallel Sinusoids with Linearly Increasing Period* (1964), created by a computer programmed by Noll to rival her recently completed painting, *Current* (1964). Noll polled volunteers about the provenance of the computer-generated image and its value. Only 28 per cent could identify the computer picture correctly, 72 per cent thought the Mondrian itself was done by a computer, and 59 per cent of the viewers preferred the Mondrian imposter to the original.²

Programmers such as Noll and other artists soon noted that computers excelled at constructing logarithmic spirals, Lissajous curves, stochastic or repetitive patterns (e.g. checkerboards), mathematically defined shapes and surfaces, and oscillating waveforms.³ When judged in terms of their artistic worth, some of the images produced through these new means appeared to imitate, mimic, simulate, and at times even surpass those created by some of the most celebrated artists of the century.


With the creation of ELIZA, one of the first chatbots to successfully impersonate a human (it was designed to simulate a Rogerian psychoanalyst), fears increased that machines capable of taking on human characteristics might be built. Joseph Weizenbaum, who had programmed ELIZA, grew so afraid about how easily people fell for ELIZA's charms that he wrote an entire book, *Computer Power and Human Reason*, to point out the dangers of further developing these machines. 'A relatively simple computer program', he warned, 'could induce powerful delusional thinking in quite normal people.'⁴ He hoped that they would not be given speech recognition abilities in the future, with which they could listen in on our conversations.

By 1984 the MIT professor Sherry Turkle still considered as 'far away' the possibility that 'AI [artificial intelligence] can make robots with superhuman powers.'⁵ Those fears, she claimed, were misplaced, as 'the eventual success of artificial intelligence researchers in actually making machines that duplicate people', mattered little compared to the already present tendency of users to anthropomorphise machines.⁶ 'People tend to perceive a "machine that thinks" as a "machine who thinks";' she wrote.⁷ A decade later, AI had gotten much better at the task of seemingly creating 'duplicate people' in certain, albeit limited, contexts. The childish proclivity to anthropomorphise things, denounced by Weizenbaum, Turkle and others, would soon prove to be a quaint concern.


'[The Internet] is growing faster than O. J. Simpson's legal bills,' noted an article in *Time* magazine in 1994.⁸ Distinguishing humans and computers became much more difficult within the framework of the Internet. On 10 October 1996, President Clinton and Vice President Gore announced their commitment to the Next Generation Internet (NGI) program launched by the US government: 'The Internet will provide a powerful and versatile environment for business, education, culture, and



Front covers of *Computers and Automation*, 17, 19 (August 1968, 1970).



Bridget Riley, *Current*, 1964, synthetic polymer paint on composition board, 148.1 x 149.3 cm, Museum of Modern Art, New York. Philip Johnson Fund, Acc. n.: 576.1964. © 2020. Digital image, The Museum of Modern Art, New York/Scala, Florence



A. Michael Noll, *Ninety Parallel Sinusoids with Linearly Increasing Period*, 1964, photographic print of a computer-generated image, 28 x 21.8 cm, Victoria and Albert Museum, London. © A. Michael Noll

entertainment, [where] sight, sound, and even touch will be integrated through powerful computers, displays, and networks. [...] People will use this environment to work, bank, study, shop, entertain, and visit with each other.⁹

Republicans tried anything they could to defeat these two IT-friendly Democrat incumbents. Collecting data on the popularity of their candidates during the primary elections could be key to getting the keys to the White House. That same month Digital Equipment Corporation (DEC)—later acquired by Compaq, and since merged with Hewlett-Packard—launched a new web service that collected and displayed polling data on the candidates. The problem with their data became

clear early on: the same person could enter an opinion into the poll many times. Worse yet, anyone could potentially create a bot to enter data automatically to skew the polls. As a precaution, the service started displaying 'an American flag in a random position on the screen, and then required the user to click on the flag before entering an opinion.¹⁰ The solution was faulty, since it was 'easy to write a program that recognizes the American flag and simulates a click; therefore, this method does not effectively restrict access by electronic agents.'¹¹

In August 1998, while speaking at the annual summit on 'Cyberspace and the American Dream' in Aspen, Colorado, Ira Magaziner, the Clinton administration's Internet czar, claimed that 'censorship and content control are not only undesirable, but effectively impossible.'¹² If regulating human Internet users appeared to be next to impossible, regulating bots proved even harder. Just as companies such as DEC were working to distinguish single from duplicate pollsters and human ones from automatic ones, others were charging on ahead creating computers and programmes that could pass the Turing Test (i.e., could impersonate humans so well that they could not be outed as such). That year the programmer and artist Peggy Weil noted that 'no matter that no AI has "passed" a Turing Test, however that might be formally debated or defined; there are plenty of "good enough" bots, online and embedded in our personal devices, to have neutralized the original question.'¹³ Weil went on to create one such bot herself, whom she named Mr Mind, taking him online in September 1998 and leaving him there for the next sixteen years. Intriguingly, he was designed to ask users what made *them* human.

Bots soon started playing more active roles on the Internet. That fall, a website launched an online poll asking users to vote for the best computer science department. Students at Carnegie Mellon University found a way to send in votes by using automated programmes. The next day MIT students developed another programme, winning with 21,156 votes, compared with Carnegie Mellon's 21,032, while other schools received less than 1,000.¹⁴ The bot-rigged poll had worked. The result most likely indicated which schools had the best computer science departments, but the poll, which expected answers from humans, had not worked in the way it had been designed.

During these years, the challenges of Internet regulation were mostly looked upon as positive, since the new medium was associated with freedom of speech, free flows of information, democracy and anti-totalitarianism. When, in 2000, President Bill Clinton discussed having China join the World Trade Organization, he downplayed that country's censorship efforts. Its attempts 'to crack down on the Internet', he said, were futile: 'That's sort of like trying to nail Jell-O to the wall.'¹⁵ Censoring humans would soon prove to be much easier than controlling bots.

While early pioneers had tried to make computers more like humans, the challenge for the next millennium would be to find a way to distinguish between the two. One of the first patent applications submitted for this purpose was designed to prevent users from creating multiple accounts automatically for online subscription services. By asking users to identify a 'graphical vali-



A. Michael Noll, *Computer Composition with Lines*, 1964, photographic print of a computer-generated image, 28 x 21.8 cm, Victoria and Albert Museum, London. © A. Michael Noll

dation pattern', the invention 'prevents or makes difficult the automatic creation of multiple users accounts through the use of programming scripts or similar techniques.' Not just any pattern would work, but only those 'designed such that it is made difficult for optical character recognition techniques to identify'.¹⁶ For the first time in history, an inventor had found an effective barrier between the human and non-human Internet users in the new millennium. But the barrier was very porous. Would art change?



Progress in making machines appear as humans was such that the acronym CAPTCHA, for Completely Automated Public Turing Test to Tell Computers and Humans Apart, was coined.¹⁷ The next year Compaq won a patent for 'selectively restricting access to computer systems'. It insisted that 'fairness is particularly important if the provider is running a lottery, or conducting a popularity contest or a poll that allows a user to make multiple entries', pointing out the needs of the advertisement industry (since 'advertising revenue is usually based on the number of times that advertisements are displayed when service requests are made'), and it explained how search engines (such as DEC's Alta Vista) were particularly vulnerable to 'page-boosting'.¹⁸ By then, the challenge was considered a fight between 'real human users' and 'automated agents' 'operating on behalf of users', the latter defined as 'a Software program, or Script generator that can mimic user accesses'. The task was urgent since 'it is well known on the Internet that many agents are intentionally designed to behave in a malicious, destructive, or otherwise annoying "anti-social" manner'.¹⁹

During the 2003 Eurocrypt conference in Warsaw, CAPTCHAs were formally introduced to the larger community of cryptographers. Its creators were more or less optimistic, or both, about finding a particular solution to the problem. 'There is no way to prove that a program cannot pass a test which a human can pass, since there

is a program — the human brain — which passes the test,' they wrote.²⁰ By equating the processing abilities of brains with computer programmes, they signaled a novel predicament of their times. Antagonists were no longer solely defined as computers versus humans, but rather as a series of 'adversaries' some with 'malicious' intentions.²¹ Those who knew how these systems worked stood in a 'win-win situation: either the CAPTCHA is not broken and there is a way to differentiate humans from computers, or the CAPTCHA is broken and a useful AI problem is solved'.²² From then on, the problem became seen as less of a fight between ontologically distinguishable parties than one of 'leveraging' two sides against each other.²³ Yahoo, Hotmail, and PayPal rushed to adopt the new tech, while Microsoft sponsored research.²⁴

The task before those in the CAPTCHA industry grew more complex every second. A growing number of practitioners understood the challenge before them as 'proving humanity', developing 'tests of humanity', 'tests that demonstrate humanity', 'proof of humanity', and a way for a user to 'prove of his or her humanity', and so on.²⁵ The risks of confusing the two — or even worse, favouring one over the other — appeared ever more ominous: 'Is it possible for real users (i.e., humans) to become locked out if the system grows convinced that the erroneous answers provided by a certain algorithm are correct?'²⁶

In 2009 Google acquired a new test designed to improve on traditional CAPTCHA, called reCAPTCHA, developed at Carnegie Mellon. In the first weeks of March 2012, only a few years after Google entered into the competition, the company rolled out a new system. Suddenly, distorted-text tests were gone, and users were asked to identify photographs of 'house numbers pulled from Google Street View'.²⁷ One victim described their surprise at what they saw. 'The one I just saw looks an awful lot like gold address numbers screwed on over vinyl siding,' which led them to grow suspicious of the real intention behind such queries. 'I think it is entirely possible Google is decoding numbers found in images from their Street View cameras. Perhaps attempting to create some new feature within Google maps,' they wrote. After asking if others had similar thoughts and experiences, they concluded: 'I'm curious what's going on now, and just what the ol' G has up its sleeve this time ...' Soon thereafter it became clear to users that valuable data could be put to use without user knowledge or consent. The use of these technologies for fighting spam had evolved to include improving image recognition software. When asked by a probing journalist working for *Techcrunch*, Google confirmed that it was 'currently running an experiment that involves using its reCAPTCHA spam-fighting system to improve data in Google Maps by having users identify things.' Another side to the entire human-bot tracking enterprise started to come into view. While 'the system is designed to cut down on spam and fraud', it 'also helps digitize the text in printed materials, like books and newspapers.' Google admitted using reCAPTCHA results to help digitise content for Google Books and the Google News archives. 'We're currently running an experiment in which characters from Street View images are appearing in CAPTCHAs,' they responded, proceeding to explain: 'We



Top: Figure originally published in David Pletcher, 'United States Patent: Application for System and Method for Validating User Registration for On-Line Services' (n.d.), 24.

Bottom: Figure originally published in Mark D. Lillibridge, Martin Abadi, Krishna Bharat, and Andrei Z. Broder, 'United States Patent: Method for Selectively Restricting Access to Computer Systems' (2001), 5.

puuuu



Figures of CAPTCHA's based on text distortion (top) and image recognition (bottom), originally published in Luis von Ahn, Manuel Blum, Nicholas J. Hopper, and John Langford, 'CAPTCHA: Using Hard AI Problems for Security', in Eli Biham (ed), *Advances in Cryptology — EUROCRYPT 2003* (Berlin: Springer, 2003), 295, 306.

often extract data such as street names and traffic signs from Street View imagery to improve Google Maps with useful information like business addresses and locations. Human users or bots, mostly unknowingly, were helping computers get better at image recognition.

Programmers had long been aware of the possibility of using the data from humans to help computers become more human and more intelligent. For half a decade already, ImageNet, an AI project launched at Stanford University, relied heavily on crowdsourcing to classify and categorise images. By having volunteers classify sets of widely available Internet images, this 'data' could become 'nourishment' for AI. Initially, ImageNet gathered 50,000 willing and able individuals to classify about one billion images and the images were then fed to the computers being trained. Immense progress could be made in AI, explained the project leader, now that 'we have the data to nourish our computer brain.'²⁸

What if these individuals started providing AI nourishment for free — without even knowing? CAPTCHA researchers had landed on the idea of 'using instances of open pattern recognition problems to build CAPTCHAs in order to benefit both online security and pattern recognition research' early on.²⁹ Google was now implementing that project on a massively unprecedented scale. The original project of differentiating humans from bots now took on a secondary importance. Vinay Shet, Google's product manager for reCAPTCHAs, explained in the company's Security Blog the reasons behind the change: 'Advances in artificial intelligence have reduced the gap between human and machine capabilities in deciphering distorted text.'³⁰ The aggressive nature of the fight was as clear as it had even been, with bots who tried to break in considered by their opponents as 'attackers'.³¹

Another breakthrough arrived on 25 October 2013, with the announcement that 'reCAPTCHA just got easier (but only if you're human)'. Psychologists had long known of a clever trick for diagnosing behavioural or mental disorders. When administering an inkblot Rorschach test,

for example, an analyst might as well forget the answer to their question about the inkblot if the subject's behaviour before and after appeared strange. Google tried a similar strategy, picking up data about users' behaviour watched. The CAPTCHA itself was no longer the central part of the test. In fact, it was more of a distraction laid out to see how users reacted to the test itself before and after it was administered. Google started focusing on the 'user's entire engagement with the CAPTCHA — before, during and after they interact with it.' The 'test' had lost all boundaries; CAPTCHAs were just one challenge in a rapidly growing bag of tricks. 'That means that today the distorted letters serve less as a test of humanity and more as a medium of engagement to elicit a broad range of cues that characterize humans and bots,' they explained.³² What is more, by 2014, computers running machine-learning algorithms developed by Google got the distorted text right 99.8 per cent of the time, while humans got a mere 33 per cent.³³ The results of such human/bot discrepancy sent computer engineers into a panic. 'Thus distorted text, on its own, is no longer a dependable test,' they concluded.³⁴ The comments section exploded with complaints from users that they were being blocked. One irate user wrote:

From my very limited dataset of just my own encounters, reCAPTCHAs human/bot pre-distinction is atrociously bad. And from my own horrible experiences with it, assuming that many of my own legitimate users will have to deal with the same, I've had no choice but to ditch reCAPTCHA entirely. I cannot afford to subject my users to this torture and scare them away forever. It's gone too far. Especially since I hear that there are bots out there with automatic CAPTCHA solvers that have a much better hit rate than myself at this point.³⁵

In other contexts, bots were not only passing as regular human users — they were passing as brilliant scientists. That year the problem was so severe that a computer-generated author became 'the world's 21st most highly cited scientist'.³⁶

It was time for an update. In December 2014, CAPTCHA tests were changed to a system based on forced confessions. Users were now asked to click on a box besides the prompt 'I am not a robot'. If used in combination with user behaviour data gathered before and after the answer was given, the new test worked like a charm. Vinay Shet called it 'No CAPTCHA.




Figure of automatically generated face HIP (Human Interactive Proof), originally published in Yong Rui, Zicheng Liu, Shannon Kallin, Gavin Janke, and Cem Paya, 'Characters or Faces: A User Study on Ease of Use for HIPs', in Henry S. Baird and Daniel P. Lopresti (eds), *Human Interactive Proofs* (Berlin: Springer, 2005), 57.

reCAPTCHA', and recounted how his team landed on the idea: 'We figured it would be easier to just directly ask our users whether or not they are robots — so, we did!' The new system was soon adopted by Snapchat, WordPress, Humble Bundle, and several others. Consumer access to shopping websites had been secured, permitting users to engage in 'some last-minute holiday shopping' after 3 December, the day it launched.³⁷

How important are the subtle changes these systems undergo on a regular basis? CAPTCHA designers (starting in 2004) would claim that 'you've probably seen them' without knowing much about them. What else do we see on a regular basis that we do not know much or anything about? What is essential to our lives, yet remains unnoticed?

It takes art as subtle as Decrauzat's to direct our attention to changes in our sensory-motor environment that remain unnoticed. Art in the age of CAPTCHA is a call to a different kind of creative action, one aimed at amplifying the near-silent echo of transformations that interpellate us in coded, abstract, and rarefied ways. Because they have no boundaries, these changes defy categorisation in terms of disciplines. They stump traditional explanations based on longstanding conventions of rational discourse, where they are seen as too minor to be taken seriously. They escape the attention of humanists as much as scientists. Yet it is through them that we can see how some of the most important metaphysical boundaries of our age, such as that between humans and machines, and art and artifice, are created and sustained.

What else? What is taking place right before our eyes that we are not noticing? Answers remain elusive. We have seen, painted, and photographed many sunsets. But only for a much shorter time have we seen, painted, and photographed the repetitive patterns of industrially produced products, or created art based on the toothed wheels and gears animating the machines around us. Only since the invention of the printing press did we become increasingly acquainted with the shape of single letters; only after typewriters left us flooded with loose sheets did we start to read blank pages; only after early computers made us aware of the strange patterns on punch cards did we begin to see them elsewhere. Today, barcodes and QR codes we cannot read reflect prices we cannot afford; pixelated images give us less rather than more when we zoom in; and scrolling text before our tired eyes reflects our present condition just as much as the divisive content on the screen. What have we become when the ring or the vibration of a new notification matters more than the tolling of the bells? When we look up at the sun, it no longer looks the same. Art in the age of CAPTCHA shows us a world much different from what it was when we last saw it.

Notes

1. Dennis Gabor, 'Inventing the Future', *Encounter* (May 1960), 9.
2. Michael A. Noll, 'Human or Machine: A Subjective Comparison of Piet Mondrian's *Composition with Lines* and a Computer-Generated Picture', *The Psychological Record*, 16:1 (January 1966), 1–10.
3. Jasia Reichardt, *Cybernetic Serendipity: The Computer and the Arts* (New York: Praeger, 1969).
4. Joseph Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation* (San Francisco: W. H. Freeman, 1976), 7.
5. Sherry Turkle, *The Second Self: Computers and the Human Spirit*, 20th Anniversary Edition (Cambridge, MA: MIT, 2005), 244.

6. Turkle, 29.
7. Turkle, 29.
8. Philip Elmer-Dewitt, 'Battle for the Soul of the Internet', *Time*, 144: 4 (25 July 1994), 46.
9. 'The Next Generation Internet Initiative: Concept Paper', (July 1997), 1.
10. United States Patent, filed 13 April 1998. Mark D. Lillibridge et al., Method for Selectively Restricting Access to Computer Systems. US Patent 6,195,698 B1, filed 27 February 2001, 2.
11. Lillibridge et al., 2.
12. Ira C. Magaziner, 'Creating a Framework for Global Electronic Commerce', at *The Annual Summit of the Progress and Freedom Foundation* (Aspen, Colorado, 24 August 1998). Since adapted and available online at http://www.pff.org/issuespubs/futureinsights/f16.1_globaleconomiccommerce.html.
13. Peggy Weil, 'The Blurring Test', in Robert W. Gehl and Maria Bakardjieva (eds), *Socialbots and Their Friends* (London: Routledge, 2017), 24.
14. The incident is described in Luis von Ahn et al., 'CAPTCHA: Using Hard AI Problems for Security', in Eli Biham (ed.), *Advances in Cryptology—Eurocrypt 2003: International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4–8, 2003, Proceedings* (Berlin: Springer, 2003).
15. William Jefferson Clinton, 'Clinton's Words on China: Trade Is the Smart Thing', *New York Times* (9 March 2000).
16. David Pletcher, 'United States Patent Application for System and Method for Validating User Registration for on-Line Services' (25 April 2000), 5–6.
17. von Ahn et al. (2003), 309.
18. Lillibridge et al., 1.
19. Lillibridge et al., 1.
20. von Ahn et al. (2003), 295.
21. von Ahn et al. (2003), 296, 307.
22. von Ahn et al. (2003), 294.
23. Daniel P. Lopresti, 'Leveraging the CAPTCHA Problem', in Henry S. Baird and Daniel P. Lopresti (eds), *Human Interactive Proofs: Second International Workshop, HIP 2005, Bethlehem, PA, USA, May 19–20, 2005, Proceedings* (Berlin: Springer, 2005), 108.
24. Luis von Ahn, Manuel Blum, and John Langford, *Communications of the ACM* 47, 2 (February 2004), 60. The Second International Workshop on Human Interactive Proofs (HIP 2005 in Bethlehem, PA) was sponsored by Microsoft and Avaya Labs Research.
25. The quotations about 'proving' come from Amalia Rusu and Venu Govindaraju, 'Visual CAPTCHA with Handwritten Image Analysis', in Baird and Lopresti, 42; those about 'tests', from Tim Converse, 'CAPTCHA Generation as a Web Service', in Baird and Lopresti, 82, 83, 84.
26. Lopresti (2005), 108.
27. Sarah Perez, 'Google Now Using ReCAPTCHA to Decode Street View Addresses', *Techcrunch* (29 March 2012).
28. Fei-Fei Li, 'How We're Teaching Computers to Understand Pictures', at *TED 2015*, min 9:41.
29. Henry S. Baird and Daniel P. Lopresti, 'Preface', in Baird and Lopresti, 5.
30. Vinay Shet, 'ReCAPTCHA Just Got Easier (But Only If You're Human)', *Google Security Blog* (25 October 2013).
31. Shet (2013).
32. Shet (2013).
33. Josh Dzieza, 'Why CAPTCHAs Have Gotten So Difficult: Demonstrating You're Not Robots Is Getting Harder and Harder', *The Verge* (1 February 2019).
34. Vinay Shet, 'Are You a Robot? Introducing "No CAPTCHA reCAPTCHA"', *Google Security Blog* (3 December 2014).
35. Daniel Saner, comment on Shet (2013), published on 2 August 2014 at 7:20 p.m.
36. Richard Van Noorden, 'Publishers Withdraw More Than 120 Gibberish Papers: Conference Proceedings Removed from Subscription Databases after Scientist Reveals That They Were Computer-Generated', *Nature* (24 February 2014).
37. Shet (2014).