

AI and the social construction of creativity

Dr Paul Atkinson 

Monash University, Australia

Dr Richie Barker 

Deakin University, Australia

Convergence: The International Journal of Research into New Media Technologies 2023, Vol. 29(4) 1054–1069
© The Author(s) 2023



Article reuse guidelines:

sagepub.com/journals-permissions
DOI: 10.1177/13548565231187730
journals.sagepub.com/home/con



Abstract

Artificial Intelligence (AI) encroaches on new terrains of human activity by dint of its efficacy and an expanding ability to autonomously incorporate information from many disciplines and sources. In this paper, we focus specifically on how AI affects the communicative practices associated with creativity. AI has the capacity to reshape discipline and taste communities by providing new content that competes with human production and by mediating between human activity and information sources. To frame these issues, we turn to the influential systems model of creativity devised by Mihaly Csikszentmihalyi (1996), which Csikszentmihalyi and Daniel Gruner (2018) recently extended to incorporate AI, redubbing it *Creativity 4.0*. The model assesses how AI affects the social structure of creative practice without overly accentuating the similarity between humans and AI, or questioning whether computational devices will replace creative jobs. The paper examines Gruner and Csikszentmihalyi's revised systems model, arguing that it does not sufficiently take into account the variety of ways that AI can be incorporated into creative practice. Prompted by a theoretical reflection on the nature of the model and the emerging features of AI, we propose a new version of the model that highlights how embedded AIs play a key role in filtering and gatekeeping, as well as the importance of generative systems in informing creative practice. We propose that any discussion of AI and the future of creative practice should look at *where* and *how* AI supported technologies are used. We examine how AI can reduce and shape the qualitative diversity of sources of inspiration drawn into the creative process, with the associated technological biases, as well as provide an emergent platform for the development of novel ideas.

Keywords

AI, systems model, creativity, field, domain, filtering, Csikszentmihalyi

Corresponding author:

Dr Paul Atkinson, Media, Film and Journalism, Monash University, Caulfield.
Email: paul.atkinson@monash.edu

Introduction

Artificial Intelligence (AI) is increasingly integrated into a range of social and technical systems, from education, news, the economy, and medicine to the arts. Rather than emerging as fully independent systems that emulate human cognition, most AI features in networked applications and devices as autonomous agents that can adapt to an environment and respond independently to new information (Franklin, 2014: 28), as well as standalone programs that perform specific tasks, from analysing medical scans to transforming photographs into drawings. These technologies are encroaching on new terrains of human activity due to their increasing ability to autonomously incorporate information from many disciplines and sources. In this paper, we focus specifically on how the widespread use of AI affects the social foundations of creative practice. Creativity does not operate in a vacuum, reliant only on the ingenuity of individuals, for it emerges within discipline and taste communities and derives from processes of information gathering and technical practice. Although questions concerning whether AI can effectively replace or simulate organic intelligence have long been a dominant theme, we look instead at the potential of AI to reshape the communicative practices associated with creativity. All AI, including those labelled 'generative AI' and AI embedded¹ in everyday platforms, has the capacity to reshape communicative contexts by providing new content that competes with human production, by changing the conditions for the evaluation of creative output, and by mediating between human activity and information sources.

To frame these issues, we turn to the influential systems model of creativity devised by Mihaly Csikszentmihalyi (1996), which Daniel Gruner and Csikszentmihalyi (2018) have extended to incorporate AI, redubbing it *Creativity 4.0*. Most importantly, the systems model of creativity focuses almost entirely on the social and communicative structures informing creative practice rather than psychological factors underpinning the production of novel ideas. The model provides a basis for assessing how AI affects the social structure of creative practice without overly accentuating the similarity between humans and AI, or strictly focusing on whether computational devices will replace creative jobs. The latter is somewhat of a distraction, for AI is already shaping communication and undertaking discrete, iterative tasks irrespective of whether or not it replaces individual workers and professions (Susskind and Susskind, 2015: 212). The paper examines Gruner and Csikszentmihalyi's revised systems model, arguing that it does not sufficiently take into account the variety of ways that AI can be incorporated into creative practice. Prompted by a theoretical reflection on the nature of the model and the emerging features of AI, we propose a new version of the model that highlights how embedded AIs play a key role in filtering and gatekeeping, as well as the importance of generative systems in informing creative practice. This should assist both practitioners and researchers in assessing how AI is used and incorporated into creative practice. Principally, we argue that the systems model provides a platform for understanding how AI and creativity operate in a communicative context characterised by a surplus of information. Of particular interest is the role of database search tools that filter information for use in creative practice as well as the role of generative AI systems in producing more content. With regard to the latter, we do not assess whether generative AI produces outputs that are better than or replace human-produced content, rather we look at how it changes the way that creative practitioners operate within the wider social structure of creative practice, including alternating from content creator to content selector and/or modifier of AI-generated material. Drawing upon the systems model, we propose that any discussion of AI and the future of creative practice should look at *where* and *how* AI-supported technologies are used. We examine how AI can reduce and shape the qualitative diversity of sources of inspiration drawn into the creative process, with the associated technological biases, as well as provide an emergent platform for the development of novel ideas.

The systems model

One can marvel at the capacity of AI to perform and succeed in completing tasks comparable to the highest levels of human intelligence and creativity, from playing chess to drawing and painting. Although these accomplishments are truly remarkable, attending to them in isolation can obscure the fact that creative activity, whether organic or synthetic, is a social act, depending on others in a network of communication. When Deep Blue defeated Gary Kasparov in the late nineties, it signalled a change in our conception of chess as an indicator of human intelligence, a clearly social construct, in addition to demonstrating the ability to beat human players at chess. Each notable act of creative activity changes the way we see and evaluate a particular discipline of sphere or activity. To consider these issues, the paper foregrounds those aspects of creativity most directly invested in a social system and involve the negotiation of meaning between individuals and cultural assemblages. Rather than outlining the psychological and intrinsic aspects of creativity, Mihayli Csikszentmihalyi's original systems model of creativity highlights the importance of artistic, professional, and academic disciplines in shaping individual practice, as well as the social conditions for regulating and assessing creativity. In this model, creativity involves the production and development of novel objects and ideas, yet novelty is not an intrinsic property that can be broken down into a set of practices or procedures for it is socially determined. He argues that there is no absolute, objective determination of creativity because it is decided upon by participants in a discipline and is therefore relative (1999: 314). This social approach avoids some of the ontological questions about creativity. Yasemin Erden (2010: 360) states that the term creativity has developed to describe human acts such as intuition and contemplation, and can only be applied to AI on the proviso that it describes a 'new language-game' with different features. Likewise, Margaret Boden (2010: 29) argues that creativity cannot be readily defined as creative outputs are necessarily '*new, surprising, and valuable*'. In addition, she argues that the loose definitions that have emerged, many of which are associated with peculiarly human attributes from consciousness to intentionality, and cannot easily be used to assess whether computers are creative. In line with these arguments, we do not seek to understand whether AI production is inherently creative or to directly compare human creative production with algorithmic production. The aim rather is to assess how AI production and filtering in creative fields can change the relationship between participants and objects of knowledge.

To give greater specificity to his argument that creativity is socially founded, Csikszentmihalyi proposes a tripartite model in which the creative individual is placed at the juncture of a social system comprising two other main components: 'a cultural, or symbolic, aspect which is here called the domain; and a social aspect called the field' (1999: 314).² The model clearly distinguishes between the individual's socially moderated creative action, and the cultural sign system or discourse that serves as the source and framework of creative production. The cultural domain refers to a culturally accepted body of knowledge or area of interest, from painting to physics, with known 'symbolic rules and procedures' against which the creative object or act operates (1996: 28). It is a form of discourse that comprises signs, beliefs and practices that inform as well as constrain creative production and includes, inter alia, university, artistic and scientific disciplines. For the act to be recognised as creative, it must make a substantial contribution to the domain and this requires the participation of another group of people, the field. The field comprises cultural arbitrators, experts within a particular domain (gallery directors, critics, review panels, historians, etc.), who assess and promote the creative acts and objects, and decide which are sufficiently innovative. The members of the field serve as 'gatekeepers' of the domain (Csikszentmihalyi 1996: 28). In short:

Creativity is any act, idea, or product that changes an existing domain, or that transforms an existing domain into a new one. And the definition of a creative person is: someone whose thoughts or actions change a domain, or establish a new domain. It is important to remember, however, that a domain cannot be changed without the explicit or implicit consent of a field responsible for it. (Csikszentmihalyi 1996: 28)

In this model, all activity is socially determined or evaluated. The creative individual works in relation to a specific domain or across domains, which is usually associated with types of work. Most people working in the domain do not make a definite creative contribution even if they end up changing the domain (Csikszentmihalyi, 1996: 37). The model does not provide a detailed framework for understanding how creative products are accepted by the field or even how individual creators develop novel ideas by drawing upon the domain; nevertheless, it is valuable in foregrounding the social and communicative system in which creative activity is embedded.

The systems model effectively spatialises and externalises creative activity by placing it in a communicative network; consequently, Csikszentmihalyi states ‘the first question I ask of creativity is not *what* is it but *where* is it?’ (1996: 27). In this approach, creativity is recognised as socio-culturally shaped and no longer hidden within the black box of the human mind. For our purposes, this spatialisation can help foreground AI’s role in the distribution of knowledge, mediating the relationship between creator, field and domain in addition to creating novel content. This raises a number of questions: Should AI be located within the domain or between the creator and domain? Does it have a role in modifying the field? Or does it even act as part of the field? Gruner and Csikszentmihalyi have sought to situate AI in a revised version of systems model, see Figure 1, they refer to as ‘Creativity 4.0’ (2018: 451). This type of ‘fourth wave creativity research’ builds on the first wave of research with its emphasis on Big-C creative individuals,³ a second wave that addresses the psychological processes inherent in creative activity, and a third wave that foregrounds social systems and culture including, of course, Csikszentmihalyi’s own systems model (2018: 450–451). So where does AI operate within this revamping of the systems model? For Gruner and Csikszentmihalyi, AI’s main function is generative, creating new content with the potential to replace the creative individual (2018: 451). Generative applications can often grab headlines when producing works that appear to rival human creative production, especially when conjoined with humanoid bodies and associated with the common fear that the robots are coming.⁴

Generative AI is now widely used to refer to particular types of programs that produce texts of various types using training data, and the most well-known are ChatGPT and DALL-E. However, the term ‘generative’ has a lineage that extends beyond the current crop of machine learning applications, and has been applied broadly to refer to the production of art using rules and procedures – from methods for musical composition to literary movements such as Oulipo (*Ouvroir de littérature potentielle*) (Boden and Edmonds, 2009:21). Due to the emphasis on procedure, the term generative has been adopted by AI researchers and programmers and is mostly associated with art generated by computers (Boden and Edmonds, 2009: 23). In using the term ‘generative’ throughout this paper, we maintain this wider meaning to refer to the procedural production of creative objects using various types of AI, not only GANs and associated forms of machine learning. In another example, story generators have, since the 1970s, been able to generate strongly plot driven stories by applying combinatorial rules and genre conventions, and most create stories that emulate existing types of narratives in line with the programming aim of mimicking human mental activity (Gervas, 2009: 58). Likewise in the 1990s and early 2000s, programs have generated humour, particularly wordplays and puns using specific rules supplemented by access to broad dictionaries (Ritchie, 2009: 73). More recently, intelligent agents can generate stories with only

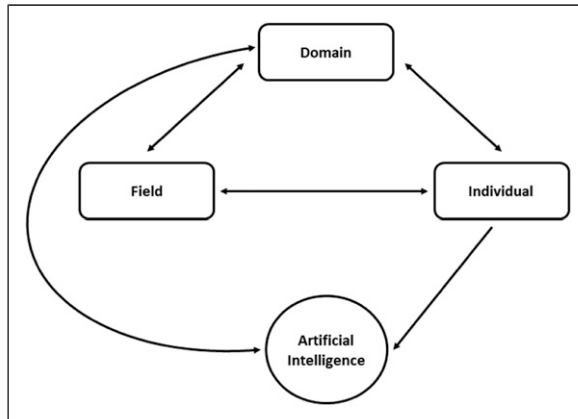


Figure 1. The creativity 4.0 model (Gruner and Csikszentmihalyi, 2018: 452) highlights the role of AI in communicating with the domain without showing how it also assists in filtering information for the creative individual.

minimal input, such as OpenAI's *Talk to Transformer*, in which the user adds a minimal amount of text and the AI continues the story by creatively manipulating material from a range of sources (Thorne, 2020: 812). Musical composition also readily aligns with generative art because the language (tones, rhythmic patterns, harmonic structures) is easily encoded and in most cases follows clear principles and rules. For example, as early as 1983, David Cope trained his program EMI (Experiments in Musical Intelligence) to compose works in the style of well-known composers. In an experiment in which an audience was asked to listen to three pieces, one by Bach and two in the style of Bach (one composed by a musicologist and the other by the program), the audiences judged the EMI composition as more likely to be original rather than the one written by the musicologist (Sawyer, 2012: 143). Numerous other programs paint, draw, and design in ways that draw comparison with human outputs. One critique is that many programs do not sufficiently take into account how audiences might receive a joke, story, or piece of music. They operate procedurally rather than socially, unlike a human creator who is always embedded within a social system. For example, jokes produced by algorithms are not particularly funny because lack contextual awareness and timing, and generated stories do not have the depth of reference that characterises human creativity (Gervas, 2009: 61). Irrespective of quality, programs can produce artistic and literary works with only minimal human input after the initial period of programming or training, and are becoming better at self-assessment due to machine learning with in-built systems of appraisal, training and regulation.

AI-generated creative output has developed outside the arts in areas such as building design, fashion, and advertising, with machine learning and generative AI providing a significant increase in the number of applications. In fashion, AI research is mainly being conducted in promotion and sales, for example, in transforming two-dimensional images into images of people wearing the garments, or for transforming sketches into detailed images of the final project (Luce, 2019: 128–129). Nonetheless, AI researchers in fashion design are exploring the use of generative models, demonstrated by Amazon's claim in 2017 that it could use GANs to create design images and concepts (Luce, 2019: 125–126). In advertising, a 2018 advertisement for car maker Lexus attracted significant attention because its script outline was developed by the AI system, IBM Watson, based on the analysis of previous award-winning ads and car ads, even though the actual video was

directed by a human. Although this has not been replicated, it marks a significant shift in the way we think about creativity in commercial contexts, demonstrating that creative content development does not necessarily require human actors. Academics are also extolling the value of AI in replacing advertising creatives. Demetrios Vakratas and Xin Wang proposed a model for a ‘creative advertising system’ or CAS that not only produces creative advertising but evaluates it (2021: 39). The system would be able to draw upon data from the domain, for example, through ‘video mining’, to evaluate the quality of the output, as well as stimulate creative practice by locating ‘white spaces’ – those areas of creativity that have not been sufficiently investigated (2021: 44). Such a fully integrated program has not yet been devised, but it shows the ambition of those working in AI to develop comprehensive systems that operate across all facets of the systems model, drawing upon the domain and supplanting the creative individual as well as the field. A challenge for these anticipated AI models is they are not truly aware of social context – what it means to live in a social system with all the complexity associated with interests, desires, and beliefs. Advertisements are communicative processes in which the advertising creative shares a social context with the addressee of the message, as well as an awareness of what it means to desire and consume goods (Scott, 1994: 468). Novelty in advertising is not solely a matter of creative production, for it also arises from an attentiveness to the way people live and how they respond to variations in social trends.

The updated version of the systems model focuses on the potential of AI to perform tasks usually attributed to a creative individual, although Gruner and Csikszentmihalyi note that this does not constitute a significant threat to human creativity because programmers are still the creative driving force and because current AI requires rule-based algorithms (2018: 451). They argue that most applications depend on programmers to determine what is inputted and derive their outputs from models, which means that the applications are not intrinsically creative (458) – an argument that does not readily apply to machine learning. From this perspective, AI is a form of secondary activity that relies on creative programming to produce creative outputs and does not have the requisite ‘moral and emotional reasoning’ to replace human creativity (Gruner and Csikszentmihalyi, 2018: 460). Gruner and Csikszentmihalyi ultimately adopt the replication argument – in which AI’s creativity is proven by its capacity to produce creative objects on a par with those produced by humans – in their revision of the systems model. This foregrounding of the ontology of creativity is a little surprising considering that the systems model deemphasises the psychological dynamics of creative production, and, like the Turing test, mainly addresses how novel variations are shaped by the field or how the creative individual draws knowledge and inspiration the domain. Due to the problems with comparing human cognition with algorithmic procedures, Boden suggests that it would be better asking ‘*what aesthetically interesting results can computers generate, and how?*’ or ‘*Just what might lead someone to suggest that a particular computer system is creative*’ (2010: 10). In the context of the systems model, a pertinent question is: *how does the introduction of computer-generated content affect the relationship between the creative individual, field, and the domain rather than can AI replace human activity in areas of creative production?*

AI-generated creative output should be positioned as a standalone element in the systems model because creativity is necessarily relational, a product of the interaction between the creative individual, domain, and field (see Figure 2); a logic that should also apply to any discussion of AI’s creative potential. This brings us back to the ‘where’ of creativity. A key site of creativity is the individual’s relationship to the domain, and the process of reworking and responding to existing ideas and content to create something new. Gruner and Csikszentmihalyi do recognise AI’s increasing capacity to incorporate new symbolic content from the domain, and therefore ‘make new combinations that are perhaps increasingly independent from the original information provided by

individual programmers', and yet argue that AI mainly operates within a closed system (2018: 451). In contrast, human creativity has a much more flexible and open relationship to the domain and can more easily shift from one domain to another (Gruner and Csikszentmihalyi, 2018: 459). Most current AI works within very specific domains and is usually associated with discrete tasks from playing Go to designing fashion. In contrast, a creative individual, even when working on a discrete task, is cognitively engaged with a broader social and cultural world and this feeds into associative thinking. This social embeddedness can serve as the foundation for what Boden (2009: 25) designates as 'transformational creativity', or the capacity to develop truly novel ideas outside what is already demarcated in a specific domain. This could refer to minor alterations of a set of principles – not quite following the rules of perspective in painting, looking to subvert audience expectations in advertising – or it may involve fundamental changes to a discipline, often operating on a metalevel – the full rejection of linearity by a painter or developing a new genre of advertisement. However, this type of creativity should not be highlighted to the exclusion of others. Most forms of creativity operate within genres, disciplines, or existing rule-based practices, and do not necessarily engage with multiple domains – for example, writing counterpoint in music or designing new kitchenware products based on existing style principles, both of which align well with the procedural functioning of AI. When talking about social systems, the internal logic of the creative activity is less important than how creative individuals assimilate information from the domain. Increasingly, we are confronted by a surfeit of information that cannot be managed within the finite limits of human attention,⁵ unlike AI which has a much greater capacity to assimilate data from increasingly open domains, largely facilitated by internet-based data trawling. Consequently, theorists cannot refer to the flexibility of human thinking without recognising the significant scope of current AI to address extremely large data sets. In short, flexibility of thought and association does not automatically increase in line with an expanding domain, for an individual cannot attend to all the information that is available. Information and creative outputs have to be filtered in order to be brought within the purview of human attention.

Filtering: From field to domain

The systems model offers some scope for understanding how the creative individual filters what comes from the domain, and how the field as a gatekeeper determines what outputs are judged to be creative. Nonetheless, the Creativity 4.0 model does not sufficiently address how these relationships are altered in response to AI, for it is assumed that AI will only disrupt the role of the creative individual through usurping the creative. Alternatively, we argue that the creative individual can operate like the field in response to content generated by AI, and that AI can also mediate the relationship between the creative individual and the domain (Barker and Atkinson, 2019). This is a significant shift that has been worked into our revision of the model which situates the influence of both embedded AI and Generative AI systems (see Figure 2). The field comprises experts in a particular discipline area and serves as a gatekeeper or regulatory mechanism for the recognition of creative acts, and also as a means of establishing the conditions for more creative output. In the Creativity 4.0 model, Csikszentmihalyi and Gruner exclude AI from any discussion of the field because they argue that it comprises human actors and is only 'indirectly' affected by AI (2018: 451). This seems reasonable because the field is mainly concerned with values, normative practices and matters of taste, primarily human issues with no absolute criteria for judging what is creative or not. From this perspective, it would at first appear that AI could not operate like a field, for how can a machine make human judgements? The question is not whether AI can properly make such judgements, again we have to move away from directly comparing human and synthetic modes of

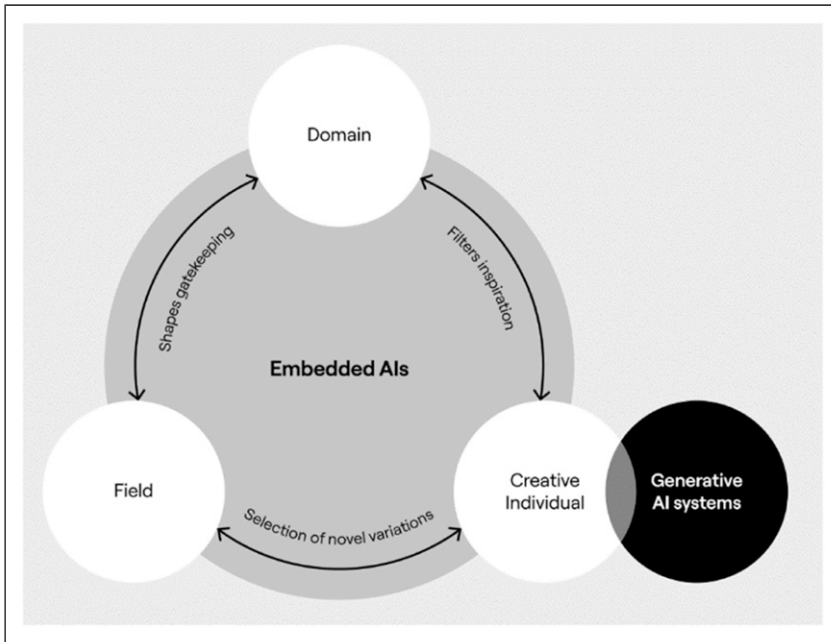


Figure 2. This adaptation of the systems model shows how embedded AI shapes the relationship between participants without producing novel content. Filtering through such devices as popularity algorithms can regulate content in creative production, as well as serve as a proxy field. The creative individual also collaborates with generative AI and in some cases selects from generated content to also act as a type of field.

production, but rather whether AI activity can directly, or even indirectly, affect what is nominated creative in the domain. One example is the use of algorithms in embedded AI to select and determine successful content in video streaming and download services, that is, YouTube and Amazon, or in music services such as Spotify. For example, Netflix algorithms condition the choices of viewers through a recommender system, drawing upon data from particular users to collectives (Pajkovic, 2022: 215–216). These algorithms make suggestions about what people should listen to or watch, thus demarcating what is popular in the domain, and, in doing so, indirectly influencing the commissioning of new works, something that is usually limited to the field. Of course, the determination of success differs from the assessment of novelty, yet both are forms of gatekeeping that shape the domain by regulating what people want to see and hear. Indeed, the field will only attend to creative works and ideas if they are within their purview, for example, production companies must lobby so that their films are brought to the attention of members of the Academy in the Academy Awards.

AI can affect creative practice by mediating the relationship between the creative individual and the domain – those symbolic materials and systems germane to a particular discipline area. Digital communication has greatly contributed to the amount of available information in the domain, and, consequently, creative individuals increasingly rely on communicative technologies to both facilitate and limit access. These digital communication technologies, Csikszentmihalyi argues, can promote novelty because they communicate and store information with greater facility than older technologies, which includes broader access to information from a wide range of other cultures and participants, creating a much more differentiated and varied domain (1999: 317–318). In addition,

computational technologies allow for accurate and long-term storage of information: ‘The more permanent and accurate the storage, the easier it is to assimilate past knowledge, and hence to be well positioned for the next step in innovation’ (1999: 318). The capacity of computers to easily replicate and store data ensures that if someone truly wants to find something, it is highly likely they can. For example, an artist should be able to access digital copies of works, they are not restricted to what is currently on display in local galleries, museums, libraries, and so on, all of which are much more directly regulated by the field.

The overabundance of available information has to be balanced against the definite limits of human cognition, and, hence, the development of filtering technologies. This was signalled by Csikszentmihalyi (1996: 42) in his early account of the systems model of creativity, where he argues that it is impossible for any individual to fully come to terms with surfeit of information in any one domain due to a ‘scarcity of attention’. Physicists cannot read all the works pertaining to their discipline, nor can artists and critics properly come to terms with all the paintings produced in any 1 year. So how do we explain what survives and what disappears? Csikszentmihalyi frames the filtering process in terms of Richard Dawkins memes (1996: 41), although he focuses on the field’s role in selecting memes and therefore does not strongly adhere to the biological connotations of the theory.⁶ The memes that survive are those that truly change the way a culture thinks or acts: ‘[t]o be creative, a variation has to be adapted to its social environment, and it has to be capable of being passed on through time’ (Csikszentmihalyi, 1999: 316).⁷ Survival is determined by the field, which includes the formation of a canon that stabilises the domain and serves as a touchstone for new creative work (1996: 42). From this perspective, the field can be as broad a population, for example, those determining whether an advertisement or song is creative and part of a canon, or as narrow as a few key experts working within an area of high technical expertise.

AI currently has a more prominent role in managing information than developing creative content, particularly in mediating an individual’s access to the domain. Based on this logic, embedded AI can operate as a proxy or even supplementary field and thus shape that is accepted into a cultural domain, a relationship we have incorporated into our revision of the Creativity 4.0 model (Figure 2). Understanding how AI affects creativity should not be limited to the production of novel content. This mediation partly determines which ideas will remain relevant and prevalent: social media feeds regulate access to content based on approximations of user interest; and popularity and personalisation algorithms in search engines fashion what is popular and therefore visible. If search engines are used to stimulate the creative process for the creative individual, the preponderance of popularity algorithms could reduce access to qualitatively diverse sources of inspiration or occasion serendipity. Early forms of search such as AltaVista were more likely to produce more varied results because they used less factors to determine ‘relevance’, and mainly relied on the preponderance of a search term. Most contemporary search mechanisms ‘consider how often the site is linked to by others and in what ways, and enlist natural language processing techniques to better ‘understand’ both the query and the resources that the algorithm might return in response’ (Gillespie, 2014: 175). Despite improvements in relevance, these search mechanisms do not definitely know why the user has sought particular information (Gillespie, 2014: 175), for example, an artist looking for unusual designs to inspire creation differs from someone looking for a design object to purchase. In the latter, popularity is a definite criterion where in the former, eccentricity should be a factor. The algorithms in search engines do not favour eccentricity, for example, Google ‘yokes utility to quality of search experience’ and in doing so maintains the status quo (Hillis et al., 2013: 54). Creativity is usually associated with forms of divergence and, to return to Boden, new and surprising ideas, whereas search algorithms require utility to deal with large amounts of information quickly and efficiently.

As such, the role of AI in filtering how the creative individual accesses the domain also emerges as a key communicative element identified in our reworking the systems model (Figure 2).

Many AI programs are designed to improve the speed and efficiency of workflow, and most of this involves extracting information from databases or speeding up routine tasks. In the field of marketing and advertising, Adobe Sensei purports to ‘collapse the time between marketing ideation and execution’ through making it easier to access existing image databases, tag content, and recognise features that may be pertinent to a current task. The program is integrated with Adobe Advertising Cloud Search to better target audiences as well as improve ‘bid optimisation, forecasting and media budget allocation’ (Adgully). Although the program claims to help with content creation, the emphasis is really on what is already known or knowable, such as the assessment of a market and searching databases, rather than the unpredictability and openness of truly creative work. Efficiency is usually associated with the increase in speed of existing processes and tasks, whereas creativity places greater emphasis on the degree to which an idea, process, or product diverges from what has been done before. Algorithmically driven applications overly valorise efficiency, in particular the efficient use of time – they seek to circumvent any activity that is deemed to be idle or does not fit the logic of the system – and this becomes the basis for self-governance (Wajcam, 2019: 332–333). From the perspective of creative practice, should programs operating in the interest of efficiency remove those elements that are nominally deemed superfluous, including idle time, which from a human perspective involves the broadening of knowledge and an increase in entropy? Creativity necessarily involves both novelty and task appropriateness, but if the latter is overemphasised, there is very little creativity. AI that is embedded into everyday applications used in the creative process extend the cognitive range of the users,⁸ and yet still operate within a means-ends logic in which all activity is organised around clearly articulable goals, which effectively overlook how the individual is situated within a meaningful social system in preference for technological notions of efficiency (Simpson, 1995: 44), and in the case of search engines, relevance. This does not mean that we should not adopt such technologies, rather we have to work out the degree to which the speeds and rhythms of AI actually complement the speeds and rhythms of human activity and foster creativity. Gruner and Csikszentmihalyi’s Creativity 4.0 model does not address the role of AI that is integrated into everyday platforms or applications as it only considers standalone generative systems. As a result, it overlooks the intersection between the creative individual and information sources. In the study of AI in creative process, understanding where and when is as important as the generation of content.

From algorithmic filtering to human selection

Even when addressing AI’s capacity to generate new content, the analysis should not be limited to the quality of that content, for the prevalence of the outputs will change the domain and the role of the field in filtering that domain. With the increase in processing power and the development of self-learning technologies (machine learning), generative AI can produce content at a much greater speed and volume than human creatives. Programs might not demonstrate the same level of planning and foresight as a creative individual, but due to the sheer number of outputs will definitely produce novel content. The capacity to produce numerous works for little cost and in a short time can directly affect how artists work: how they incorporate AI-generated material into their own works, and, with a surfeit of content, how they select what is useful or valuable. With generative AI systems, the focus shifts from the selection and filtering of human-produced content by AI, to the human selection of AI produced content. This is already occurring in small pockets of creative activity, and it is likely to increase. For example, the program Amper is currently used by song

writers to assist in creating backing music. Amper generates a number of backing tracks based on the selection of parameters ranging from instrumentation to genre, and the artist then chooses the backing music that best suits the song they are writing. Taryn Southern used the program for her album *I AM AI*, and she ‘sometimes rejects as many as 30 versions of each song generated by Amper from her parameters; once Amper creates something she likes the sound of, she exports it to GarageBand, arranges what the program has come up with and adds lyrics’ (Love, 2018: 55–56). In addition to significantly reducing the costs of producing an album, the technology assists the creative individual in realising a work and differs from earlier programs, such as David Cope’s EMI, because the artist selects, reuses, and curates AI-produced material. In advertising and marketing, AI could significantly reduce costs by automatically creating new ads to suit different regions without refilming (Campbell et al., 2022: 28). In this role, generative AI produces new content in a collaborative relationship with the creative individual, who is the arbiter of what will be used as the final work.

Programs are increasingly coming into the market that automate intellectual and creative practices, from large language models to image generators. What distinguishes many of the programs is their ability to internally self-assess the quality of the output without human intervention. One of the most talked about examples are text-to-image generators, which use strings of words as prompts to generate novel images. Some of these programs, such as *Pixray* (<https://replicate.com/pixray/text2image>), use GANs (Generative Adversarial Networks) to produce new content. In such models, the generative component of the program produces samples that are checked in a discriminator module against real samples, and the system works by trying to produce generative content that can be mistaken for human-produced content (Nelson 2020). Google’s *Imagen* and DALL·E establish a relationship between text and image by drawing on existing data sets designed for image recognition software, such as the COCO visual data set, or automatically by scouring the web for images using the alt-text data (a feature of most websites) to associate a word with a type of image. This data is modified through a ‘diffusion’ model, where images used in training the AI lose their definition through the addition of Gaussian noise and then new images are developed by subsequently reducing the noise and adding detail (Ho, 2021). Using such processes, Google’s *Imagen* can produce images recognised by AI visual detection systems and can be confused for actual photographs by human observers and rival the illustrations produced by graphic artists, although the main role at this stage is co-creative as opposed to operating without any human intervention, a relationship identified in the overlap between the creative individual and Generative AI systems in Figure 2. For example, AI was used to generate the 11 June issue of *The Economist*, as well as design the cover for *Cosmopolitan* magazine using DALL·E based on the following prompt, ‘wide-angle shot from below of a female astronaut with an athletic feminine body walking with swagger toward camera on Mars in an infinite universe, synthwave digital art’ (Liu 2022). The chosen image fitted all aspects of the brief including the use of the low and wide angle shot – the surface of Mars appears to curve around the boot of the astronaut, which is the most prominent feature of the image – so it is not just a matter of producing a recognisable object, but also rendering the object in a way that adapts to the viewer’s perspective.

These systems work independently of human activity except for the design, the text prompt, and the setting up the machine learning training system. They produce content that aligns with their internal training, and, as has been noted, this can be underpinned by a range of technical biases, including stylistic biases. The more that these types of programs fill the space of design production, the more a type of aggregate and statistical reasoning will shape what is considered to be good design within a domain. Zachary Kaiser (2019: 177) argues that principles of design are increasingly reduced to mathematical and symbolic forms that can be read by computers and algorithms,

including such aesthetically difficult categories as beauty. This process of quantifying and encoding these aspects of art and design is not a mere description of how they operate, with all the variability of aesthetic categories, but of prescribing them. Machine Learning will ‘flatten’ design by seeking the ‘beautiful’ through a ‘ruleset’ without truly understanding the historical ‘situatedness’ of aesthetics (Kaiser, 2019: 177). This is somewhat true of the current text-to-image generators. When typing the prompt, ‘Van Gogh style fifties dress’ into *craiyon*, all the images are very similar and use Van Gogh’s *Starry Night*, the most popular work, for the dresses’ patterns. The dresses are very conservative in cut and correspond to a middle-brow notion of taste. The reproduction of the style occurs without reference to the conditions under which Van Gogh worked or any aesthetic attention, but there again, the merchandising of Van Gogh’s work in most gallery shops is equally indifferent to art history.

Considering these programs in terms of the broader social construction of creativity, novel variations often arise through co-creation between individuals and this should now include the collaboration between creative practitioners and AI – a factor that is not addressed in the current creativity 4.0 model. Putting stylistic bias to one side, these programs can still produce incredibly novel content due to the variability of human input in the text prompts, and the process of translation across two modalities, text and image. Texts can combine elements in ways that are not usually apposite to image production – which was the basis of many of the creative ideas in art movements such as Dada and Surrealism. The *Imagen* website (<https://imagen.research.google/>) includes a range of sample images where the arbitrariness of the text prompt underpins the novelty, for example, ‘An alien octopus floats through a portal reading a newspaper’ or ‘A giant cobra snake on a farm. The snake is made out of corn’. DALL-E is able to generate images not predicted by the programmers because it uses a very large data set rather than human evaluators (Ramesh et al. 2021: 8821). The program can also perform ‘zero-shot image-to-image translations’ where the image is reworked into different forms such as photographs transformed into drawings, combined into a single image, or translated into a new style or palette (Ramesh et al. 2021: 8828). In all cases, the constant process of translation – text to image or from one style to another – creates novelty. The main criterion is that the images are legible to human audiences. Notably, DALL-E creates a set of images for each text prompt, and if the prompt is reinserted, it will produce another set, and it is at that point they are subject to selection, where the creative individual operates like the field. The program, of course, struggles with abstract forms (ideas, concepts) because they are less likely to have readily accessible visual types. Using the freely available cut-down version of DALL-E, *craiyon*, (<https://www.craiyon.com/>), and inserting the aforementioned prompt, ‘an old American car flying through the sky in the form of a spaceship’, produces an array of nine spaceship images, most of which are in a cigar shape against a neutral blue sky. The preponderance of the cigar shape shows a stylistic bias, with dependence on common images of UFOs and 50s car styles. However, there is some variation in the type of cigar shape and peculiar features, for example, oddly placed taillights. The images are all legible and even if they do not quite meet what the user wants, they could serve as ideas for designers to use in their own work. In a quite prescient article, Boden (1995: 21) claimed that AI will play a role in mapping out new ideas by demonstrating what is possible within a particular style or concept. It will be an adjunct to the process of mental mapping by creating concrete examples of these ‘conceptual spaces’ and by correcting human experimentations with design. Text-to-image generators could actually perform a similar role by experimenting with visual concepts, and constantly translating images into new types of design. Whether this can be properly called ‘creative’ does not really matter, for irrespective of how the images are generated, they will be incorporated into the broader social system of creative practice, changing how designers, artists, and illustrators work.

Conclusion

Both embedded and generative forms of AI are increasingly able to create content that resembles human production in areas as broad as image and music production. Undoubtedly, there is scope for AI technologies to replace individuals in particular areas of work. Commercial photographers could lose work because text-image generators can produce quality stock images; why use Getty or Alamy images with all the copyright restrictions when a new image could be directly created from a well-crafted prompt? However, the context determines how the technology is integrated into a particular domain. Will AI directly supplant human creative production in music and the visual arts? Not necessarily, for in the arts there is already an excess of production with many more artists than the market demands. Creative labour in this domain is, in part, underpinned by scarcity, with the field playing a key role in controlling the number of works passing through the galleries and creating market value. We have to appraise the use of AI within the broader system of creative communication. It should not be driven by tech companies posing solutions to non-existent needs. Instead, creative industries should articulate their own aims and then look for the best means of addressing them. This can be articulated through Gruner and Csikszentmihalyi's Creativity 4.0 model because they demonstrate that creativity depends on a whole range of social actors, not just the isolated individual experimenting with novel ideas. The question then shifts to the type of roles that AI will perform with respect to the model, and how this will change human systems of creative production. The Creativity 4.0 model provides a broad platform on which to consider these issues, but Gruner and Csikszentmihalyi's focus more on the replacement of creative individuals than on how AI mediates their relationship with the domain or provides a platform for creative practice. In adapting the model, we have clearly articulated the difference between generative applications and AI that is embedded into everyday platforms, such as Search. Generative applications might replace a creative individual, whereas embedded AI usually mediate between the creative individual and the domain or field. In some cases, especially those in which popularity underpins the evaluation, the technologies might serve as an extra field, creating memes and overly foregrounding worldwide popular content over niche local content. Programs analysing markets and trends, including stylistic patterns, push production towards what is already known under the logic of task appropriateness, whereas generative technologies might alter creative practice by producing ancillary content that can be incorporated into a human creative works. The rapid production of new content will require creative individuals to select outputs that best suit their own vision, and modify them in their own work, in which case, the creative individual is operating as a subset of the field. Technologies contract or expand the domain depending on how they are used and their inherent technological biases. The analysis of creativity will involve determining what aspects of the task are undertaken by AI, whether generative or embedded, and the degree to which human creative practice involves the selection rather than the generation of material, and future research must consider how to better cultivate creative divergence rather than just efficiencies in production.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

Paul Atkinson  <https://orcid.org/0000-0002-5250-0424>

Richie Barker  <https://orcid.org/0000-0001-5908-8973>

Notes

1. From this point forward, we use the term embedded AI to describe these types of systems.
2. It should be noted that Csikszentmihalyi's use of 'field' somewhat overlaps with social-theorist Pierre Bourdieu's use of the term, who talks about how persons of taste play a role in determining the value of art (1996: 292–93).
3. Big-C creativity refers to the role of exceptional creative individuals in developing new objects and ideas, which is underpinned by a notion of the creative genius.
4. The robot Ai-Da is the most well-known recent example (see <https://designmuseum.org/exhibitions/ai-da-portrait-of-the-robot>).
5. Many theorists have examined the relationship between digital technologies and attention. Bernard Stiegler talks about the role of the 'programming industries (film, television and digital media) in competing for attention, leading to such conditions as 'Cognitive Overflow Syndrome' (2008: 98). Katherine Hayles presents a more sanguine perspective, arguing that our brains can develop to deal with an 'information-intensive environment' (2012: 100).
6. The idea is directly borrowed from the work of Richard Dawkins, who argues that the activity of memes resembles many aspects of gene behaviour, including the capacity to move from one mind to another (similar to how a gene occupies bodies) and to replicate and spread throughout a culture (the meme 'can be said to propagate itself') (2016: 249). However, the analogy with biology should not be overemphasised, as memes do not operate like genes in an evolutionary context, for they rely on human action for replication (Shifman, 2014: 12).
7. Memetics has been adapted to the analysis of Internet communication, where this spread of information can be easily tracked. However, the popular use of the term meme does not entirely fit the academic definition, for it treats memes as short-term trends and fads, whereas academic study bases its argument on enduring trends that lead to significant cultural change (Shifman, 2014: 13).
8. Embedded AI effectively augments human activity: see Clark and Chalmers (1998) on the extended mind hypothesis and more recently, Katherine Hayles (2017) on machine-human assemblages.

References

- Adgully Bureau (2018) Adobe advertising cloud brings Sensei power to search advertising. *Adgully*, 12.
- Barker R and Atkinson P (2019) A new master of the cultural domain? Google's role in the transactive memory systems of advertising copywriters and art directors. *Media International Australia* 172(1): 61–73.
- Boden MA (1995) How artificial intelligence can help our creativity. *Convergence: The International Journal of Research into New Media Technologies* 1(1): 20–22.
- Boden MA (2009) Computer models of creativity. *AI Magazine* 30(3): 23–34.
- Boden MA (2010) *Creativity and Art: Three Roads to Surprise*. Oxford: Oxford University Press.
- Boden MA and Edmonds EA (2009) What is generative art? *Digital Creativity* 20(1–2): 21–46.
- Bourdieu P (1996) *The Rules of Art: Genesis and Structure of the Literary Field*. Cambridge: Polity Press.
- Campbell C, Plangger K, Sands S, et al. (2022) Preparing for an era of deepfakes and AI-generated ads: A framework for understanding responses to manipulated advertising. *Journal of Advertising* 51(1): 22–38.
- Clark A and Chalmers D (1998) The extended mind. *Analysis* 58(1): 7–19.
- Csikszentmihalyi M (1996) *Creativity: Flow and the Psychology of Discovery and Invention*. New York: HarperCollins.

- Csikszentmihalyi M (1999) Implications of a systems perspective for the study of creativity. In: RJ Sternberg (ed), *Handbook of Creativity*. Cambridge: Cambridge University Press, pp. 313–335.
- Dawkins R (2016) *The Selfish Gene*. Oxford: Oxford University Press.
- Erden YJ (2010) Could a created being ever be creative? Some philosophical remarks on creativity and AI development. *Minds and Machines* 20: 349–362.
- Franklin S (2014) History, motivations and core themes. In: K Frankish and WM Ramsey (eds), *The Cambridge Handbook of Artificial Intelligence*. Cambridge: Cambridge University Press, pp. 15–33.
- Gervas P (2009) Computational approaches to storytelling and creativity. *AI Magazine* 30(3): 49–62.
- Gillespie T (2014) The relevance of algorithms. In: T Gillespie, KA Foot, and PJ Boczkowski (eds), *Media Technologies: Essays on Communication, Materiality, and Society*. Cambridge: MIT Press, pp. 167–193.
- Gruner DT and Csikszentmihalyi M (2018) Engineering creativity in an age of artificial intelligence. In: I Lebeda (ed), *The Palgrave Handbook of Social Creativity Research*. Palgrave Macmillan, pp. 447–462.
- Hayles NK (2012) *How We Think: Digital Media and Contemporary Technogenesis*. Chicago: University of Chicago Press.
- Hayles NK (2017) *Unthought: The Power of the Cognitive Nonconscious*. Chicago: University of Chicago Press.
- Hillis K, Petit M, and Jarrett K (2013) *Google and the Culture of Search*. New York and London: Routledge.
- Ho J (2021) High fidelity image generation using diffusion models. *Google AI Blog*. Available at: <https://ai.googleblog.com/2021/07/high-fidelity-image-generation-using.html>
- Kaiser Z (2019) Creativity as computation: Teaching design in the age of automation. *Design and Culture* 11(2): 173–192.
- Liu G (2022) The world's smartest artificial intelligence just made its first magazine cover. Available at: <https://www.cosmopolitan.com/lifestyle/a40314356/dall-e-2-artificial-intelligence-cover/>.
- Love T (2018) OK composer. *The Guardian Weekly*, 2 November, pp. 55–56.
- Luce L (2019) *Artificial Intelligence for Fashion How AI Is Revolutionizing the Fashion Industry*. Berkeley: Apress.
- Nelson D (2020) What is a generative adversarial network (GAN)? Available at: <https://www.unite.ai/what-is-a-generative-adversarial-network-gan/>.
- Pajkovic N (2022) Algorithms and taste-making: Exposing the Netflix recommender system's operational logics. *Convergence: The International Journal of Research into New Media Technologies* 28(1): 214–235.
- Ramesh A, Pavlov M, Goh G, et al. (2021) Zero-shot text-to-image generation. *Proceedings of Machine Learning Research* 139: 8821–8831. Available from: <https://proceedings.mlr.press/v139/ramesh21a.html>
- Ritchie G (2009) Can computers create humor? *AI Magazine* 30(3): 71–81.
- Sawyer RK (2012) *Explaining Creativity: The Science of Human Innovation*. 2nd edition. Oxford University Press.
- Scott LM (1994) The bridge from text to mind: Adapting reader-response theory to consumer research. *Journal of Consumer Research* 21(3): 461–480.
- Shifman L (2014) *Memes in Digital Culture*. Cambridge: MIT Press.
- Simpson LC (1995) *Technology, Time, and the Conversations of Modernity*. London: Routledge.
- Stiegler B (2008) *Taking Care of Youth and the Generations*. Stanford: Stanford University Press.
- Susskind R and Susskind D (2015) *The Future of the Professions: How Technology Will Transform the Work of Human Experts*. Oxford: Oxford University Press.
- Thorne S (2020) Hey Siri, tell me a story: Digital storytelling and AI authorship. *Convergence: The International Journal of Research into New Media Technologies* 26(4): 808–823.
- Vakratsas D and Wang X (2021) Artificial intelligence in advertising creativity. *Journal of Advertising* 50(1): 39–51.
- Wajcman J (2019) The digital architecture of time management. *Science, Technology, & Human Values* 44(2): 315–337.

Author biographies

Dr Paul Atkinson teaches within the Media, Film and Journalism program at Monash University, Australia. He has published widely on a range of media, from cinema and animation to dance and comic books, in addition to his work on philosophy, technology and time. He examines aspects of the relationship between time, creativity and aesthetics in his current book *Henri Bergson and Visual Culture: A Philosophy for a New Aesthetic* (Bloomsbury, 2020).

Dr Richie Barker Deakin University, Australia Richie Barker is a communication academic with the Faculty of Arts and Education at Deakin University, Australia. His research focuses on theories of creativity in the professions, and has written a number of articles on the role of new technologies in mediating the creative process.