


ChatGPT, Culture and Creativity simulacrum and alterity

Alan Dix

Professorial Fellow Cardiff Metropolitan University
and Director of the Computational Foundry, Swansea University

Notes on keynote at Creative AI Research Conference 2023

<https://alandix.com/academic/talks/CAR2023-keynote/>



Abstract

Over the years many of the ‘red lines’ of artificial intelligence have been crossed: challenges that were deemed to require uniquely human understanding. In 1997, chess fell as Deep Blue defeated Kasparov; then, twenty years later, AlphaGo beat Ke Jie, the world’s top Go player. Arguably, game playing can be considered artificial and formal, not representing the rich, nuanced nature of human intelligence embodied in the real world. However, large language models have challenged these assumptions, producing dialogue and texts that appear human – passing the Turing test. Furthermore, the text and poems generated by ChatGPT and images created by DALL-E appear almost creative.

Has the last bastion fallen or is it merely the babbling of ‘stochastic parrots’? Is AI the ultimate charlatan peddling plagiarism or instead the child’s cry that reveals the emperor’s clothes of human creativity to be sham? And what does it mean to be creative anyway?

I will attempt, if not to answer these deep questions, at least to lay down some pointers. We will test the limits of the myth of the individual innate genius with inspiration gifted by the muses; and explore the way creativity is always embodied in culture and technology. Yet, while artists and philosophers debate, the child draws on.

ChatGPT and the end of everything

including
human creativity

In the last few months the intensity of debate about artificial intelligence has become a frenzy driven largely by the phenomenal success of large language models (LLMs) such as ChatGPT. On one side are the doom-mongers warning of human extinction due to AI super-intelligence [HB23]; on the other side those who say ChatGPT is no more than a ‘stochastic parrot’ repeating back human knowledge with no real understanding [BG21].

In one sense the latter are right. Large language models are simply trained to predict the next word in text. However, to do this they are trained on billions of words of human-generated text, more than most of us hear in a lifetime and certainly covering far more topics. The models also include ways to capture longer term associations in the text, not merely simple statistical properties such as that ‘is’ often follows a proper noun. This certainly gives the impression of a certain level of ‘meaning’ or even comprehension¹. Whether this is real ‘understanding’ or merely the impression of understanding is a matter of debate; indeed it is at the heart of philosophical discussions around Searle’s Chinese Room.

There is little disagreement about the potential benefits of AI – we have already seen new medical advances as well as more aids for coding, machine translation and more. There is also widespread concern about the potential practical dangers including deep fakes and discriminatory bias. Indeed as far back as 1992, I predicted dangers of social, gender and ethnic bias in black-box machine learning [Dx92]; problems that are now all too common.

In addition, there is a growing angst that in some way the advent of this AI will undermine human intelligence. In practice, it may replace human jobs [Sw23] in much of what has been considered ‘knowledge work’, not least those of the journalists writing stories about these dangers!

¹ The OpenAI technical report of GPT-4, [OA23], includes detailed discussion of benchmarking against various types of quizzes and tests that would normally be assumed to require human knowledge and understanding.

However, is this also true for creativity? If AI can create paintings in the style of Gauguin or music in the style of Bach, is there any future for human culture?

Rather than answer this question directly, I will instead use the challenge of AI to force critical consideration of the nature of creativity, and then through that the way that AI may or may not fit within it.

technological determinants

Recent advances in AI have been driven by two technological factors: the availability of massive computational power and the availability of vast corpora of data including human communications. Neither of these are about AI in itself, but rather about other aspects of the global digital industries.

Looking first at the latter, *big data*, this is partly due to the impact of the web itself, with large amounts of publicly available information, from the near complete works of past novelists to bulletin boards and individual blog posts. Initially this data was exploited by the large search engines, such as Google, but this also led to massive exercises in further proprietary data collection, such as Google Street View imagery. In addition, social networks, such as Facebook and Twitter created, almost from scratch, their own data sources focused more on inter-personal data.

These are self-reinforcing processes. They build on publicly available data, including the apparently democratic open-data initiatives of many governments, but then augment this with closed data, enabling lucrative business models, notably around advertising. This both creates the financial means and business imperative to collect yet more data, and in the process establishes emergent monopolies in different areas of digital life.

The emergence of big data drives the second technological factor, *big computation*. The need to process these volumes of data led naturally to the development of vast cloud computing with associated data centres. This is a physical infrastructure necessitating new network communications and the siting of data centres where power and cooling are cheapest. However, it is enabled by novel algorithms, such as MapReduce, that allow computation to be distributed over thousands of computers [DG08].

Again, this becomes self-reinforcing as the availability of computation allows new applications, such as autonomous vehicles, and new business models, such as Amazon's cloud computing services.

This background is important, both technically to understand how AI has become so powerful recently, and also because of the way in which it is part of increasingly centralised and self-reinforcing digital industry.




We'll first think about creativity using three Gs:

- **g**od-like?
- **g**enerate and filter
- filling the **g**aps

After doing this we will turn to two key terms of post-modernism: *simulacrum* and *alterity*.



Authors often spend considerable time on the first words of a book, words that set the scene and the tone of everything that follows. Some have become so well known that they have entered popular culture in their own right. Think of “*Christmas won’t be Christmas without any presents*” from Louisa May Alcott’s *Little Women*, or “*It was the best of times, it was the worst of times*” beginning Charles Dickens’ *A Tale of Two Cities*.




Whether you are religious or not, the Bible has shaped Western civilisation, and globally it is the most translated, purchased and read book. Genesis, the first book of the Bible, is also the first book of the Torah, and the opening words are:

In the beginning God created ...

Just let that sink in. The first words of the most widely read book in the world, and the cornerstone of Christian culture declare that the first act of God is to create. Furthermore, the first attribute we learn of God, before omniscience, omnipotence, omnipresence and all the other ‘om’s, is *creativity*.


This may not be the first impression one might get from most US evangelical preachers, nor post-Reformation Protestant theology. But there is no denying the written word. This is utterly radical.



Just as radically, Genesis goes on to say that humankind was created in God’s image, so by extension creativity is an essential, or even the primary, human characteristic. Few would

disagree with this, whether coming from Darwinian, neurological or psychological backgrounds, and certainly no one who spends time with small children would doubt the universal creativity of the four-year old.

Yet often this is not the first image one has of 'creativity', especially in Western contexts. In direct contrast to the universal creativity of every person, we often have the image of the 'creative' person as one set apart, god-like, not in the sense of every person, but as special, lauded.



We see this cult of the individual in many fields, including design. Think of Arthur Fry, the inventor of Post-it notes, realising in a flash of inspiration that a failed glue could be a totally new product. Indeed the idea of 'inventor' conjures up images of the lone genius scribbling, not necessarily in an artist's garret, but alone in a paper-filled office, or bent over bubbling test tubes.

Another classic design story is Philippe Starck's "Juicy Salif" lemon squeezer. The story goes that at lunch in a pizza restaurant on Capri, whilst squeezing lemons over calamari, he has a flash of inspiration, a picture in his mind of something squid-like, which then led to the award-winning design ... albeit not so good for actually squeezing lemons.



Notice that both these stories, Arthur Fry and the Post-it notes, and Philippe Starck's lemon squeezer, are about not only the uniquely gifted individual, but also a 'flash of inspiration', the muse falling from on high. The Greeks envisaged muses in the Olympian pantheon, each

the source of different kinds of musical, poetic and artistic creativity, beyond the potential of mere humanity.


Arguably, this is a very Western-centric viewpoint, especially in the fine arts; in contrast, both Eastern art and Western craft traditions focus more on the development of skills through practice and apprenticeship. Although now, through mass media and social media, if anything the cult of the individual has become global and ubiquitous.



Yet, if we dig a little deeper, even these stories of moments of inspiration for unique individuals become a more complex. They do involve moments of inspiration and people ready for creativity, but as part of a bigger picture.

In the case of the Post-it notes, it was Spencer Silver, working at 3M, who first discovered the glue that didn't stick properly, but peeled back off. As a new glue formulation, it was a failure, and most people would have simply discarded the formulation and moved on to the next experiment. But Silver felt that there was something here; it might be useful even though he didn't know what for. Then, as he visited different parts of 3M, he talked about this failed glue. It was two years later that Arthur Fry was in the audience, remembered a problem he'd had with pieces of paper marking pages in a hymnal, and spotted the potential [3m23, MM23].

There are sparks of inspiration here: Silver's hunch that this was useful and Fry's recognition of a potential use. However, it is not just about the individuals, but about teamwork, listening to other people, connecting threads between work and life. Note also what this says about the internal innovation culture of 3M, where Silver could feel able to spend two years telling people about a failure.

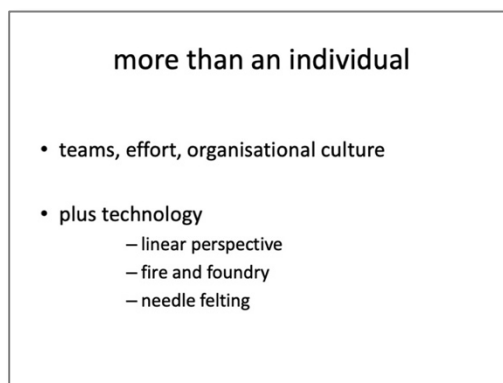


Starck's lemon squeezer is closer to the classic flash of inspiration moment. Indeed the moment is captured in his sauce stained pizza-mat [Ho14]. Starck, who was on holiday at the time, had been commissioned by Alberto Alessi, who recalls the mat arriving by post:

"I received a napkin from Starck; on it, among some incomprehensible marks – tomato sauce, in all likelihood – there were some sketches. Sketches of squid. They started on the left and as they worked their way over to the right, they took on the unmistakable shape of what was to become the Juicy Salif lemon squeezer." [Be17]

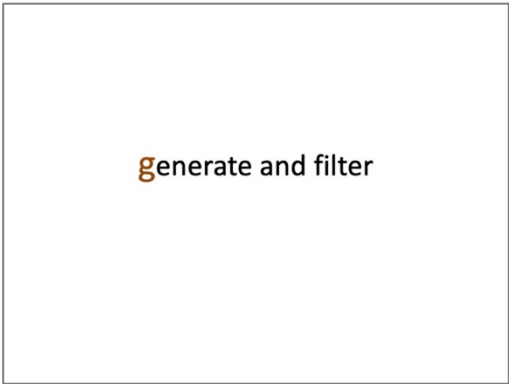
However, looking more closely at the mat [Ho14, Be17] the order is less clear. In the centre there are images of a standard table-top glass lemon squeezer (conical top with integral glass dish), with some sketches of the dish part standing on three small legs. As these are closest to a standard design, one presumes they began the process. Beside these are images of the conical squeezing part, with a spiral rather than straight fluting and with two strands of fluting unfurling to provide two legs and a third conical support point – this seems closer to a hand-held squeezer, albeit supported. It seems that this latter image then transmuted into the final form, with any parallels to the squid appearing late in the process, if anything.

That is, rather than a single instant of inspiration, this records an inspired, but step-wise, evolution from the standard shape to something very different. Furthermore, this would have stayed as a doodle if not for the collaboration with Alessi.




In short, even in the classic stories of individual inspiration, we see the importance of personal effort and persistence, *and* the way that this is richly embedded in and enabled by teamwork and organisational culture.


Note also that technology, while not obviously present in these stories, often plays a crucial part in shaping or enabling creativity. Think of the way the discovery of linear perspective transformed the nature of art for hundreds of years; or the emergence of photography, and later filmography, as media for forms of creativity that were previously not merely impossible but near unimaginable. Similarly fire and smelting has enabled metallic arts, needle-felted animals, and also the casting of the “Juicy Salif” lemon squeezer in aluminium. Of course, industrial processes for paper and chemical production underpin the Post-it note. More recently



generate and filter



The complete collected poems of William Wordsworth run to over 1100 pages, but while he is one of the great poets in the English language, only a small number of works, such as “I Wandered Lonely as a Cloud” or “The Prelude”, are well known beyond literary circles.



Similarly, the great composers, such as Bach, Mozart and Beethoven, produced hundreds of pieces of which only a small number are regularly played as part of the canon.




You may have heard of Searle's Chinese Room [Se80]; here is Alan's Creativity Room.

Imagine there is a door cut into a mountainside. In the centre of the door is a letterbox and below it a hatch into which, each day, the attendants pass crate upon crate of bananas and a very large steak.



For long periods nothing happens, but then, occasionally, out of the letterbox, a folded piece of paper falls, and on it is a poem or aphorism of intense meaning or great beauty.


Those outside imagine that the door leads to a small hermit's cell, within which a constantly hungry reclusive poet breeds fruit flies as a hobby whilst composing these works of genius.



The truth is dramatically different. Instead, the door leads to a vast cavern, within which lie ... row upon row upon row of typing chimpanzees.

In aisles between them a single watcher walks, or maybe a small team who share the steak each day.

They watch the pages produced and discarded by the chimps. Some are meaningless, some start tantalisingly well, but then after a few more random keypresses dissolve into the absurd "*to be or nurgle posit*". But occasionally, they spot something that has promise: "*to be or to become ...*"; and if it excites, intrigues or moves them, they fold it up and post it out through the door.



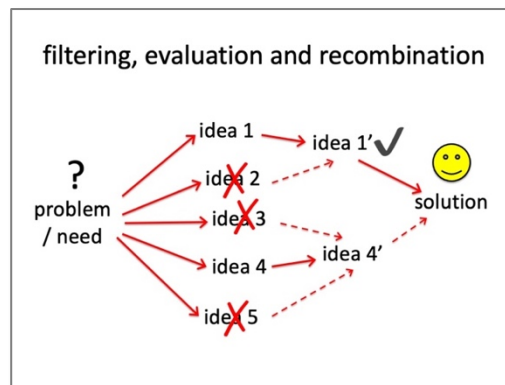
So where is the creativity? It is merely a room full of chimps randomly typing.

Is there none? Those outside would disagree.

Is it in the minds of those who walk the aisles? But surely they are merely critics, not creatives?

Maybe it is the room as a whole: a creative system rather than a creative individual?

If you know Searle's Chinese room, you can see the parallels between creativity and intelligence here.



One (partial) model of creativity is *generate and filter*. From some form of problem or need, whether practical or aesthetic, a variety of ideas are generated. Of these some are discarded immediately, and others retained, perhaps modified and then the modified ideas further filtered down until one or more survive as the preferred concept, solution or work of art.

This is certainly how many in the fields of graphic or industrial design work: initially creating many sketch ideas and then, either individually, with others in their studio, or in consultation with the client, honing them down to a single concept for production.

Think now of Wordsworth or the played canon of the great composers [We99]. There is clearly creativity in the original production. Furthermore, Wordsworth will have discarded poetic ideas and certainly worked on poems before publishing them, that is he already engaged in evolution and filtering.

However, open the complete works and choose a poem. It may, if you are lucky, speak direct to your heart. More likely it will be accomplished, polished, well crafted, but not great, nor stunning. A book of the selected poems of Wordsworth or of any other poet, will comprise maybe a quarter of the full corpus, and even then it is an even smaller number that are frequently read. Similarly with musical canon, traditional or popular: how many songs of the Beatles or Queen do you know?

A critical part of the creative output of 'Wordsworth', alongside the poet himself, is the community of critics, editors and general public who have filtered his works down to those we see as representing the man today. 'Wordsworth' is not merely the man who lived in 19th century Westmorland, but the moniker of a social and cultural process that led to the works we read today.

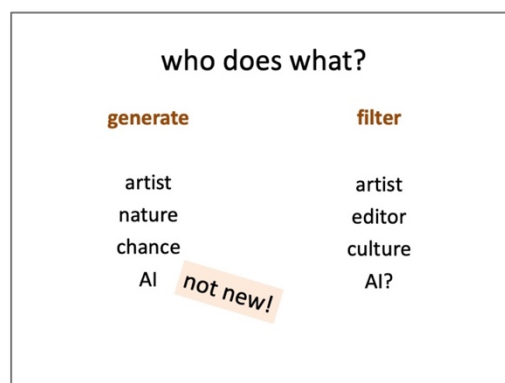


Different arts differ in the balance between generation and filtering. This is most obvious in photography, where a large part of the skill, accomplishment and creativity of the photographer is in the framing of a scene, or the choice of a final image from many taken during a shoot. This is not to say that there isn't a process that leads to this point, choosing a spot to wait for the right light on a landscape, or positioning lights to achieve the right highlights for a portrait. However, photography acknowledges and accepts that selection and choice are a key aspect of the creative work.

In fact this creativity through framing predates the photograph. In Rydal, where Wordsworth lived for much of his life, there is a small waterfall in the grounds of Rydal Hall. In 1688, at the edge of the pool below the waterfall, a tiny building, the *Grot*, was constructed, with a window that frames the perfect snapshot of this ever-changing view.

Three hundred and thirty years later, if, after waking that morning in 1988, Tracy Emin had simply taken a photograph of her bed and displayed that, it would probably have created little if any comment – there might have been discussion about its taste or merit, but not about whether fundamentally it was *art*. Yet, if you think of the translocation of her bed from bedchamber to gallery as framing, a capturing of the moment, then it has an aesthetic in precisely the same tradition as Daguerre or Leibovitz.

Oddly if Tracy Emin had resorted to sculpture or clay and *made* her bed that morning, there would likewise have been little discussion of whether it was art. Of course, if she had made her bed that morning we would never have known about it at all!



As we have seen, the roles of generation and filtering can be found in many places.

- **generate** – the *artist's* imagination in spontaneous ideas; *nature* in landscape or microcosm; *chance* as used in Dada and even as early as Mallarmé in the late 19th century [Ed21, Ma97].
- **filter** – *artists* themselves as in Wordsworth editing his own poems; *editors* both curating, but also often modifying works; *culture* and *critics* deciding which elements of a body of work to preserve and reproduce.

If we look at computation and AI, there is a long-standing literature of computer-generated poetry and art; in almost all cases heavily selected and curated. That is the computer is the *generator*, with a human acting as *filter*. At its simplest this may be semi-random recombination of words and phrases. This is in some ways 'better' than the randomly typing chimps, as the randomness is usually constrained at least to real vocabulary and often to some form of linguistic structure, either through use of formal grammars or probabilistic techniques such as Markov processes based on co-occurrence statistics.

Large language models, such as ChatGPT, are in this sense simply the latest incarnation of probabilistic techniques, using vast language corpora so that the statistical structures capture rich linguistic patterns. Certainly the use of these to 'create' stories and articles usually involves a combination of human selection and prompts to edge an incipient work into something that at least satisfies a brief or maybe says something new.

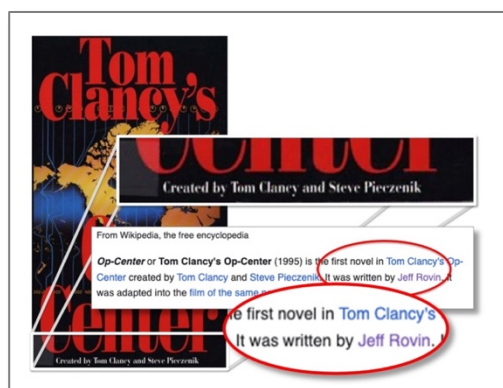
In some ways this reflects the way that a physical material such as a piece of wood or stone, has a role in shaping and setting the limits on what can be produced, but is also managed and manipulated by the artist. That is the physical materials have an element of agency, and similarly the artistic use of AI can be seen both in terms of a *new material* as well as a new actor.

However, one of the other tools of big AI, especially in image generation, is *adversarial learning*, where one neural network learns to critique or evaluate the outputs of another. For example, one neural network might try to generate realistic images of landscapes, and the other tries to distinguish generated images from images of real landscapes. Over time the generator learns to be better at creating realistic images, and it is this which is usually released as part of a product. However, meanwhile, the adversary AI has learnt to become better at distinguishing the real from the artificial, and indeed may itself be deployed, for example as part of social media filtering to detect bots.

A critique neural network can equally be trained to distinguish good and less good works of art, or novels, and indeed there are already services producing AI-generated analysis of your writing, based on exactly this principle (e.g. authors.ai). This critique will be trained within the established canon of literature or art, and hence not recognise unusual or transformative art, but then, that has also been the story of human critics!

In short, creativity is often the result of an ecology of roles, and AI could play, and indeed has played, different roles within this process.

filling the gaps



I read thrillers. The best are usually by sole authors, but I also read the mass-produced literature by authors such as Tom Clancy. If you look at the cover of Tom Clancy's "Op-Center", it is indeed his name you see writ large, but in smaller writing, at the bottom of the cover it says, "Created by Tom Clancy and Steve Pieczenik", and if you look inside, you see it is actually written by Jeff Rovin.

If we think of the 'author' as the one who pens the words we read, we may initially feel this is a bit of a cheat. However, Tom Clancy undoubtedly will have had the main role in creating the world of the Op-Center series of novels, including the characters and settings. It is likely that he will have worked with Steve Pieczenik in fleshing out the details of the plot of this specific novel, and then Jeff Rovin 'filled in the gaps', the actual words, but based on a rich corpus of Clancy's literature to guide the style and genre of writing.



Tom Clancy is standing in a long artistic tradition dating back at least to the Renaissance [Wa14], where the artist is not working alone, but leading a team. It is likely (but contested) that Michelangelo painted the entire ceiling of the Sistine chapel personally, but he certainly had assistants to help him. When it came to the sculptural work that he preferred, his assistants did more than assist in practicalities, they engaged in sculpting the stone. A painting by Collaert in the 15th century depicts an artist's workshop or *Bottega*; the artist is at the centre, but surrounded by apprentices, not merely grinding paints but actually painting.

The principal artist such as Michelangelo would determine the overall artistic direction of a series of sculptures, perhaps work personally on the central figures, approve and guide the work of assistants, and carve the fine details that turn the accomplished to the masterly.



Artists Workshop. Jan Collaert I (ca. 1530–1581)
 (photo: The Elisha Whittelsey Collection, The Elisha Whittelsey Fund, 1949.
<https://www.metmuseum.org/art/collection/search/659725>)



ChatGPT?

Many are using ChatGPT in exactly this way: starting with bullet points to establish the main flow and structure, then using prompts to determine the genre of writing, but letting ChatGPT act as a ghost writer, filling in the gaps ... just like Michelangelo's assistants, or the relationship between Tom Clancy and his team of co-writers. ChatGPT will not put Tom Clancy out of a job, but the days of the jobbing writer may be numbered.



creativity?
outcome of a
socio-technical-cultural
process
with
multiple interlocking roles

As a final word on this discussion of the three 'G's of creativity, it is clear that, rather than being the inspired work of individuals, creativity is more often the outcome of a rich socio-technical-cultural process with multiple interlocking roles. It is already the case and likely to continue to be the case computer tools and AI can function in some of these roles. In any project it seems, at present, unlikely that every role will be automated, but increasingly likely that at least some will.



As well as a general sense of impending doom, artistic fears of AI and critique of AI have been driven by two, almost mutually contradictory, arguments.


On the one hand are worries by artists that their intellectual property is being infringed by AI [PS22].

Some of this is about focused copying, such as the way Dall-E or other diffusion models can create works in the style of a specific artist [RD22]. Whilst early examples were of long-dead artists, such as renderings of modern topics in the style of Rembrandt or van Gogh, it was clear that this could be applied to living artists, and indeed already has been with tracks becoming viral in the style of still-performing singers and groups [Sa23].

Even more focused is the use of digital copies of an actor's appearance, voice and movement patterns to drive a fully artificial digital twin instead of the real actor. As I write, Hollywood is paralysed by strikes about this very issue that include many big-name actors [Re23,St23]. Who owns your image, your mannerisms? For many this is a philosophical or privacy issue, for actors it is a matter of basic livelihood.

For some time before the actors joined the strike, screen writers had already been taking industrial action. Their worry, and the worry of many artists, is more diffuse. It is less that their specific works are used to create a film score or picture in their own style, but that the amalgam of countless artists' works, including their own, is used to create new, entirely artificial works that displace them both artistically and financially.

On the other side, those critiquing the possibility of true AI creativity point to exactly the same aspects of AI. It is never more than derivative, adapting, recombining or building on the works of others with no true originality or novelty, no more creative than a photocopier.

-
- 
- History of Art
 - intertextuality
 - genres
 - Eastern artistic traditions
 - apprenticeship

However, in some ways we can say “thus it has always been”.

There can be direct copying or emulation, including art forgery or tribute bands. However, most influences are more indirect.

I have a huge glossy volume entitled the “A History of Art” [Ja77] from a heady day exchanging a book token won as a Mathematics Essay Prize while at university. It starts with the earliest prehistoric art and runs all the way to the modern day (at least to the 1970s when it was written), including some aspects beyond purely Western art. I’m not sure if I ever read it cover to cover, but, in the long-lost leisure of student life, I certainly did start at the beginning. The introduction talks about novelty as part of what defines art. It shows “Le Déjeuner sur l’Herbe”, the well-known 19C painting by Manet, in which two (fully clothed) men and one (near naked) woman idle over a picnic in the park. This is regarded as a great work of art, but then the book shows an earlier Renaissance painting with a group of people in near identical poses. As one begins to question the novelty of Manet’s painting, given this, the author then shows an image of a Greek sculptural frieze, which appears to be the original inspiration ... unless there is a yet earlier unknown work.

The message was clear: *all* art is derivative, but it is still art.

Indeed, the term *intertextuality* celebrates the way that each generation of literature references the literature of the past, sometimes explicitly, sometimes more indirectly.

Some artists deliberately try to separate themselves from prior influences in order to establish their own artistic identity; whilst others, notably the pop artists of the 1960s, celebrate the cacophony of past and current creative voices. Of course, the former are the products of the artistic milieu that has surrounded them since childhood; and the latter are now finding themselves facing IP lawsuits, not unlike ChatGPT!

We all live within genres and artistic traditions, both general and specific. From apprentices in a Renaissance *bottega* to students at art school, we soak in the influence of past works and current teachers. Should JK Rowling be able to claim royalties from every subsequent story of a schoolchild wizard, or Mary Shelley’s descendants from *Twilight*? Indeed, are either Shelley or Rowling original, given the traditions of English novel writing from Fielding to Austen, let alone folklore and fairy tales that probably started round the first Neolithic campfire.

Given that every artist is not singular, but part of this diffuse stream of artistic styles and influences, are ChatGPT and other large language models any different? That is, should we consider them (a) any more accountable and (b) any less creative because of this?

You will have your own views on these questions!

For the second question, I certainly feel that the ubiquity of artistic recapitulation means that the dependence of AI on past corpora does not, in itself, fundamentally undermine the potential for AI to be 'creative'.

For the former question, I feel more hawkish. The way in which big industry controls and exploits technology changes the fundamental power dynamics of creative culture.

There is a difference between the cultural influences on a sole artist that have arisen during their personal life course, and the systematic mining of all culture. It changes the scale of attribution and analysis.

If another group commercially exploits a Beatles song without permission, the courts will rightly rule in the Beatles' favour (estates and individuals) even if the song can be attributed to no single member of the Beatles and the sound the new group produces is not identical frequency-by-frequency with the original. If a large-language model reproduces and replaces the overall work of a field of artistic endeavour, is this not the same, especially if done for commercial reasons?



Of course, art is more than imitation. Invention may be 90% perspiration, but the other 10% is crucial.

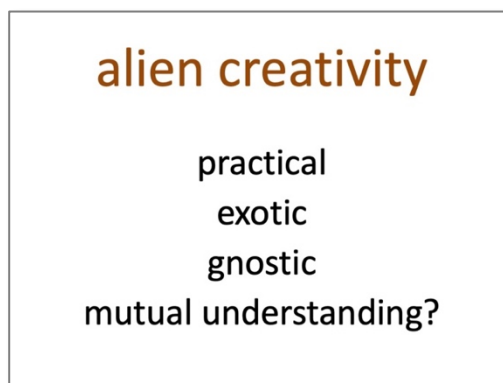


I recall a talk in the early days of computer-aided music, in this case not in the sense of the computer composing, but rather computer interfaces that made it easier for novices to get into playing and indeed improvising. One interface used a 2D layout of squares, rather like a large chessboard. I don't recall the exact layout; I think one direction moved in tones, the other in minor thirds. Crucially, the layout meant that when you moved in simple geometric patterns, the notes that were produced were always melodic; rather like the way that the pentatonic scale means that you can strike windchimes in any way and still create pleasing sounds.

On the screen, as well-known tunes were played, you could see the squares light up in simple patterns. But then the presenter played the Beatles song 'Michelle'. It started off following the 'rules', but then as the haunting '*ma belle*' played, the pattern went off-piste, an unexpected turn, momentarily breaking the simple geometry, but then returning.

Moments of genius are often like this, the small variations, infringements, or deviations that are not utterly other, but still shock and tantalise. Sometimes these are unrecognised or rejected by the critical establishment of the time, especially when they are more major or more sustained, such as the early Impressionists or Cubism.

Of course, the very nature of AI models based on large corpora is that they do recapitulate and recombine. One can see how they could compose new easy listening, accomplished, but uninspired, but not these turns that transform the mundane into the magical.



Sometimes the initial violation may be accidental, a missed note, or glue that doesn't stick, but recognised as aesthetically or practically of value. This is the microcosm of generate and filter, and hard to distinguish from directed alterity, if indeed it is different.

It is easy to present AI with vast numbers of alternatives, accidents or missed notes, but can we imagine AI being able to detect which are unusual, but still good?

If we think at a practical level of engineering or science, this is not only possible, but has already happened, with new drug discoveries, or for that matter the discovery of novel strategies in Go. Crucially, while AlphaGo, the first deep neural network to defeat a Go grandmaster, was initially trained on human games, AlphaZero, its successor, was only given the rules of the game – what makes a good game – and learnt from scratch [Si17].

As noted already when discussing generate and filter, adversarial techniques create networks that are designed to critique, to decide what is good. If these are taught using good and less good literature, good and less good paintings, good and less good music, it is clear they could be (and are being) used to help novices distinguish and critique their own work, and by extension notice which mistakes or accidental quirks make something still good, from the perspective of the already accepted canon.

What is less clear, but still possible, is whether the generalisations learnt during this process might enable the AI to grasp deeper notions of human aesthetics, and hence distinguish or create things that are truly different from the norm and yet also pleasing, challenging or interesting to a human. Certainly Go players have found some of AlphaZero's strategies meaningful and (apparently) insightful.

One could also imagine adversarial AI, maybe primed on human works or maybe not, generating a completely different 'aesthetic': things that it considers good and novel, but in very different ways to human aesthetics. While this sounds far-fetched, it may already be beginning to happen! One of the potential limitations of large-language models is due to the way they are trained on web corpora, yet already large tracts of this web material, such as news articles, are themselves being written by tools such as ChatGPT. As the models effectively learn from each other, they will create both alternative realities (as false news memes do within the human web) and alternative styles of presentation ... an artificial aesthetic. This aesthetic will not deviate entirely from the human. Web crawlers will attempt to distinguish human-written material in order to focus training on this, and certainly the underlying need to generate advertising revenue will ensure that computer-generated ideas of what is 'good' writing connect at least to the psychological triggers that encourage click-throughs ... an extreme extrapolation of human-generated media.


Maybe, however, computer learning from computer-generated material will lead to whole esoteric sub-genres that are only appreciated by computers – a form of *alien creativity* with its own aesthetic.

In human culture this has often been the case, both between generations and between countries or parts of the world. Indeed, for many people the movements of fine art or more experimental music are entirely gnostic, meaningful only to those 'in the know'.

However, we can come to appreciate works from completely different aesthetic traditions. Indeed, when re-listening to the original Beatles recording of Michelle, I found the rhythm

patterns a little odd and jarring even, until they sank in ... perhaps I was relearning an appreciation of the aesthetics of late 1960s pop or maybe the Eastern melodic influences on the Beatles' work.

One day soon you might find yourself in conversation with a ChatGPT-successor music critique, which is trying to explain the meaning of a new musical AI-genre that is alien to human ears.



So, does AI spell the end of human creativity, or is it merely a more sophisticated level of cut-and-paste?

If we are focused on the god-like view of momentary inspiration of rare, gifted individuals, then the bar is high; especially if we think about those critical differences, the alterity that lifts the accomplished to the amazing. However, we have seen that in fact creativity is usually a much more complex process including multiple people, technologies and social situations. In such a picture sophisticated digital tools are already being used, from spell checkers to CGI, and AI will add to these, forming part of a creative ecosystem.

Furthermore, it does seem likely that for some of the roles in this, AI *could* and thus, in the economic steam-roller of the media industry, *will* replace humans. This is already happening and will accelerate, particularly the competent, but not inspired, work of the traditional artist's apprentice or modern ghost writer.

For the B-movie or factory-line book the whole process may be automated, producing works that are not just as good as those of humans, but even better, because they are able to draw on so much inspiration of the truly excellent. In the coding domain Albrecht Schmidt, in a keynote at EICS 2023, described experiments with AI coding, finding that for most tasks AI performed at the level of a mid-range undergraduate student [Sc23].

For the more exceptional talents, the picture was different. Schmidt found that when allowed to use AI as a tool to help their own coding, the mid-range students were not significantly helped, possibly because they lacked the critical skills to assess the adequacy of computer-written code. However, he also found that the work of the best students was lifted, presumably because the AI was able to do the 'grunt work' of coding. In the 1970s, coding was more like the Renaissance artist's *bottega*, with a systems analyst guiding the process and a small team of programmers filling in the details. This disappeared partly

because programming tools improved, displacing the lower-level programmer. In some ways AI may be doing the same thing at a higher level.

In the artistic world only those who have broken through, already recognised as successful, can afford (or at least their publishers or promoters can afford) to have the assistance of teams, such as Tom Clancy's co-writers and ghost writers. Imagine if every aspiring novelist had such a team, if every painter had a *bottega* or every film student the virtual resources of a Hollywood blockbuster?

As with many things digital, one can imagine the democratisation potential, but as is usually the case, the danger is that the imbalance of computational resources means that instead this will entrench the already wealthy and powerful. The massively data- and computationally intensive nature of frontier AI means that for this, more than many digital technologies, money and power win.


This certainly creates a technical challenge for those in universities and a policy challenge for governments, to ensure that the fruits of AI's ability to enhance human creative potential are not reserved for the few.



Although we started with a critique of the god-like view of the lone genius, much of the discussion has edged towards such a view, from Wordsworth to the Beatles. Margaret Boden talks about H-creativity vs. P-creativity [Bo90]. H-creativity is at the 'historic' or cultural level; the thing that is novel across humanity as a whole, recognised by many as being of special value. P-creativity is about the personal creativity of each and every person; the 'God's image' form of creativity that is universal. The writers of the Torah began with the creation of the entire universe, but still appreciated the creativity of those who cast and carved the Ark of the Covenant.

Just because cars travel further than people, doesn't mean we do not run. Just because we cannot aspire to be John Lennon, doesn't mean we do not sing. Across the world people continue to play chess and Go.

It may be that some day, even some day very soon, AI will create works of art, music or literature that will outshine even the greatest human. However, that will not make a single child's painting less precious.



References

- [3m23] 3M (2023). About Post-it® Brand. Accessed 14/8/2023.
https://www.3m.co.uk/3M/en_GB/post-it-notes/contact-us/about-us/
- [BG21] Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. (2021, March). On the dangers of stochastic parrots: Can language models be too big? 🦜. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. pp. 610-623.
- [Be17] Carlotta Berta (2017). Juicy Salif | Design Lemon Squeezer. *unprogetto blog*, 24 February 2017,. Accessed 26th July 2023. <https://www.unprogetto.com/en/juicy-salif-design-lemon-squeezer/>
- [Bo90] Boden, M (1990). *The Creative Mind*. Weidenfeld and Nicolson, London, UK.
- [DG08] Dean, J. and Ghemawat, S. (2008). MapReduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1), 107-113.
- [Dx92] A. Dix (1992). Human issues in the use of pattern recognition techniques. In *Neural Networks and Pattern Recognition in Human Computer Interaction* Eds. R. Beale and J. Finlay. Ellis Horwood. 429-451. <https://alandix.com/academic/papers/neuro92/>
- [Ed21] Eden, M. (2021). Randomness and chance as writing methodology: Is there a case for Dada in the academic essay?. *Journal of Writing in Creative Practice*, 14(1), 79-106.
- [HB23] Geoffrey Hinton, Yoshua Bengio, Demis Hassabis, Sam Altman and over 600 signatories (2023). *Statement on AI Risk*. Center for AI Safety. Accessed 4/7/2023.
<https://www.safe.ai/statement-on-ai-risk>
- [Ho14] Hobson, Benedict (2014). "Philippe Starck's Juicy Salif was "the most controversial lemon squeezer of the century"". *Dezeen*, 9 July 2014. Accessed 26th July 2023.
<https://www.dezeen.com/2014/07/09/movie-alberto-alessi-juicy-salif-controversial-lemon-squeezer-philippe-starck/>
- [Ja77] H. W. Janson (1977). *A History of Art*. Thames & Hudson
- [Ma97] Stéphane Mallarmé (1897). One Toss of the Dice Never Will Abolish Chance. *Cosmopolis*, May 1897.
- [MM23] MoMA (2023). *Post-it® Note*. MoMALearning. Museum of Modern Art, NY, USA. Accessed 14/8/2023/. https://www.moma.org/learn/moma_learning/art-fry-and-spencer-silver-post-it-note-1977/

- [OA23] OpenAI (2023). GPT-4 Technical Report, *arXiv* :2303.08774. <https://arxiv.org/abs/2303.08774>
- [PS22] Dex Parra and Scott R. Stroud (2022). The Ethics of AI Art: The Case Of Dall-E 2 [Case Study]. *Media Ethics*, Fall 2022, 34(1). <https://www.mediaethicsmagazine.com/index.php/browse-back-issues/219-fall-2022-vol-34-no-1/3999403-an-extension-of-the-artist-dall-e-2-and-the-ethical-challenges-of-ai-art>
- [RD22] Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). *Hierarchical text-conditional image generation with clip latents*. arXiv preprint arXiv:2204.06125.
- [Re23] Regan Morris (2023). The mega-strike that could take down Hollywood. *BBC News*, 12 July 2023. <https://www.bbc.co.uk/news/world-us-canada-66140336>
- [Sa23] Chris Saunders (2023). This scarily realistic AI-generated Drake and The Weeknd track is going viral. *Hunger*, 17 April 2023. <https://www.hungertv.com/editorial/this-scarily-realistic-ai-generated-drake-and-the-weeknd-track-is-going-viral>
- [Sc23] Albrecht Schmidt (2023). *Speeding Up the Engineering of Interactive Systems with Generative AI*. Keynote at ACM Engineering Interactive Computer Systems (EICS 2023).
- [Se80] Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417-424.
- [Si17] Silver, D., et al.: Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354–359 (2017)
- [St23] Emily Stewart (2023). Nobody wins in the Hollywood strikes, including billion-dollar studios. *Vox*, Aug 23, 2023. <https://www.vox.com/money/23843034/wga-sag-aftra-strike-amtp-studios-netflix-warner-bros>
- [Sw23] Mark Sweney (2023). BT to axe up to 55,000 jobs by 2030 as it pushes into AI. *The Guardian*. Thu 18 May 2023. <https://www.theguardian.com/business/2023/may/18/bt-cut-jobs-telecoms-group-workforce>
- [Wa14] Ian Wallace (2014). The Evolution of the Artist's Studio, From Renaissance Bottega to Assembly Line. *Artspace*, 101. June 11, 2014. https://www.artspace.com/magazine/art_101/art_market/the-evolution-of-the-artists-studio-52374
- [We99] Weber, W. (1999). The history of musical canon. *Rethinking Music*, chapter 15, pp.336–355.