DEEPAK SOMAYA AND LAV R. VARSHNEY

Ownership Dilemmas in an Age of CREATIVE MACHINES

Developments in computational creativity are leading to a new era of innovation. Intellectual property rights must keep up.

Recently, a computer-generated work of art titled "Edmond de Belamy, from La Famille de Belamy" was sold by the auction house Christie's for $350,000 to an anonymous bidder. Purported to be the first auctioned portrait generated by artificial intelligence (AI), the work, produced by the French art collective Obvious, fetched hundreds of thousands of dollars more than works by Andy Warhol and Roy Lichtenstein offered in the same auction. Almost immediately, authorship of the painting was contested. An artist and programmer named Robbie Barrat claimed on Twitter that Obvious produced the painting using an algorithm he had created and shared online. The issues raised by this dispute may be the first of many profound questions of ownership, attribution, and intellectual property rights for the burgeoning field of computer-generated artwork.

As AI makes seemingly inexorable progress into every field of human endeavor, it may be inevitable that smart machines will become central to creativity and innovation, activities often considered to represent the highest form of human intelligence. In fact, computational creativity—an emerging field concerned with algorithms that can produce creative or inventive artifacts—has already made significant inroads in many real-world applications. In endeavors as varied as cooking, literature, fashion, circuit design, and drug discovery, AI systems are now able to produce ideas and artifacts that meet the criteria of novelty, surprise, and usefulness that lead to them being judged as creative by human experts.

These AI technologies challenge the fundamental building blocks of existing intellectual property (IP) laws and institutions, which are misaligned with AI-driven innovation on multiple fronts. IP rights are intended to provide incentives for innovators to engage in creative endeavors and to bring the fruits of these activities to society, while simultaneously balancing the need for market competition and dissemination of new knowledge. IP laws put human inventors and creators at the center of the creative process, reflecting deep-rooted assumptions about the inherent humanness of creativity. These assumptions have now been overturned by advances in computational creativity.

Algorithmic dreamin'

Psychologists define creativity as the generation of a product or service that is novel and judged to be appropriate, useful, or valuable by a knowledgeable social group, and simultaneously generates a measure of surprise, beauty, or amazement. To produce innovative songs, paintings, writings, or other artifacts that meet these standards, computational creativity systems use a variety of algorithmic techniques, including genetic algorithms, simulated annealing, stochastic sampling and filtering, and deep neural networks. The particulars of how these systems function is not as important as understanding that they are producing work that appears to be the product of creative thinking.

Tests analogous to the Turing test for assessing intelligent behavior by machines have been recently proposed by AI researchers such as Mark Riedl, whereby algorithms can be judged "creative" if they exhibit behaviors that human observers consider to be creative. A number of systems have now created artifacts in specific domains that pass this test. Examples that have captured public attention include Google's Magenta system, which composes novel and pleasing music, and IBM's Chef Watson system, which produces new and flavorful food recipes.

The Magenta system uses a form of deep learning called generative adversarial networks (GANs), which rely on training data to create a model of artifacts in the creative domain. These data provide the network with notions of quality and how different conceptual pieces go together. The algorithm learns how to compose songs by "listening" to existing, human-authored compositions. Then the system produces novel compositions by sampling from the learned model.

The Chef Watson system uses a different approach. It explicitly models how humans perceive different food tastes (what experts call the "hedonic psychophysical properties" of ingredients) and their combining rules. Remixing existing recipes generates many possibilities; the algorithm selects the most surprising and flavorful ingredient lists suggested by the model. The system then determines recipe steps and ingredient proportions. The design and use of this creation system involved human collaborators, both in structuring the data used to train the system and in cocreating the recipes it developed. The celebrity chef James Briscione described Chef Watson as an innovation partner, saying, "Watson forced me to approach ingredients without any preconceived notions of which ingredients pair well together."

Creative computational machines can operate autonomously, such as systems that write (serviceable) business and sports news or (quite bad) poetry. These systems require little human oversight to produce their creations. However, they can also operate semiautonomously, producing novel artifacts in collaboration with human experts, as illustrated by Chef Watson and older systems such as AARON the painter, which was hand-coded by its creator, Harold Cohen, and the Automated Mathematician, an early AI system that its creator, Douglas Lenat, claimed could produce mathematical conjectures.

Margaret Boden, a scholar who has written extensively on the potential of computational creativity, suggests two areas in which individuals with deep domain knowledge can enhance innovation in conjunction with creative machines. First, such expertise may be important for defining the conceptual space and specifying procedures to explore that space's creative potential. These inputs may need to be iterative, as the richness of human understanding about a problem domain may be difficult to articulate in a single instance and the domain may itself evolve, requiring creative systems to be refined and modified over time.

Second, domain expertise may be important for incorporating knowledge about human—often synonymous with "customer"—values or tastes, which are quite difficult to identify and especially difficult to express in computational form. In an illustration of these ideas, expert chefs were found to significantly enhance the inherent innovation in Chef Watson's recipes by infusing the recipes with the human chefs' own deep and tacit understanding of cooking methods and customer tastes. When this human-machine collaboration is successful, or when a fully autonomous system creates a desirable consumer product, the results can be quite profitable—and may raise difficult questions about appropriate IP protections.

Intellectual property policy

Intellectual property regimes seek to foster innovation by stimulating discovery and by fostering the dissemination and use of creative artifacts. "Use" implies not simply employing an artifact for some purpose, but also the creative reuse of artifacts in cumulatively developing more creative products, such as the transformation of artistic creations (e.g., remixing in visual art and music) and the cumulative development of new technologies (e.g., inventions that build on prior ones). IP protections give innovators a private incentive to develop new creations, but at the social costs of monopoly, since they grant innovators exclusive rights to their creations for some defined period.

Machine-based innovations raise a number of challenges for these classic trade-offs embedded in IP policy. For one, AI-driven innovation takes place on a spectrum of autonomy, ranging from fully autonomous algorithms to various forms of semiautonomous operation in collaboration with humans. IP institutions such as the patent system, which grants exclusive rights to the inventor(s) of a product or service, are often not aligned with the complexities of such collaborations, or with the relationships among AI creators, data owners, and domain experts. And as a general-purpose technology, creative AI is likely to produce outputs in sectors and industries that vary widely in the institutional contexts within which innovation occurs; the corresponding role played by IP in fostering innovation will differ as well.

The IP rights relevant to computational creativity are copyrights (which are our focus here) and patents, but they also include related "lesser" rights such as design rights and utility models. A central tenet of both copyright and patent law is that an IP right stems from and vests in a human creator or inventor. According to the US Copyright Office, a work must be the product of "human authorship" to be entitled to copyright protection. Similarly, US patent law requires that each patent application should name the inventor, defined as "a person [who] contributes to the conception of the invention." These provisions would appear to constrain the available options of IP protection for machine-created innovations. The issue, at least for the foreseeable future, is not that AI deserves some kind of moral personhood. Instead, the question is how one assigns ownership over AI-created (or cocreated) artifacts in such a way as to strike the right balance in IP policy between private innovation incentives and social knowledge recombination.

Some commentators and scholars have argued both conceptually and empirically that purely computer-generated creative artifacts are a myth. When it comes to IP protections, they argue, the work-for-hire doctrine, in which an employer owns all rights to the creation of an employee, can be used to assign copyright to owners for creative works produced by AI. Similarly, the current copyright and patent case laws appear to offer options for machine-created innovations: not assigning IP rights and letting new advances fall into the public domain, or assigning the rights to some human cocreator or owner/operator. However, another option, whereby the creator of the AI system becomes a co-owner of IP rights awarded for the AI's creations, also deserves serious consideration, particularly in situations where monetary incentives are needed to motivate the creation, upgrading, application, refinement, or recombination of AI algorithms.

Continued on page 84→



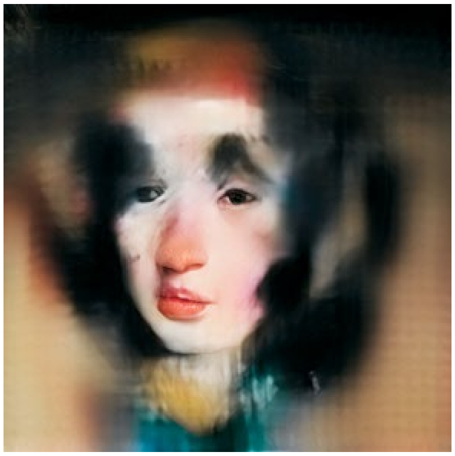

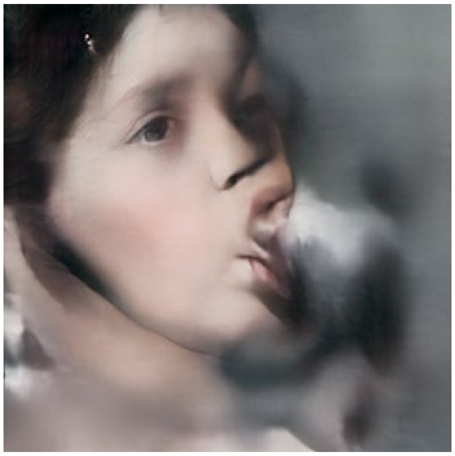


Mario Klingemann

Mario Klingemann describes himself as "an artist and a skeptic with a curious mind." His work spans and incorporates a wide range of tools and technologies, including neural networks and deep learning, computer code and algorithms, and artificial intelligence and generative art (work produced by autonomous or semi-autonomous systems). The driving force behind his evolving aesthetic is an interest in the idea that artificial intelligence has the potential to generate surprising new images and perspectives. As technology advances, the question then becomes for the artist, "What does one hope to find?"

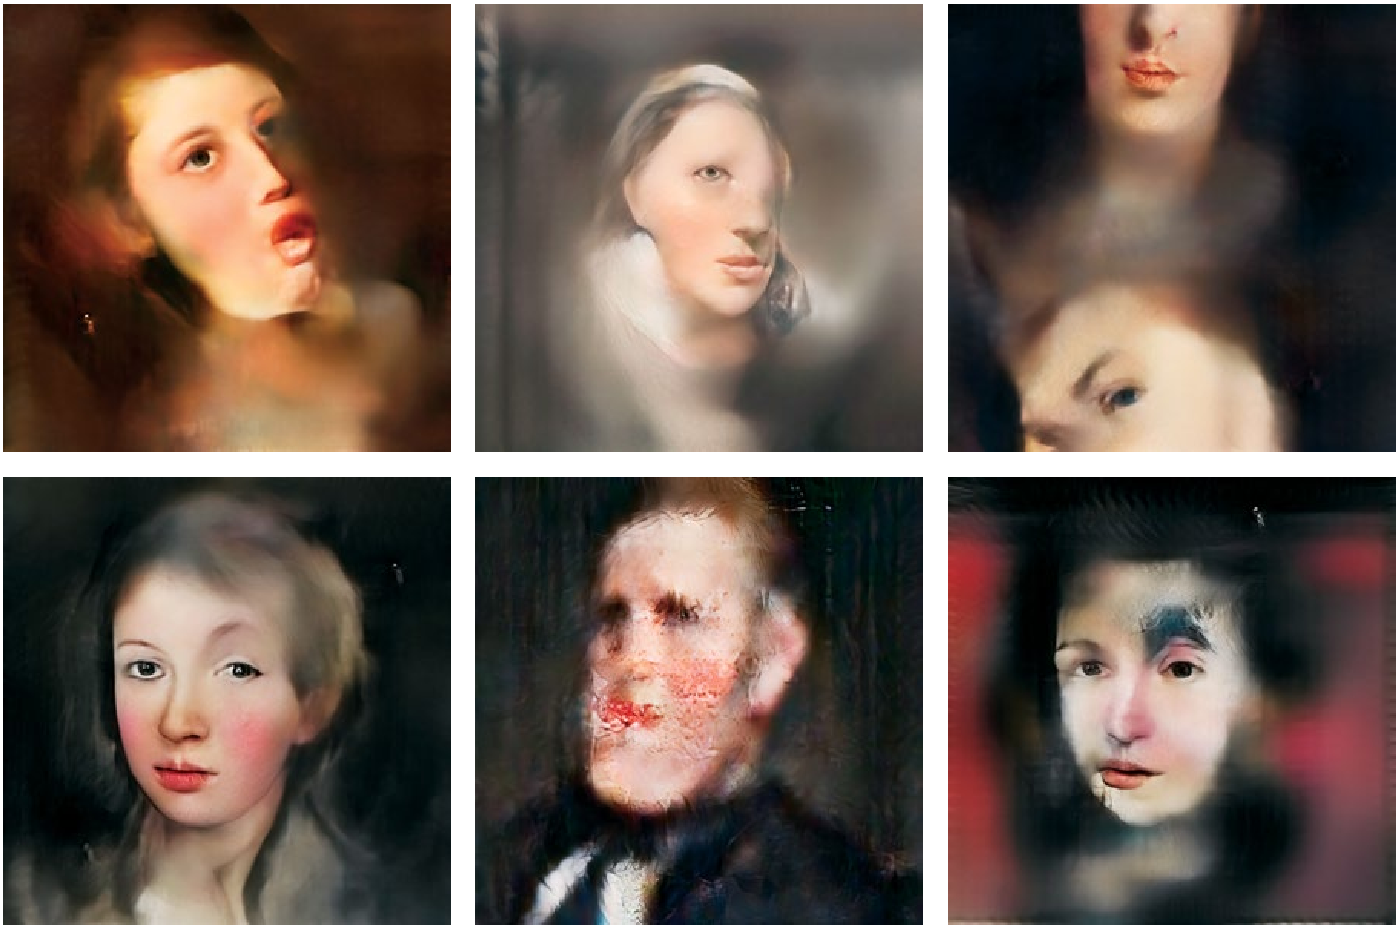

All images by MARIO KLINGEMANN from the series *Neural Glitch*, 2018

Neural Glitch is a technique Klingemann started exploring in April 2018, in which he manipulates generative adversarial networks, or GANs, which are a class of machine learning system that use training data sets to create new data. Due to the complex structure of the GANs' neural architectures, the introduced glitches cause the models to misinterpret the input data in interesting ways, some of which could be interpreted as glimpses of autonomous creativity. According to Klingemann, "One interesting aspect of this process is that on one side the same input data can yield very different results depending on the glitch, whilst at the same time different input data, transformed by the same glitched model chain, will result in a coherent style and show the same semantic misinterpretations."

As a pioneer in producing art with artificial intelligence, computer learning, and other technologies, Klingemann's work has been shown at the Museum of Modern Art in New York City and the Centre Georges Pompidou in Paris. He has worked with a variety of prestigious institutions, including the British Library, Cardiff University, and New York Public Library. He is currently an artist in residence at Google Arts & Culture.

Domain differences

Before discussing options for IP ownership more fully, considering the domain in which the AI is operating illustrates how various creative domains have different levels of IP protections and relationships with IP regimes. In some domains, pervasive IP rights can potentially stifle innovation by making it time-consuming and often very costly to innovate by building on others' creative work. In a recent example, the artist Alexander Reben commissioned an artwork based on several images he believed to be randomly generated by the AI algorithm GANbreeder, which mashes together multiple input images in different ways. In this context, the economic incentives needed to create new images may be small, and the potential creativity that might be unleashed by recombining images in different ways is apparently substantial. Imagine how challenging it would be for an AI artist using GANbreeder to obtain licenses to hundreds of images created by others (which might themselves be derived from other images) before the artist could even begin to explore creative ideas, to say nothing of images used as training data for machine learning models.

Thus, some applications of creative AI may run into the so-called anticommons problem, where the proliferation of strong IP protections can stymie innovation within a field because innovators find it too costly or even impossible to access all the rights needed to create new artifacts. The potential for an AI anticommons highlights the fact that IP policy for AI creations cannot be isolated from the institutional features of the domain in which innovation occurs.

Another example of the domain-specific character of innovation is IP-negative spaces, wherein IP rights have traditionally been very limited or absent. These spaces include fields such as fashion, cuisine, tattoo artistry, graffiti, financial services, and sports. Nonetheless, significant innovation occurs in these domains, presumably because there are other innovation-promoting factors at play. For example, creation in some settings may rely less on extrinsic rewards and be driven by the intrinsic motivation to produce socially enriching and meaningful creations. Such intrinsic motives may even be socialized (and reinforced) in the creative communities within such domains.

Innovators in IP-negative spaces may also derive value and motivation from the fame and respect of peers, which may offset (to a degree) their need for financial rewards. Moreover, due to these reputational benefits, and simply by being first to market, these innovators may earn returns from their innovations even without strong IP rights. For example, fashion designers might rely on their personal reputations as innovators and trendsetters to run profitable businesses, despite having no formal IP protection for their innovative designs.

A final set of important IP policy issues stems from the potential of AI technologies to rapidly produce creative artifacts, in turn creating opportunities for supercharging innovation through the use and recombination of these artifacts in novel ways. All things being equal, the ability of innovators to take advantage of these opportunities by building on others' AI creations will likely be stymied by a policy that supports strong IP rights (for the usual anticommons reasons). One can easily imagine a musically inclined AI system composing thousands or millions of tunes, with an opportunistic rights-holder suing anyone who remixed an AI-produced melody. However, where such rights may be necessary to provide appropriate incentives, there could be other policy tools that help create and support transactions in AI creations.

Ownership of IP

So who should own the intellectual property rights for creations generated autonomously or semiautonomously by AI systems? Should anyone? There are three perspectives regarding creative AI and the IP ownership regime needed to support it.

First, computational creativity may be viewed simply as a tool, which enhances creative or innovative output but has no inherent impact on downstream IP rights. This approach aligns with the point of view that robots should be treated simply as other technological tools that might enhance creativity and productivity (e.g., software packages or printers). The IP resulting from AI would then vest in the human coinventors or collaborators of creative machines that operate in semiautonomous mode. Such an approach may be appropriate for settings in which the role of these collaborating humans is quite significant, particularly in adapting the creative output to the tastes and values of the intended audience and when the context itself is one that is not characterized by significant IP-negative features or anticommons concerns.

Second, the creative output of AI may be inherently unsuitable for IP protection, such that it automatically enters the public domain. A non-IP policy of this kind may be appropriate for AI systems that innovate in fully (or substantially) autonomous mode, where the role of downstream human cocreators is relatively minor and therefore need not be substantially incentivized. Denying IP rights to creative-AI artifacts from such systems might also be compatible with IP-negative spaces, such as fashion or tattooing, where the potential anticommons problems of strong IP protections would otherwise be significant.

Lastly, upstream developers of the creative AI systems could be granted IP rights over innovations produced by their AI, potentially in addition to the IP rights granted to users or operators of the system. To use the example of an AI-produced painting, the algorithm developer and the artist using the AI to create a painting could both be granted IP rights. In the two ownership regimes described above, the incentives for upstream innovators of creative machines would derive entirely from IP rights in the AI systems they created, and they are unable to secure any contractual extension of these rights when selling these systems. However, allowing upstream AI innovators to co-own the creative output of their AI systems would better calibrate the innovators' incentives to the ultimate value their products generate. This alignment may be especially valuable when upstream incentives are needed to encourage the development of better computationally creative technologies for a given domain, such as portraiture or musical composition. Such a policy likely will not be appropriate for IP-negative spaces, where one of the goals of IP policy is to prevent an IP anticommons from impeding innovation.

Although these principles for IP ownership over AI creations may be appealing, there are practical impediments to implementing multiple policies targeted at different domains. Notably, firms may try to game the system and seek out ownership regimes that best suit their private interests. For example, a software company that builds creative AI for the fashion industry may lobby for inappropriately strong IP protections for its system's creations, creating an anticommons problem where none previously existed. It may be difficult to design bright lines that demarcate domains within which each regime is implemented.

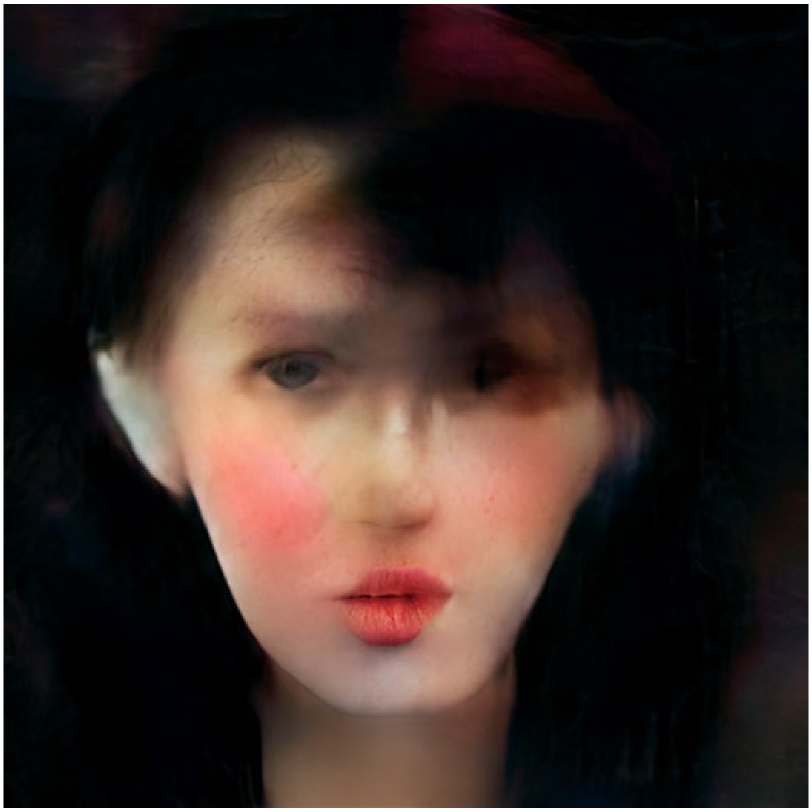

Robot remuneration

One potential policy response is the creation of a sui generis IP right, tailored to the needs and contexts of creative-AI innovation. Among the features of such an IP right might be a much shorter duration of the right, which would put AI creations more quickly into the public domain and allow other innovators to build on or recombine them. Copyrights in the United States, for example, last 95 years, a duration that has its roots in the lifespan of human authors and their descendants. Such a long period for IP rights is anachronistic for the fast-moving fields likely to be created by AI-driven innovation.

Policy-makers may also examine approaches that directly facilitate and support transactions in IP-protected creative artifacts. Among the possibilities are a default licensing regime whereby all such IP is made available to others (for a price), and a (possibly blockchain-enabled) market-creating mechanism that tracks usage of IP and estimates its relative value to putative or actual licensees. The more useful and unique the creative artifact is for potentially generating further innovations, the greater the price it would command. If an AI-generated song is such a hit that it becomes a key ingredient in producing new chart-topping remixes, the original would be automatically licensed but at a higher price than other, less consequential AI-generated music. Such a default licensing system that encourages recombination and ensures a modicum of fair reimbursement to upstream creators would strike the right balance for supporting innovation in many domains.

The distinction made between licensing and sale in IP law may create perverse incentives to treat all transactions in digital AI creations as a license. IP owners lose many rights under the "first-sale doctrine" when they sell a product (after the initial sale, the new owner of a copy can resell it without permission from the copyright holder). But IP owners are able to use licensing contracts quite flexibly and even obtain "reach-through rights," which give an inventor rights to downstream uses of the product in subsequent innovation. Why then would any IP holders want to sell a digital product when they could license it? These incentives may need to be realigned by preserving some rights for innovators even when they sell AI systems, and at the same time constraining the use of some types of licensing contracts that interfere excessively with the recombination of innovations and creative artifacts produced by those AI systems.

Of course, our policy suggestions are tentative and meant to spur further discussion. Creative artificial intelligence, whether fully autonomous or in collaboration with humans, shows significant promise for increasing the pace of innovation in a variety of endeavors. Intellectual property policies and institutions must keep pace if these innovations are to be properly incentivized and rewarded.

***Deepak Somaya** is a professor of strategy and entrepreneurship, and a professor of law (by courtesy), at the University of Illinois at Urbana-Champaign.* ***Lav R. Varshney** is an assistant professor of electrical and computer engineering, and an assistant professor of computer science, neuroscience, and industrial engineering (by courtesy), at the University of Illinois at Urbana-Champaign. He is also a principal research scientist at Salesforce.*