This is a section of doi:10.7551/mitpress/10913.001.0001

Beyond the Creative Species

Making Machines That Make Art and Music

By: Oliver Bown

Citation:

Beyond the Creative Species: Making Machines That Make Art and Music
By: Oliver Bown
DOI: 10.7551/mitpress/10913.001.0001
ISBN (electronic): 9780262361750
Publisher: The MIT Press
Published: 2021

The open access edition of this book was made possible by generous funding from University of New South Wales.

The MIT Press

3 Creativity and Culture

The brain is merely Nature's latest means of self-preservation.
—W. Ross Ashby[1]

Everything is the way it is because it got that way.
—D'Arcy Thompson[2]

Our species is the only creative species, and it has only one creative instrument, the individual mind and spirit of man. Nothing was ever created by two men. There are no good collaborations, whether in music, in art, in poetry, in mathematics, in philosophy. Once the miracle of creation has taken place, the group can build and extend it, but the group never invents anything. The preciousness lies in the lonely mind of a man.
—John Steinbeck[3]

Variables external to the individual must be considered if one wishes to explain why, when, and from where new ideas or products arise and become established in a culture.
—Sami Abuhamdeh and Mihaly Csikszentmihalyi[4]

Not until Vespucci recognized that the so-called West Indies were part of an entirely different continent did Columbus's almost superhuman efforts get retrospectively revised as a "discovery." And if in the fullness of time it turns out that it was Erik the Red who really discovered America, that "discovery" will be as much a result of scholarship and politics as a result of Erik's travels.
—Mihaly Csikszentmihalyi[5]

From Individuals to Distributed Creative Systems

Creativity in Humans, Nature, and Machines

This book began with reference to a familiar vision of robots walking among us, perhaps indistinguishable from us, perhaps quirky reflections of ourselves. In such visions, it is simple but effective to treat the world as if there are two types of things: humans and other objects. Humanlike artificially intelligent robots make good material for fiction, conceived through the interaction between protagonists in the medium of language. In the narratives of films like *Ex Machina*, they pop into existence from the world of mundane objects in the hands of a lone genius. The glacial gradualism of technological advance is skipped over; it just doesn't make good movie material. Inevitably, however, the distinction between humanlike intelligence and simpler algorithms must reveal some kind of fuzziness at the boundary. Today's AI systems are largely (but not entirely) utilitarian tools with only restricted autonomy, but as AI gets more powerful, the gap between people and things is being gradually filled in with a spectrum of machines of varying degrees of intelligence and autonomy.

For many theorists of agency in the social sciences, this is a common and problematic dualism. Andrew Pickering, for example, refers to it as asymmetrical dualism, according to which "the world is a knowable machine, while … humanity calls the shots as the only locus of genuine agency."[6] Pickering describes this as "the natural ontological attitude of modernity, the way we all tend to approach the world."[7] For the anthropologist of art Alfred Gell, according to this default perspective,

> whereas chains of physical/material cause-and-effect consist of "happenings" which can be explained by physical laws, which ultimately govern the universe as a whole, agents initiate "actions" which are "caused" by themselves, by their intentions, not by the physical laws of the cosmos.[8]

At the same time, the discussion in chapter 2 worked with the broad simplification that individual human agents perform creative acts in relative isolation, given that they have the cognitive machinery to create. An alternative perspective that will be explored in detail in this chapter is that individuals are always operating situated in networks of creative activity, making *contributions* to a process of creativity. That is, creativity really only happens at a super-individual level which cannot be reduced to individual cognition. While the cognitive machinery of individual brains is critical to this process, it is not the only ingredient, and anyway leads to effects that are only manifest and observable at the network level. Breaking the dualism of human-object relations is important to developing this framework and also in building unbiased models of creative processes that *admit* a creative role for machines, of all kinds, that fall short of humanlike cognition and social situatedness.

A social perspective on creativity is most notably associated with the work of Mihaly Csikszentmihalyi, in his call for researchers of creativity "to abandon the Ptolemaic view of creativity, in which the person is at the center of everything, for a more Copernican model in which the person is part of a system of mutual influences and information."[9] The result is a subtle shift from the question of how a given entity *is creative* to what role an entity plays in contribution to a larger creative process. Such a shift in stance allows for the study of social aspects of creativity, which as we will see are not insignificant. But it also means that nonhuman entities—including advanced AI systems, but also dumb algorithms, as well as everyday objects and materials—can more easily be treated as contributing to accounts of creative processes, without having to be understood as creative agents themselves.

This chapter proceeds through three sections. Firstly, I undertake this review of creative systems as collective phenomena. Secondly, I look specifically at how evolutionary thinking informs the understanding of creativity. Thirdly, I convert this into a useful set of concepts that helps us analyze different creativity scenarios, in particular locating artistic practices and movements among human creative activity.

Socially Situated Individuals

It is widely acknowledged that groups of people putting their heads together enhances the effect of creative search. Dean Simonton's studies of scientific creativity[10] reveal a pattern of distributed creative effort, looking at the relationship between quantity and quality of a researcher's output, which supports his blind variation approach to creativity. Simonton's data shows that increasing quantity increases the chance of producing a quality output but also increases the chance that the researcher produces junk. "It is significant," he says, "that those who publish the most highly cited works also publish the most ignored works, so that quality is a probabilistic consequence of quantity. In fact, the ratio of high-impact publications to total output—the ‘hit’ rate—is uncorrelated with total output. This has been called the *equal-odds* rule.” This lends evidence to the view that, as far as individual creative success is concerned, chance is an important factor in the success of creative work; raising one’s quantity of output is one of the most accessible strategies for increasing the chance of making a breakthrough, and increasing output becomes a genuine part of an overall creative strategy.[11] For Simonton, this means that the same strategy of blind variation and selective retention that works for an individual also works across a social group.

Over a large enough body of output, given an underlying stochastic process, the chance of quality output becomes increasingly certain. In Simonton’s view the creative capacity of individuals is revealed over long rather than short time scales, and creative productivity would appear to be more deterministic at a population level than at an individual level. Note once again an important clarification: this does not mean that chance is the only factor. This argument has (almost) nothing to do with a million monkeys managing to write a good play; the researcher must be an expert in the field and must apply methods appropriately. The production of quantity must be production at a certain threshold of quality. In the sciences, peer review and other forms of quality control further feed into this, influencing the quantity of output in the first instance, setting basic checks and balances that are reliable enough to successfully constrain the blind search toward a sensible subset of candidates. Simonton’s work tallies with Thomas Kuhn’s[12] influential account of the scientific process. Scientific progress takes place, Kuhn argues, as a series of periods of normal science, punctuated by rare, high-impact scientific revolutions. Normal science tends to be more formulaic and may be less prone to chance effects on success. But the paradigm underlying the stable progress of normal science becomes strained by the increasing demands of explaining new data, necessitating revolution. As we will see, other domains of collective activity experience such cycles of stability and disruption.

Thus a socially situated perspective illustrates one very simple way in which creativity scales from individuals to groups. Simply put, the process of creative search that any individual engages in works perfectly well as a distributed process. Ten people searching for a solution to a problem will be more effective than one. One individual from the ten may independently find a winning solution, but it would be mistaken to infer in hindsight that the other nine people were redundant. That would be to ignore the stochastic nature of creative search. As well as distributing the process of generation, this introduces a hierarchy of levels at which evaluation is conducted. In Simonton's example of academic work, the individual researcher generates and evaluates outcomes, which are further iterated through the peer-review process, and then again are placed in the market of long-term collective evaluation, in which some work will sink and other work will swim. In Csikszentmihalyi's words, "It has been said that 99% of all new ideas are garbage, regardless of the domain or the status of the thinker. To sift out the good ideas from the bad, another system is needed."[13]

But except when faced with some imminent task in a restricted domain—say, extracting a dropped piece of clothing from behind a wardrobe with a coat hanger—being socially situated means, among other things, not only performing creative tasks but first of all making choices about which tasks to take on. Much of the research into the psychology of creativity therefore deals with a more holistic approach to the strategies involved in creative success. As Csikszentmihalyi puts it, "We cannot study creativity by isolating individuals and their works from the social and historical milieu in which their actions are carried out. This is because what we call creative is never the result of individual action alone."[14]

The study of individuals as problem seekers as well as problem solvers—choosing where they will situate themselves to perform creative work—leads to theories of creative individuals as situated in a broader social context, involving a wider set of strategies for getting by in a world where creativity is not only a way of solving imminent problems but a valuable social commodity. Robert Sternberg's investment theory of creativity,[15] for example, draws a parallel between the world of financial investment and creative work: creative individuals are those who "buy low and sell high" in the world of ideas. Some may choose to avoid the mainstream and go for the higher risk outside chances. In doing so they take on the risk of failure but also increase the possibility that they will hit upon a more radical, creatively explosive outcome with greater rewards.

The parallel with investors helpfully frames a fundamental relation between problem-solving capability and *chance*. In a hypothetical example by the "black swan" investor Nassim Taleb, related by *New Yorker* writer Malcolm Gladwell,[16] we are asked to imagine ten thousand investment managers winning and losing money on the markets. Imagine that each year, each investor wins or loses money completely by chance, with an even chance of winning or losing. Follow this thought experiment through for ten years (you could easily model this in an Excel spreadsheet); the statistics have it that nine of those investors will have made money every year in a row. While this might give the impression to an observer that these nine investors had some special competence not possessed by the others, this number is the consequence of a model in which all investors are equally skilled and knowledgeable, and winning is a completely chance event. Furthermore, it is possible in some circumstances that investors with higher risk strategies will come out richest. Someone will have become fantastically rich, far richer than could be possible by taking only the safe bets. The point of this thought experiment is to highlight how easy it is to miss the large number of people, employing *exactly the same strategy*, who are worse off than if they had played a safe game (this common misjudgment is known as *survivorship bias*, focusing only on those who have survived and thrived in some risky situation and ignoring those who haven't). By analogy, while we tend to assume that those who are more creatively successful must be more creatively capable, a situated view reveals an alternative possibility, with a greater emphasis on chance.

This is the backdrop to the problem of creative search for individuals, a reminder that there need not be as strong a correlation between success and capability as we might expect. With investors as with creative work, skill and strategy will nevertheless improve your chances. Investors have knowledge about what they are investing in, both specific to the field of expertise—understanding how IPOs work, say—and more generally in terms of approaching risk—understanding how to spread bets or how to manage long- and short-term investment strategies properly, the factor in which Taleb was particularly prescient.

And so with creative work, approaches to risk management and uncertainty in the world of creative ideas, combined with a certain domain-specific expertise, are dominant factors in understanding how individuals best set themselves up for creative success. In a similar way, David Perkins[17] compares the risk involved in creative work with that of gold diggers trawling the land for gold deposits. He dubs the conceptual spaces that we search for creative outcomes "Klondike spaces," after the gold rush of the Klondike region of the Yukon in Canada. The analogy works on several levels. Firstly, upon hearing about the gold of the Yukon the first step of the hopeful gold digger is to head to that region, rather than doubling down on their efforts to find gold in their own backyard. This is obvious but important, because any creative person starts by making several apparently very uncreative steps: choosing a field, getting to know the area, learning from others, trying to imitate their methods. Once on the gold fields, the problems of search begin proper. Now that you've traveled to an area known for gold, you'd hope that you would find it everywhere you dug, but at this finer resolution it turns out that gold in one place does not necessarily mean gold in an adjoining area. Perkins thus describes the features of creative problems through this metaphorical landscape: there is the "wilderness of possibilities," the simple fact that high-impact outcomes are rare; the "clueless plateau," the "lack of signs that point to the gold," even once you've found traces of gold already; the "narrow canyon of exploration," the fact that you are generally forced by practical constraints to restrict search to a specific local area, which may yield no results; and the "oasis of false promise," that it is overwhelmingly tempting to carry on looking in your current location.

Perkins builds this analogy with evidence from studying a variety of creative practitioners, where he witnesses the kinds of heuristic strategies that suggest this kind of terrain. Creatively successful people, his evidence shows, don't tend to work on just one thing but have a number of separate activities on the go, keeping each up in the air in the hope of a breakthrough. You could say they are spread betters, and following Sternberg's metaphor, they are gauging not only the chance of success but the impact: comparing hitting the jackpot on a wild outside chance with building a number of safe but small wins.

As with Wallas's four-stage model, this broader view of creative practice can be framed as a series of iterated phases of production and reflection, in which both the goal and the strategy are reviewed. Todd Lubart shows further how artists structure such processes around their specific relation to a given task. For example, they perform more exploratory association tasks after breaks when they are coming to their work cold, rather than performing evaluation tasks then.[18]

The evolutionary computing researchers Ken Stanley and Joel Lehman develop this argument from a computational perspective,[19] claiming that creative discovery can't follow a targeted approach. In a targeted approach, you think you have a good clue about what the solution to your problem

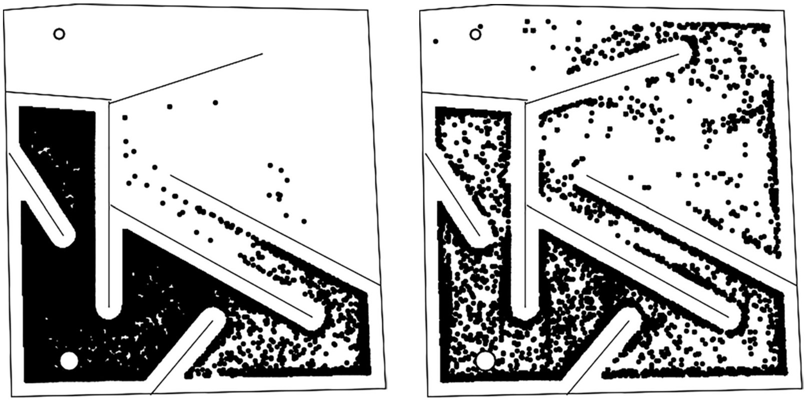

Figure 3.1
Lehman and Stanley's original demonstration of the effective power of novelty-based search, even for traditional goal-based search tasks. In both images, the large white dot (bottom left) is the starting point for an agent that is trying to find its way to the small white dot (top left) through the maze, within a given time limit. The black dots represent the final resting place of the agent in many different trials. In the left panel the agent search strategy was evolved using target-based evolutionary search: agents that got nearer the target were deemed fitter and more likely to seed future generations. However, the shape of the maze means that fitness might be rewarded for agents that are close to the target as the crow flies but not in reality. In the right panel, novelty search is used: fitness is rewarded to agents whose resting position is novel—that is, dissimilar to previous agents. The images show that the novelty search strategy is actually more effective at reaching the target. Image credit: Joel Lehman and Kenneth Stanley.

might be, and so you orient your efforts in that direction. They claim, instead, that even when searching for a specific solution to a specific problem, it is often more effective to conduct the search based on the novelty of solutions rather than on their suitability to the problem. Creative problems, in short, are like mazes, in which the path to the exit is almost never *actually* in the direction of the exit, and heading in that direction is at best no better than choosing any other direction and at worst counterproductive. That is by definition: we have never explored this territory before, so we could not possibly know its structure. This is one reason why a terrain metaphor is potentially misleading: in physical terrains you can generally (though not always) walk toward something and you'll get there.

A consequence of this tangled path to discovery is that the history of discovery is full of people solving problems they hadn't set out to solve.

And indeed it is common for innovation to work this way around: explore a space of possibilities, such as new materials or chemicals, and use your exploration to discover where applications lie. Designers and pharmaceutical companies engage in this kind of solution-first thinking quite systematically. The tale of the Post-It note is an example with mythical weight. Dr. Spencer Silver's goal was to make a new kind of strong glue. What he came up with was the Post-It note glue, which, despite his failure, he noted had interesting properties: it was low-stick, reusable, and pressure sensitive. For a long time, however, nobody could see a viable application. Silver's perseverance in seeking a "problem for his solution" eventually led to the collaboration with Arthur Fry in the development of the Post-It note.

It is amazing to think that something that is now so ubiquitous in offices and homes took longer than five years of puzzling to invent *after* the main technical challenge had been solved. In the field of design, where artistic and goal-directed technical factors play an equal role, such an approach to creativity is widely recognized and formalized in creative strategies. For example, in product design, one can approach concept ideation from a problem-first perspective, a scenario-first perspective, or a technology-first perspective. All are valid. In the latter two cases, we talk of "need finding."

Collective Creativity Is Greater Than the Sum of Its Parts

This already gives us good reason to understand human creativity primarily as a group process, but so conceived, this is still a view of group creativity occurring only as the sum of its individual parts, which are readily separable; in algorithmic terms this is parallelizable search. But there is good reason to believe that the creativity of human groups does not simply involve summing the creative efficacy of individual humans in a structured hierarchy. Groups structure creativity in various ways, leading to Csikszentmihalyi's appeal to correctly view creativity as a multi-actor dynamic process, and leading us to look at concepts of individual and collective agency.

In the simplest case, social processes can be seen to organize groups around specific goals, to which the creative power of distributed search can be applied. Consider, for illustration, the case of the solving of the problem of longitude, elegantly documented in Dava Sobel's popular history book *Longitude: The True Story of a Lone Genius Who Solved the Greatest Scientific Problem of His Time*.[20] John Harrison, a modest but highly skilled clockmaker, developed a series of clocks during his lifetime—throughout the eighteenth century—that made dramatic improvements to the accuracy of timekeepers overall and specifically in the face of adverse climatic conditions—changes in temperature, air pressure, and humidity. These improved clocks, which consisted of a large number of distinct innovations, meant that naval navigators could properly tell their location at sea, which had huge implications for the safety and efficiency of naval operations, and hence a staggering economic benefit and competitive advantage.

We can say with relative certainty that Harrison was a creative agent who transformed clock design with dramatic impact on the world. Yet even though in the book's subtitle Harrison is declared a lone genius, Sobel's account points to an alternative agency responsible for the achievement. The British government had in fact set up a prize for solving the problem of longitude. A tantalizing sum of money (£20,000) and a precise test were detailed in the award. As well as recognizing Harrison's agency, would it be equally fair to call the government a creative agent, establishing an intention (to solve longitude) and implementing a sort of meta-method (define the criteria for success and offer the reward)? The government's strategy was well adapted to the exploratory nature of creative discovery, exploiting a divergent search process, a little bit like the genetic algorithms we use today. Rather than betting on a solution (the problem was longitude—it was clear that better clocks could offer a winning solution but not the only possible solution), they displayed a purposeful ambivalence to what form the solution should take, focusing instead on devising the most accurate test of success, a strategy in keeping with the issues of creative problem-solving described above.[21] The government achieved its goal, demonstrating the efficacy of such a creative meta-method, now familiar in a range of market-minded devices at the heart of modern industry, such as patent law and crowdsourcing strategies. I will certainly not devote any time to laboring the argument over which was the true agent; the point is that they can both be seen to contribute to the creative outcome in their own right. But we can note some properties that distinguish these respective agents. Firstly, as mentioned, the British government was ambivalent to the form of the solution, whereas Harrison had all his eggs in one basket, so to speak. Secondly, the longitude prize exerted a force on individual actors, altering their choice of problem by incentivizing them. Along with his formidable intrinsic motivation, Harrison was further directed toward a specific task by the promise of the prize.

That different social configurations, such as governments that set prizes, embody different mechanisms of social creativity is lived out in today's cities, which compete to be powerhouses of creativity. For example, Charles Landry's influential book *The Creative City*[22] outlines how cities can be structured to stimulate knowledge transfer, serendipity, and evaluation, and urban planners talk of things like "serendipity management."[23] Richard Florida's work on the creative class[24] discusses how policies, such as those of increased inclusion, can support the nurturing of creative capacity. Fields like design thinking and organization science seek to untap the creativity of organized groups of individuals, pitching social creativity as greater than the sum of individual human creative acts, all because we can steer, stimulate, and structure individuals' behavior.

Although it is not without predecessors and peers, the work of Mihaly Csikszentmihalyi consolidated the social strand of creativity research. Csikszentmihalyi rejects the attribution of creativity strictly to specific individuals: "Where was Mendel's creativity? In his mind, in his experiments, or in the use his results were put to by later scientists? The answer, it seems to me, must be that it is to be found in all three."[25] In place, he outlines a systems-based view, the Domain Individual Field Interaction (DIFI) theory, comprising these three core elements. The domain is the area of activity in which the individual performs creative acts. Take abstract expressionist art, for example. The domain can be thought of as a passive repository of knowledge, concepts, and precedent work relating to the subject matter. The individual acts upon this domain, transforming it through some form of creative production. The field is the community surrounding that domain. This might involve other abstract expressionist artists or artists from the wider domain of visual art, critics, collectors, galleries, and the gallery-going public. The field's primary role in this conceptualization is to evaluate and select elements within the domain, but this is treated as equally active in the task of creation, and it may perform other roles as well.

Csikszentmihalyi's model frames specific networks of interaction: the individual must interact with the domain, becoming an expert in it before being able to thoroughly engage with it: "The information that will go into [an] idea existed long before the creative person arrived on the scene. It had been stored in the symbol system of the culture, in the customary practices, the language, the specific notation of the 'domain.'"[26] The field may influence the success of individuals and the acceptance of elements into the domain. An important part of the model is the collective effort that is involved in identifying and validating a discovery, regardless of who made it: "Not until Vespucci recognized that the so-called West Indies were part of an entirely different continent did Columbus's almost superhuman efforts get retrospectively revised as a 'discovery.' And if in the fullness of time it turns out that it was Erik the Red who really discovered America, that 'discovery' will be as much a result of scholarship and politics as a result of Erik's travels."[27] Historians of science have also illustrated how social evaluative interactions underlie those domains as well. Simon Schaffer presents an image of collective creativity that involves not only validation but also the mythologization of success:

> Much recent work on discovery and invention in the sciences demonstrates that retrospection and celebration play key roles in the production of discovery. Because discoveries acquire their status as the result of subsequent work with the relevant community, the "fetishism" of discovery is therefore the consequence of the whole process through which change is analyzed, debated, and assessed.[28]

He reiterates a view that "innovation is perhaps less a product of uniquely creative individual attitudes … and more a matter of being pivotally situated during those comparatively brief passages of industrial history when the balance of collective choice can be tipped one way or another."[29] One example is Edison, a prolific and widely acclaimed inventor, who can be seen to manage, promote, and negotiate as much as cogitate on the design of electric devices, channeling the work of many people and placing himself historically. Kekulé is also discussed, and his famous reveries are shown to have a degree of social construction to them, a collective retelling of the story of the problem and its discovery. Schaffer's view is that "discovery is as much a judgmental process in culture as it is the intrinsic property of some isolated mind." But he also maintains that we should not be disparaging of any such mythologizing but rather remain open to its potential importance, a driver for individual action. In the creative arts, he draws on the work of Ernst Kris and Otto Kurz,[30] who witness a similar construction of "culture heroes," with a common narrative of nurtured genius.

The effect of Csikszentmihalyi's model is to diffuse the obsession with individual creativity and provide a way to take account of the nontrivial social effects beyond the sum of individual creative actions. In a similar way, studies of creativity in design have shown how in collaboration individuals adapt toward each other forming new collaborative units of creative productivity[31] that involve skills like negotiation, including both persuasion and the ability to remain open. These may be manifest in different individuals' personalities in complementary ways.

But this is a challenging proposition. Once we start to view creativity as something that happens in groups, some assumptions about what a creative agent is become complicated. What Csikszentmihalyi refers to as "the generative force of the field" is a different type of creativity from that of the individual, and in the next section I consider how this approach is part of a wider movement in the social sciences to truly confront the dualism mentioned at the beginning of this chapter. Such a distributed view has real consequences for how a creative computational system might be understood. In Csikszentmihalyi's model, not only must lone humans be understood as components in creative systems, rather than creative systems in their own rights, but they must also be understood as performing different roles within the larger creative system. Some roles may be minor, and not all roles require the sophisticated creative machinery required by the kinds of ideal creative systems imagined in what he calls the Ptolemaic view. As such, a computer program might be tasked with performing any one of these roles, and this needn't require advanced AI. The work of computational creativity practitioners like Cohen and Cope illustrated how farming out content generation to a machine can produce effective and surprising results, even if the system is being highly mechanistic. In other circumstances, such as in building analytics and other goal-directed design tasks, computers can be very good at making judgments by which to evaluate things.

Questioning Agency

Parallel to Csikszentmihalyi's model is a wider effort, crossing academic disciplines, to address the dualism that leaves us stuck with the idea that individual humans can exhibit genuine agency while other systems, including *collectives of individuals*, cannot. This mission of addressing asymmetrical dualism in the social sciences was introduced at the beginning of this chapter through the thoughts of Andrew Pickering and Alfred Gell. The basic gist of this movement is to rethink agency, the capacity for action in the world, as something that can be manifest in more temporary and loosely bounded organizations such as classes, companies, movements, and so on.

When I discussed the British government just now as such an agent, it was relatively easy to picture this organization as if it were a single individual with a coherent goal. But a complete account of any social organization never just amounts to a group of people. Instead it must be fully understood through nonhuman elements that give it its capabilities—books, rooms, rules, maps, machines, and so on. There is also not likely to be a common goal among the people (and perhaps the objects); in Sobel's account of the longitude prize, the drama takes the form of numerous conflicts among entities both within and outside the body we have just identified as an agent.

Such questions have for some time now preoccupied social scientists and philosophers. Most famous of the proponents of a revision of the application of agency are the actor-network theorists. In a well-known passage,[32] actor-network theory figurehead Bruno Latour narrates the central concern of this movement through the example of the US gun control debate. "Guns kill people!" gun control advocates say, to their opponents' riposte, "People kill people, not guns." For Latour this is the false dichotomy between a materialist and a sociological account. Instead, Latour proposes that "a third agent emerges from a fission of the other two. … Who, then, is the actor in my vignette? *Someone else* (a citizen-gun, a gun-citizen). … It is neither people nor guns that kill. Responsibility for action must be shared."

Thus Latour sets up a resolution in the idea that the gun-citizen is a vague, temporary actor who is capable of actions (including forming realizable intentions) which are not accessible to either entity on its own. That is, networks of interacting entities temporarily merge to form sites of possible action, at which point the component entities are no longer the same as they were before. He calls such changes "translations." The resolution sounds like a sleight of hand: critics have complained that *it was only the person* who could have formed the intention, chosen to acquire the gun, pulled the trigger. And yet that person's path through life was in fact a constant interaction with a world of influences, changing course at every stage based on what is deemed possible, what is to hand. Latour's main contention, then, is that "the twin mistake of the materialists and the sociologists is to start with essences, those of subjects or those of objects," and that this is doomed to failure because "neither subject nor object (nor their goals) is fixed."

This analysis has proven particularly powerful in the discussion of the contemporary scientific method, where Latour's main work was focused, pointing to "a new actor … inert bodies, incapable of will and bias, but capable of showing, signing, writing, and scribbling on laboratory instruments before trustworthy witnesses. These nonhumans, lacking souls but endowed with meaning, are even more reliable than ordinary mortals, to whom will is attributed but who lack the capacity to indicate phenomena in a reliable way. … In case of doubt, humans are better off appealing to nonhumans."[33] This allows a framework in which both the actions of scientists and the action of the objects they study come together with a shared responsibility for scientific outcomes, something that a purely sociological account struggles with.

A consequence of this revised view of agency is to replace human intentions with real tangible outcomes as the fundamental marker of agency; everything that happens in the world is given equal attribution to an agency that caused it. If one building falls down because of an earthquake, while another is demolished by developers, it does not allow that in one case agency was involved whereas in the other case it was not.

In a similar vein, Lambros Malafouris investigates "material agency" through the description of a potter at a wheel, creating a pot from clay. The potter's hands are in a direct and dynamic interaction with the clay. What happens when you indulge in a detailed analysis of this interaction is that causality reduces to our physics-eye view, the view alluded to by Turing in his treatment of the Lovelace objection.[34] As the potter, wheel, and clay interact there is no way to describe the potter as originating the action. Neurons are firing, but they are being influenced by the visual and tactile senses, responding to the shape and behavior of the clay. There is no clear beginning to the process and we cannot identify the process of origination that is supposedly the possession of the potter and not the clay or the wheel. Malafouris concludes, "If human agency *is*, then material agency *is*, there is no way that human and material agency can be disentangled. Or else, *while agency and intentionality may not be properties of things, they are not properties of humans either: they are the properties of material engagement, that is, of the grey zone where brain, body and culture conflate*." (original emphasis)[35]

This conclusion is drawn from zooming in very close to a specific act of embodied creation to reveal the impossibility of true origination. The same can occur by zooming right out to witness longer-term cultural effects. The potter may have intended to produce a pot of a certain style, but all of this may well be scripted (Latour's term) by the potter's cultural precedents and demands—even the intention came from elsewhere. The potter, with his motivations, norms, aesthetics, and cognitive equipment, is the product of millennia of cultural evolution and hundreds of thousands of years of human biological evolution, as well as the things that were said to him or that he saw on the news that morning.

This is a potted summary of a vast topic. Variations of this essential idea can be found cutting across many different strands of late twentieth- and early twenty-first-century thinking. Related thinking can be found in the extended-mind hypothesis of Andy Clark,[36] which views technological artifacts as *aspects* of our cognitive capability: from eyeglasses for improved perception to notepads that allow us to iterate visual ideas or offload information from short-term memory, to pocket calculators that perform other magical feats of cognition.

Likewise, in design it has proven fruitful to understand humans as users of objects in this way. Relevant strands of thought include distributed cognition and situated cognition, and a bundle of practices that have become known as third-wave human-computer interaction (HCI). Susanne Bødker notes in third-wave HCI an increased attention paid by designers to the "webs-of-technology" or "artefact ecosystems" in which human action takes place: "We never design single, monolithic devices or systems, but technology that must be seen and used in relation to many other devices, applications and systems. Webs-of-technology are used to describe ubiquitous interaction as a process of negotiation between the users and the technology, focussing on the availability of technology and interpretability of services."[37] It is understood that "users' shared capacities and experiences are not just based on individual acting and learning in the world; rather, they are bound to shared practices, joint activities, and so on. … It is against this background that the relationship between the user and the artifact exists." Another design-focused anthropologist, Paul Dourish, argues in a similar way to Malafouris that "embodiment is not a property of systems, technologies, or artifacts; it is a property of interaction. … In contrast to Cartesian approaches that separate mind from body and thought from action, embodied interaction emphasizes their duality."[38]

For designers, social theorists, and historians, the challenge of constructing a workable framework for describing and designing for systems that do not privilege human actors is a monumental task. Despite the huge body of theoretical work dedicated to this effort, only lightly touched on above, a clear and universally accepted picture is still lacking. I will draw two key points from the above summary.

The first is that it can be effective—and many would say it is essential—to view creativity as occurring through interaction, and for the agency responsible for any creative act to be defined by a shifting boundary, perhaps drawn around individuals, groups of individuals, individuals interacting with materials and machines, and so on.

In our daily lives, it is practical to think about humans as originators of things, or more generally as causes of external effects, to think of ourselves as fully autonomous agents. We affect the world in ways that we have generally planned. Even if you agree with Turing in principle, in your everyday life you don't give up claiming credit for coming up with a good idea or cracking a fine joke. Rightly so; as a system you are capable of behaviors that are irreducible to their lower-level components. Only in the sophisticated configuration that is *you* could that joke have come into the world.

But despite our prototypical tendency to point to ourselves, as well as to great individuals of art and science, as creative exemplars, we should always do so with caveats, scare quotes, and question marks. Likewise, asking if a computer program is creative will always require a situated view that takes into account that system's interaction with the world; the preferred question is whether and how any computer program *contributes* to creative outcomes.

The second point is that, as Latour frequently reiterates, as different entities collide in the formation of temporary agents, it is not only their capacities that change. Everything is in flux, including *goals*, which is the reason we can't just describe outcomes in terms of individual intention. This is important to the idea established in the previous chapter that a creative act is one in which something of value is created. The question that remains is, when we ascribe value to something we have to address the more complex problem of "valuable for whom or what agent?"

The Lovelace Objection Revisited

These ideas lend a new perspective on Lovelace's question of origination. In his response to Lovelace's objection, Turing's view bears a resemblance to Malafouris's challenge to human agency. He reminds us that in a classical view of systems, everything is at some low level just following a cold, hard path of determinism—just simple billiard balls smashing into each other and bouncing around, a causally uncomplicated state of affairs. To trace the cause of anything, run back the tape and see where the billiard balls went.[39] For Turing, as we have seen, the idea that we are originators of things is no less dubious than the idea that computers could one day be.

Somewhat harder to stomach, Turing also argues that machines routinely produce surprising behaviors, taking an indifferent attitude to whether such surprises are simply due to errors—accidents that result in behaviors that aren't there by design (being errors, you can say they aren't really there for any reason whatsoever). Bringsjord, Bello, and Ferrucci proposed one specific hypothetical example of such a machine error in their response to Turing, which was discussed briefly in the previous chapter.[40] They imagine that, due to a hardware error, a production-line robot sticks a spare tire on a car in the wrong place, and as a result a designer is inspired to produce a new car design based on what the machine has accidentally initiated. For Bringsjord, Bello, and Ferrucci, this is a clear-cut case of the machine *not* being creative: "Would we want to credit the malfunctioning robot with having *originated* a new auto?" they ask. "Of course not," they reply, but why not?

It is understandable to reject the errant robot's creativity a priori, being so dysfunctional in its genesis, but applying criticism derived from Latour, Malafouris, and company, we might question if this prejudges that one part of a holistic creative process is more important than another. From a standard perspective on creativity, the errant robot is not creative because it did not know what it was doing (and like Bringsjord and colleagues, applying our creativity prototype detectors, we would be inclined to agree). But on the other hand, there is no way of explaining the particular outcome that occurred without reference to the particular actions of the errant robot. Importantly, as Simonton says of artists like Michelangelo—*fundamentally* creative in the sense that they make things that wouldn't have otherwise been made—it may well be that had the errant robot not done what it did, such a car design would never have seen the light of day. Perhaps the new car design sparks a butterfly effect chain of radical change in the use of cars and the design of cities. Now, imagine that, so pleased are they with their accidental discovery, the designers decide that they need *more* errant robots, to stimulate more accidental discoveries. They build factories with random variations added to the production-line robots, which now routinely perform blind variation. Then we have what might look a bit like an act of Darwinian evolutionary creation: a "random mutation" in the robot causes its "genes" future success. As Csikszentmihalyi notes, since the full gamut of outcomes is not prespecified at the moment of such a potentially creative act, it will be necessary to revisit the attribution of creativity later on anyway.

This could be taken as a profound dilemma. But it is a tedious trap to be stuck in. There is no need to quibble over the creativity or not of an errant robot, only to recognize that multiple heterogeneous actors are involved in pretty much any creative process. To be clear, I certainly do not wish to claim that the robot is creative, only that it is a part of the temporary entity through which (and only through which) the new car design is created. Likewise, this is not a call to dismiss the power of human intelligence in the creativity of distributed creative systems. Human brains are systems that possess great creative capacity. Without them, airplanes and symphonies would not exist, and neither would the new car design. Instead, we need a different way to think about creativity, to construe it as an entirely distributed process.

Analyzing Interactions between People and Things: Functions and Affordances

This analysis leads to a somewhat negative stance in which default assumptions are broken down but little is left to grab on to. But as the product of this thinking, various concepts are available for describing the relationship between components within distributed creative systems in general and specifically, for us, in relation to their creativity.

Of greatest importance to these relational questions in human-artifact interaction is a cluster of concepts associated with the notion of *function*. The act of the errant machine discussed earlier was seen as creatively problematic because of its lack of any creative objective, whether designed or emergent. While the designer comes with a complex of function-oriented properties—goals, motivations, objectives, intentions—the errant robot is literally *malfunctioning*; it has not performed whatever function we might attribute to the robot by its design. It has not devised new goals or evolved new behaviors, it is simply broken. We find it nearly impossible to imagine that something under these circumstances could be an agent in the creative process, let alone the originator of anything, precisely because of this failure of function.

But on closer inspection we find that a rejection of creativity on the basis that it does not arise from a clear function would taint many acts of creation, above all in seemingly agentless processes such as emergence, evolution, or autopoiesis; these fulfil the concept of *generative creativity* that can work even when there is no objective or preemptive function. This is particularly relevant in creative arts where any discussion of the utility of the product is inevitably fuzzier than in contexts such as design and science, as I will discuss in the next chapter.

In discussing agency and interaction in different systems comprising humans, human-made artifacts, and even other natural biotic and abiotic elements, we have encountered several different types of functions. Biological agents exist and act in the context of an evolutionary logic. Often it is valuable to understand their actions and traits in terms of how they serve the survival of their genes, and this generally describes the function of biological traits. The predator has front-facing vision to focus in on its prey, whereas the prey has wraparound vision to keep an eye out for predators. Both are good for that particular species' particular niche and its associated needs.

Meanwhile, we can talk about human-made artifacts having functions ascribed to them by design (the hammer is for driving nails into surfaces) or by some creative appropriation (the hammer, by virtue of it being small and heavy, makes a good paperweight). Aristotle distinguished final and efficient causes—the purpose for something existing versus the immediate sequence of events leading to its existence. But final causes are problematic in the context of natural evolution: evolved things don't have purpose, as such, even though they appear to.[41] These are all functions, but it is valuable to distinguish them, as Dan Sperber does,[42] following Ruth Millikan, using the separate term *teleofunction*. Teleofunction refers to the kind of function we see in biologically evolved systems, a function that arose and serves only to continue its own existence. The function of a bird's wings is to fly, which has arisen and exists because of its success at continuing the phenomenon of flying birds (not "in order to" but by continuation of what it has done with success previously). Importantly, however, teleofunction can be found in artifacts and social practices too, which may arise because they are useful to us, but may equally arise not because we need or want them but because they *exploit us*. This is perhaps best known in the concept of the meme (this term is discussed below, but like Sperber I dislike the confusion this term brings).

Also, Sperber notes, we actually live among strange hybrids of biological and human-made, functional and teleofunctional artifacts. Sperber suggests examples of these such as seedless grapes; they are of biological origin but adapted through interaction with humans to exchange one biological teleofunction (seeds) with another (being tasty to humans), and henceforth reproduce through the human act of grafting. Other cultural phenomena can be broken down accordingly. They can also be analyzed in terms of who is being served in the functional interactions taking place. Compare the goals of a company that produces an unhealthy but addictive soft drink product and the goals of one of its customers. The drink and associated advertising materials form a nexus of artifacts that these agents interact with.

These various considerations begin to shape an effective toolkit for thinking about artifacts. From another direction, James Gibson's concept of affordances is a powerful tool for thinking about how the world around us might more specifically guide our actions. Although better known nowadays as a theory concerning artifacts, Gibson developed the theory of affordances to understand behavior in evolved ecological contexts. "The verb *to afford* is in the dictionary," Gibson explained in 1977, "the noun *affordance* is not. I have made it up. I mean by it something that refers both to the environment and the animal in a way that no existing term does. It implies the complementarity of the animal and the environment."[43] Affordances are the interactional relations of potential between an animal and its environment. A water bug can skate over the surface of water. A lion cannot. This of course influences how water bugs and lions approach a pool of water. To the water bug, water affords skating.

The concept of affordance is beautifully obvious! It is a marvel that the simple derivation of a noun from a verb could have had such a remarkable effect on our thinking to be considered one of the great intellectual contributions of the last century. Affordances help us deal with the most mundane and hidden of interactions. Gibson refers to the ground as "stand-on-able." Something's possible affordances can form a lengthy and perhaps very abstract list: "Air affords breathing, more exactly, respiration. It also affords unimpeded locomotion relative to the ground, which affords support. When illuminated and fog-free, it affords visual perception. It also affords the perception of vibratory events by means of sound fields and the perception of volatile sources by means of odor fields."[44]

What the theory of affordances does, as applied in design, is help to enmesh this two-way interaction between things and people, indeed between all types of entities. Gibson's world consists of complex layers of species interactions, negotiated through affordances and growing in complexity through evolution's exploitation of these affordances. Things get interesting when the complex cognitive abilities of one species actually become an affordance for another species (*bee eyesight affords attracting bees using visual stimuli, as performed by flowers*). This is interesting because it is not always the more cognitively sophisticated species that is doing the exploitation; cognitive complexity can be hijacked, which happens in many instances where a plant or fungus uses an insect or animal to spread its seeds or spores, or with any sort of biomimicry. These examples, where the manipulator is not a cognitive entity, while the manipulated entity is, provide another example of why we should admit a more fluid theory of agencies.

Likewise, in the world of human artifacts, many affordances are straightforward. A skateboard affords rolling down a hill. In the Gibsonian sense, this means it *can be done* (whether you recognize it or not—if you don't, it is simply a hidden affordance). Donald Norman's translation of the theory of affordances into the world of design[45] then poses issues of how we can manipulate and establish affordances in our built environment—not just what the affordances of something are but how you present the affordances of an object to a user in a way that improves usability. Many poorly designed things, Norman observed, fail for this reason—there is some deficiency in the *perceived* affordances.[46]

But there are also situations where an affordance needn't be perceived at all for the artifact to function. Consider the difference between two types of electric hand dryer: one is button operated while the other has a proximity sensor, which is designed to detect when something is placed under the dryer. These are near identical systems. The proximity sensor is a binary switch which is, from the point of view of the hand dryer's electrics, functionally identical to the button. But whereas in one case there is a button that affords pushing—you are clearly the agent operating the machine—in the other case there is a "smart" hand dryer that senses when a pair of hands is waiting to be dried. In the latter case, you are really just pushing an invisible button (everyone has experienced waving their hands around in front of such a device failing to activate it, wishing they could *perceive* the button). Indeed, the proximity sensing hand dryer is at the somewhat humble end of smart technologies. But the point remains that the way we conceive of the relationship between participants in this interaction can shift quite some way with only a subtle shift in the basic conditions.

The theory of affordances provides a low-level description of how our agency relations with artifacts exist as two-way exchanges. The affordance of a designed artifact is at once something that a user actively *chooses* to act upon and something that influences their behavior by indicating various action possibilities. Indeed, along with usability comes the broader concept of user experience. This might include seductive or manipulative features of an object's design that are not functionally beneficial to the user. In other cases, the affordances are put to the service of someone who is not the immediate user. Adding a heavy weight to a hotel room key is a well-known strategy for hotel managers to reduce the instance of lost keys (an example discussed by Latour). Thus, at a stretch, weighted hotel keys afford for the hotel staff "being left at reception," but they are a nuisance to the hotel's customers (though perhaps less of a nuisance than a hefty fine).

Affordances form a central pillar for Vlad Glăveanu, a contemporary proponent of an inclusive distributed approach to creativity grounded in social psychology. For Glăveanu, "creative phenomena cannot be properly understood and studied outside of their integrated and dynamic totality defined as a creativity complex."[47] Replacing Rhodes's four Ps of creativity—product, person, process, press—he offers his own five As—actor, audience, action, artifact, affordances—that constitute the units of analysis of a "molecular" perspective, with affordances providing the relational component. Glăveanu is at the same time keen to stress the forced nature of such a framework: "to analyse creative activity means, in essence, to engage in a form of 'violence' towards the lived experience of creators, to segment the wholeness of a creative experience into parts and units that might, in the end, be foreign to it."

For Latour, the interaction between the larger network structure and the individual is also seen as one in which more or less active or passive relations may play out. He uses the term *script* to refer to the way in which a situation may dictate what any individual might do at a given time. The individual can either stand under the script (be passive in the unfolding of events) or stand above the script (enact a form of agency, redirecting events).[48]

Alfred Gell[49] applies similar thinking about agency to cross-cultural art production, and in doing so offers what I think is a slightly more palatable terminology that aligns with a modern view of agency but offers a handle on the classical view at the same time. *Primary agents*, he says, are those "self-sufficient" agents that have all the stuff of intentions and cognitive power.[50] He also wryly refers to these as "pukka" agents—in other words, the real deal—satisfying the classical view of human agency. *Secondary agents*, then, are artifacts that don't really do anything on their own but can *act* as agents in various ways, meaning simply that they have an effect on their environment. They might act as delegates for the agency of their maker—as the weighted keys of the hotel management have an effect on the behavior of hotel guests. In Georgina Born's interpretation, "the artifact mediates and relays social relations across space and time."[51] But they could equally have effects that were not intended, when something takes on a *life of its own*, such as whenever something malfunctions.

Affordances, scripts, and secondary agency can all be seen as ways of capturing how individuals are subjugated under complex distributed creative systems, and how artifacts have the potential to play a part in that subjugation. Having worked through concepts of agency, autonomy, and function across different contexts, we have a better chance of breaking down and describing computational creativity systems without getting stuck in asymmetrical dualism: believing that human individuals form distinct and privileged sites of agency. The value of these terms is that they are not specific to computationally creative systems and so help establish a continuity. The breakdown of the binary distinction between humans-as-agents and artifacts-as-non-agents carries over to the *timeline* of emerging computationally creative systems,[52] which can instead be understood as systems that can perform in different ways in the formation of temporary distributed creative agencies. For example, we have already looked at the artist, Cohen, and the system, AARON, in terms of a coupled interaction that defines a new collaborative agency, mediated by the various affordances of the technology and the artistic domain, and involving a nontrivial set of interacting teleofunctions alongside more traditional imposed functions. None of these concepts really shifts our technical understanding of what AARON is doing under the hood, but they provide a more structured view of how the system, despite being "merely generative" and exploratory, feeds creative outcomes. The terminology now introduced means we can move beyond asking whether or not AARON (or an errant robot) is a creative agent and begin to build a rich profile of the system in terms of its relationship to agency and function.

Bringing this back to the distinction between generative and adaptive creativity, very many processes, consisting of different human and nonhuman elements, can perform generative creativity. Computational systems such as AARON are very easily accepted as performing creative functions insofar as they are understood as producers of constrained variations: generators of unscrutinized novelty operated by a human actor. Fewer systems have the more narrow properties of adaptive creativity. Humans can be seen as such systems, with the machinery not only to generate novelty but evaluate it and tailor this process of generation to specific needs. In the above I have started to outline how, while much of human creativity is adaptive creativity, there is also a lot of creativity that is generative and accordingly deserves to be evaluated by different criteria. Individual humans compose melodies and paint paintings, but they do not create movements or perform the work, on their own, of turning their artistic outputs into symbols of cultural significance. This is done by large groupings of people working in distributed and less coordinated ways. Such groupings in the world of art are largely less purposeful than the British government setting the longitude prize. They enact generative creativity because they do not act with such common, unified purpose; what comes out of them cannot be said to have been created with a pre-established objective or common value.

What Created the Creative Species? An Evolutionary View

This leads to a further detour to consider the generative, emergent, purposeless nature of this cultural level of activity, where cultural processes are understood not only in terms of human agents being adaptively creative but at the same time via the lens of Darwinian evolutionary theory, which is itself purely generative. Culture embodies different aspects of these distinct processes, and here I look toward the more mechanical, systems-based perspectives that fill out this framework.

The blind variation approach to creative search, discussed at the beginning of this chapter, and the notion of teleofunction encountered later, both relate unmistakably to Darwinian evolutionary theory.[53] Besides these associations, bringing Darwinian evolution into the discussion may seem like an arbitrary addition. But consider that Darwin's theory provides the link between two of the domains that have been central to the discussion of origination,[54] between blind mechanical physical processes and primary agents. It binds the deterministic view of the world, appealed to by Turing and others, to our everyday experience of intentions and will. In this sense it is a good way to frame the understanding of function necessitated by the previous discussion.

By Darwinian, do not think of the straw man caricature of a narrow adaptationist view, epitomized by the reduction of every biological trait to one adapted gene or another. Think instead of a rich and pluralist array of evolutionary phenomena, which has natural selection at its core, fleshed out by a vast array of other hypothesized evolutionary effects, including but not limited to: sexual selection, group selection, niche construction, the Red Queen effect, kin selection, ontogenetic and phylogenetic interaction, gene-culture coevolution, the Baldwin effect, genetic drift, symbiosis, and punctuated equilibrium. It should be accorded the generality found in the dictum attributed to D'Arcy Thompson that "everything is the way it is because it got that way."[55] Darwin's own research wasn't just about showing that natural selection worked as a mechanism, but constructing complex and very specific historical narratives to explain speciation events and adaptive pressures: could such and such a seed have floated on water or been carried by a bird in order to get from one island to another? His world involved many different interacting forces, with a great deal of happenstance in the emergence of evolutionary outcomes.

If creativity, like agency, should be defined by effects rather than presupposed intentions, then Darwinian evolution is a quintessentially creative force, and it is also one that our prototypical creativity detectors might grant creativity. It has generally been a challenge to reconcile evolutionary creativity with human creativity, because one is intentional and cognitive and the other is absolutely not. One is conceived of in terms of its value to us, while the other is seemingly value-ambiguous, and because the latter is simply more familiar (prototypical) than the former. The distributed creativity view developed in this chapter, by contrast, is inclusive of evolutionary emergence.

The theory of evolution by natural selection is based on a simple naturally occurring stochastic mechanical process that has three components: heredity, variation, and selective retention. To begin, we require replicators, structures that can copy themselves. The theory can't explain the emergence of replication, because it depends on its preexistence. But from the existence of simple replicating molecular structures, everything else follows: with replication comes heredity (the copy is similar to the source) as well as likely sources of variation (a simple copying error may introduce differences between the source and the copy). It then also follows that we will witness selective retention; with any accidental difference between two individuals, we may find that the difference enables one individual to have better chances of copying itself than the other. This fitter one's progeny will, by this definition, become more common and increase the occurrence of whatever trait led to the improved fitness. Furthermore, even the simplest replicating molecules consume some kind of locally finite resource, at least the raw chemical ingredients needed to construct them. When more or less successful individuals are forced to compete for these resources, the less successful individuals will be outcompeted, potentially to extinction. Thus from just the existence of self-replicating molecules, an extraordinary creative force arises, which Darwin used to explain nothing less than the origin of species.

The theory of natural selection emerged at a time when we were beginning to see the world as the product of quite radical interactions between micro- and macro-level processes. Darwin followed in the footsteps of Charles Lyell, whose seminal studies of geology[56] first portrayed the earth as a dynamic system, its volcanoes and earthquakes the signs of gradual processes that over time shaped mountains and seas. Likewise, Darwin's theory shows how the minute daily interactions of biological forms ultimately influence the forms themselves, given long enough.

An immediate consequence of an evolutionary perspective is that life on Earth is revealed to have followed a path of increasing complexity, marked by certain major innovations in the design of biological forms. John Maynard Smith and Eörs Szathmáry chart these innovations in a treatise on the "major transitions in evolution."[57] These include the transition from simple replicators to chromosomes, from asexual clones to sexual populations, from solitary individuals to colonies with nonreproductive castes, and the emergence of culture.

As with the success of individual species, each of these innovations was a winning design, definitively successful and never to be undone once discovered. We might add various other minor innovations such as skin, eyes, flight, and locomotion, that may have been lost at various moments to certain species but once introduced into the world have flourished and resisted eradication. But the major transitions are particularly special in that they each, in effect, change the rules of evolution. Sexual reproduction introduces the idea of a common gene pool and genetic recombination. This allows the diffusion of the best traits through an evolving population, making evolution more effective and adaptive, and also leading to completely new forms such as the baroque excesses of sexual display, principles of exogamy and endogamy, and complex social structures.

More recently, researchers in fields such as complexity science and artificial life have sought to understand more abstract principles by which blind processes may produce high complexity. Stuart Kauffman, most famously, has developed a general theory of the emergence of order,[58] based on studies of self-organizing systems, of which natural selection may be understood as a specific case. In parallel, the biologist-philosopher team of Humberto Maturana and Francisco Varela popularized the notion of an "autopoietic system," literally self-creating, which is defined and reinforced by the maintenance of an organization through a dynamic process.[59] Biological cells are exemplary autopoietic systems, involving a metabolic process that maintains a membrane, separating the cell from the outside world. By analogy, multicellular systems and other configurations such as social groups can be described in autopoietic terms, as was attempted by Niklas Luhmann in a detailed application of Maturana and Varela's work to the social sciences, particularly the domains of artistic production.[60] Autopoiesis not only describes the immediate maintenance of a system's structure but the long-term emergence of that organization through a process of structural coupling between elements. In the social domain, it provides a mechanism to think of something as being completely function-free yet complex and actively self-sustaining, sitting at the boundary between a process of emergence and one of formation by natural selection.

Cultural Evolution

Following the trend of increasing complexity, human culture rewrites the rules of evolution in yet another radical way. Culture is the product of human populations accumulating knowledge across individual boundaries. Any individual organism with the cognitive capacity to learn about its environment can accumulate knowledge. An animal will learn its territory, learn to recognize other animals, work out what is good or bad to eat and where to find it, and can also be creative, working out how to do things that it did not know how to do innately, like appropriating an object as a tool.

This body of knowledge is generally built up over a lifetime and lost when that animal dies. But it is not always lost; many species learn from each other in a variety of more or less sophisticated ways. A young bear first learns its way around a territory with the simple strategy "stay close to mother." This is an example of the vertical transmission of knowledge—vertical meaning that it follows the genetic transmission from parents to offspring. Social learning, here, is a crude blend combining a social factor without a learning component (stay close to mother), and a learning factor without a social component (build a map of the environment).

Such simple social learning can be horizontal or oblique as well—that is, from members of the population who are not our direct ancestors. Blue tits in Britain have been shown to learn from each other that they could peck into milk bottle tops to drink the milk. But, as with bear cubs, all it takes is for the tits to follow each other around and to randomly peck a lot at things around them. The learning still consists of an individual discovery; each individual blue tit, engaged in blind search, chances upon the milk bottle and discovers the reward. But this is all facilitated by the "teacher" bird who, by hanging around near milk bottles, constrains the blind search to set up the conditions for the learning event to occur. In this scenario, neither bird needs to have any intention to teach or to learn, nor even any kind of model of what the other bird is doing or thinking.

In humans, social learning has evolved into a dedicated and sophisticated set of cognitive capacities, including thinking about what another is thinking—their goals, intentions, and understanding of a situation; understanding another's actions in relation to our own; learning complex sequences, including making things; and of course the use of language and symbolic culture to create the ultimate tools of social transmission: shared environments of the mind. Social learning is a superb strategy for gaining knowledge. The useful knowledge that individuals work out for themselves is generally learned at a cost. Exploring a new territory or tasting strange foods involves taking risks that would be better avoided, and trial and error is, if nothing else, time consuming. So social learning makes sense for individuals; if you see someone else eating a strange food and not getting sick, then you have evidence that it is safe. These benefits, many have argued, provide a selection pressure that has driven the evolution of complex social learning machinery in humans.

Thinking about what someone else is thinking, known as applying a *theory of mind* (ToM), is one important ingredient for human social learning. The neurophysiological underpinnings of ToM include circuitry that connects seeing an action performed and performing that action oneself. The same circuits are active when you see someone grasping something, as when you grasp something. Such simple mechanisms are hypothesized to underlie the transition in humans toward collective modes of action, by enabling a form of shared cognition that ultimately underpins language and culture.

Along with the mechanisms that enable learning from other individuals come mechanisms for steering learning in the right direction. In a species that learns socially, a new demand arises: to sift through all of the people, and all of the things they do, and work out what is worth learning and what is not. A sensible starting point for this is to learn to do everything that our parents do. Vertical transmission, from parents to offspring, is a relatively safe bet because it perfectly tracks the genetic transfer of traits. If it works for your parents, then it will probably work for you, besides which, your parents are likely to be readily available and helpful, being interested in your success and willing to support your learning.

Unlike genetic inheritance, vertical cultural transmission speeds up the rate at which a given lineage can adapt, because it can add in the developments made through individual learning. This mirrors Lamarck's theory of evolution; Jean-Baptiste Lamarck, before Darwin, posited that evolving species could pass on learned traits to their offspring through biological reproduction. He was wrong when it comes to the genetic transmission mechanism, but this particular model of evolution still has a place in our wider understanding of evolutionary theory.[61]

Social learning can be even more powerful if the transfer of knowledge occurs horizontally or obliquely, not just vertically. For this, Robert Boyd and Peter Richerson[62] propose heuristics for the successful social learner, such as to focus on behaviors that are very common in the population. It is reasonable to assume that a common behavior is a successful behavior; if there are many people around adopting that behavior, then at worst it is not fatal, and at best, those people have adopted it because it improves their lot. This tendency to copy common behaviors is known as "frequency bias" and can be stated simply as the rule that the more common a trait is the more likely we will be to adopt it. Another strategy is to focus on learning primarily from successful individuals, whose success, we might presume, may be down to the success of their behaviors. This accords with the fact that celebrities are paid large sums of money to advertise products.

The powerful cognitive machinery of social learning enables culture and technology, which, like biological evolution, has a cumulative nature. Once an effective innovation has been made, its ability to reproduce, whether via biological reproduction or cultural transmission, means that it is likely to stick around. Each innovation also provides the basis for future innovations to be built upon it. Some innovations, like writing or telephones, even speed up and bolster the copying and accumulating processes itself.

This cumulative process has been described as a "ratchet" effect,[63] a ratchet being a device that allows movement in one direction but not the other. On this basis it seems reasonable to suggest that the human mechanisms underpinning social learning, not creative discovery, are the real powerhouse of human creativity, because they underlie the working of the ratchet. Without it, we would not accumulate knowledge, and each individual human life would consist of solving the same problems. By analogy with biological evolution, we see that as long as the ratchet is working, then even just random thrashing about will do to drive progress: every so often we advance a notch on the ratchet. Thus even human creativity may not actually require that there is a particularly powerful individual human capacity for creative thought.

This is not to say that cumulative culture is just another manifestation of Darwinian evolution. There are competing views of how cultural evolution or cultural replication relates to natural selection. For example, Richard Dawkins[64] proposed that if the gene is the unit of biological reproduction, then we could talk of an equivalent unit of cultural reproduction, for which he coined the term *meme*. While the term became popular for a while as an academic subject of study, others reject such terminology because there is no apparent need to talk in terms of units in the context of culture, and little more than confusion to be found in hypothesizing that some cultural thing akin to a gene actually exists and performs any such function. Why bias what the nature of cultural evolution might be with such a leading metaphor? Nevertheless, whether you use the term meme or some other framework for thinking about cultural evolution, the application of the ratchet effect concept to human culture is a common thread and has the essential features of a blind search process.

Some have suggested that biological and cultural evolution form a continuum of accumulating technologies. Kevin Kelly, for example, dubs the entire system the *Technium*,[65] and argues that its evolution toward increasing complexity is a far greater force than anything we may do to exert our will upon it.

For Brian Arthur, an economist responsible for the theory of increasing returns, technologies beget other technologies.[66] Specifically, "nature possesses many sets of phenomena. … We can say that novel phenomena provide new technologies that uncover novel phenomena; or that novel technologies uncover new phenomena that lead to further technologies. Either way, the collectives of technology and of known phenomena advance in tandem."[67] As for Kelly, whose book on the subject is entitled *What Technology Wants*,[68] technological determinism supersedes human agency, and human ingenuity is merely a part of the mechanism by which this grand holistic process proceeds. As with Ashby's portrayal of the brain as just another technology invented by nature, this is the hardest of hard technological determinism. There are good cases against taking such a deterministic stance, but also strong evidence that we are set on a course of continuing complexification witnessed in the evolutionary major transitions.

Arthur observes that unlike in biological evolution, technology is capable of combinatorial evolution. Plants and animals can't hybridize, except within the limited scope of sexual reproduction operating within a species (and the tricky concept of horizontal gene transfer). But in technological combinatorial evolution, subsystems can be replaced and rearranged, and a solution from one domain can be brought to bear on another domain. This means that technological evolution experiences increasing returns: the more technology is out there, the more ways there are for elements to be recombined and lead to new innovations, and so the more quickly technology evolves.

Arthur also specifies four separate mechanisms of technological evolution:

> Innovation consists in novel solutions being arrived at in standard engineering—the thousands of small advancements and fixes that cumulate to move practice forward. It consists in radically novel technologies being brought into being by the process of invention. It consists in these novel technologies developing by changing their internal parts or adding to them in the process of structural deepening. And it consists in whole bodies of technology emerging, building out over time, and creatively transforming the industries that encounter them.[69]

In short, if biological evolution is a powerful creative force, then cultural evolution combines this with additional power in the form of human agents' problem solving, planning, adopting, and so on. Nevertheless, these various technologically determinist views strictly place the power of emergence over that of planning and cognitive creative agency; essentially the same creative force found in nature, purely generative, is seen here to drive technological change, just with human brains as the conduit through which the emergence takes place. Artistic processes can be treated in the same way, with the individual subsumed by the system.

Niche Construction

The result of powerful social learning, Boyd and Richerson suggest, is a situation in which social learning is deployed so effectively, and thus adopted so extensively, that auxiliary and potentially adverse effects begin to arise. For example, if too many individuals focus too much effort on learning from others, then there is a risk that incorrect knowledge or dysfunctional behavior can filter into the population, especially if we are inclined to copy the most popular behaviors. More generally, human culture represents the most dramatic instance of a species constructing its own environment, both in terms of the alteration of the physical environment and also in the construction of a cultural "virtual" environment within which human concerns play out. In Richerson and Boyd's words, "We do seem to have cut our way to our extraordinary adaptive success dragging a canoe-load of junk behind us."[70]

Opportunities abound for risky positive feedback to occur in Boyd and Richerson's model: the frequency bias means that popular traits become yet more popular, whereas unpopular traits might struggle to get off the ground even if they are beneficial. The hypothesized preference to learn from successful people might mean that those people's status grows greater under the attention awarded them by prospective learners. Under such conditions, strategies for success might shift from innovating objectively useful behaviors to behaviors that have a greater chance of being acclaimed, useful or not.

A framework that helps us make sense of this is niche construction.[71] The concept of niche construction came from the idea that in order to understand the circular complexities of biological evolution fully we should pay equal attention to the process of individuals modifying their environments, as to the environment shaping organisms through natural selection. Unlike earlier theories of biological coevolution, niche construction pays greater attention to the adaptability and potential agency of abiotic elements, much like the distributed creative systems of human-artifact interaction discussed above.[72]

A canonical example is the effect of rainforests on their abiotic environment, creating clouds and drawing down the temperature. Such climatic effects can be seen to happen on a global scale, and indeed we know that the earth's temperature is lower than it would be if it had no life on it. The various species making up this environment both contribute to its creation and over evolutionary time continue adapting to the new conditions that are created. Other species also adapt to the newly constructed niche and might further drive the changing conditions for evolutionary adaptation. The result is commonly described by physicists as "far from equilibrium," a situation that is sustained not because everything has settled into its resting state but because a complex network of dynamic processes acts to maintain a far more complicated state.

Without involving the modification of the environment by species, niche construction theorists argue, an understanding of the evolution of those species is incomplete—they did not just adapt to a fixed environment, but instead coevolved along with their environments. As this idea has picked up, pitched by its proponents as the "neglected process in evolution," it can be seen everywhere. Nature is constantly creating new types of niche. Rainforests establish a canopy and thus a shaded mineral-rich environment under the canopy. The wood of trees uniquely defines the niche of the woodpecker, and so on. Rich ecosystem complexes form, in which species contribute to the niches of other species in an endlessly creative process.

Accordingly, the human case is niche construction on steroids, with an extensive reconfiguration of our physical environment—clothes, buildings, spectacles, submarines—and the creation of culturally constructed virtual realms of activity—religions, stories, law, morality, social roles—both seemingly subject to a long-term cumulative trend. Niche construction helps us think of these phenomena as coevolution between ourselves and our niches. For example, if we communicate via language, then selection pressures might involve being able to handle complex grammar, or having a loud voice, or being able to remember complex stories. If we wear clothes, then selection pressures might involve the fine motor skills required to manipulate materials in the fabrication of clothes, while the need to have warm body fur might become less important. If our lives are structured according to concepts of magic and religion, then selection pressures might be related to abilities to perform in certain ways that befit special roles: act piously, do magic, demonstrate leadership, remember names and life histories, or at least convince others of these abilities. Each of these selection pressures builds upon the constructed niche created by earlier biological or cultural evolution. In theories of the evolution of contemporary civilizations, subsistence agriculture established the conditions for cities to emerge. Spoken language established the conditions for written language to emerge, and so on.

Niche construction also provides a useful framework for thinking about the division of labor. Consider the activities that contemporary humans engage in: politicians, comedians, thieves, cooks, con artists, carers, priests, monks, musicians, and computational creativity researchers. Each of these can be seen as building upon layers of constructed niches and further creating niches for others—social media officers require courses on the latest trends in social media, for example, creating a niche for the social media educator or social media trend analyzer.

Creativity as the Situated Production of Novelty

This diversion into the cumulative creativity of natural distributed processes, in their biological, cultural, and technological guises, has taken us a long way from the prototypical view of the lone creative human. Definitions of creativity focused on the human actor, attached to novelty and value, don't tell us everything we need to know about creativity in these large-scale distributed contexts. Nature produces novelty, but there is no agency to which it has value; there is no intention underlying the creation of species. But this is also true of some long-term sociotechnical processes that may be viewed as autopoietic and not driven by a designer or long-term purpose. This can be seen most explicitly in Arthur's view of ideas and technologies recombining in new combinations.

Under a perspective that admits this distributed view of creativity, we can't strictly always say that creative processes produce things of value, only that they produce new things. This may be seen in the distinction between individual artistic creation that may largely be a product of some artist's intention, and an artistic movement that cannot be said to be the product of any one person but of a collective, emergent process, even if every work in the canon of that movement was the product of adaptive creativity. I may value certain styles of music, but I am myself the product of a process of coevolution with those styles, my value system being shaped over my lifetime and informed by a lineage of other music listeners before me. That is, the value of the style does not preexist the style itself. If listeners' value systems are not objective but contingent on myriad contextual factors that influence them (as we will see in the following chapter), then when viewed over a long enough time scale these listeners' tastes are as much the product of a creative process as the music they enjoy. Again, this may be best viewed by considering how human actors or other creative elements are situated in a broader system, both influencing and being influenced by it. The situatedness of different elements defines different frames with which we can view different forms of creative process.

We have now gone long enough without specifically considering artistic behavior. In the following chapter, the above ideas are developed in their application to this domain of activity. The result is a multifaceted view of artistic creativity, with particular focus on the sociocultural and its emergent nature, held up by forms of evolutionary feedback. At last, in chapter 5 we arrive at a discussion of the existing algorithmic methods being used in computational creativity.

The open access edition of this book was made possible by generous funding from University of New South Wales.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Bown, Ollie, author.
Title: Beyond the creative species : making machines that make art and music / Oliver Bown.
Description: Cambridge, Massachusetts : The MIT Press, [2021] | Includes bibliographical references and index.
Identifiers: LCCN 2020017931 | ISBN 9780262045018 (hardcover)
Subjects: LCSH: Technology and the arts. | Arts—Technological innovations.
Classification: LCC NX180.T4 B68 2021 | DDC 700.1/05—dc23
LC record available at https://lccn.loc.gov/2020017931

10 9 8 7 6 5 4 3 2 1