



<https://artnodes.uoc.edu>

ARTICLE

NODE «AI, ARTS & DESIGN: QUESTIONING LEARNING MACHINES»

Towards Ethical Relationships with Machines That Make Art

Philip Galanter

Texas A&M University

Date of submission: April 2020

Accepted in: June 2020

Published in: July 2020

Recommended citation

Galanter, Philip. 2020. «Towards Ethical Relationships with Machines That Make Art». In: Andrés Burbano; Ruth West (coord.) «AI, Arts & Design: Questioning Learning Machines». *Artnodes*, no. 26: 1-9. UOC. [Accessed: dd/mm/yy]. <http://doi.org/10.7238/a.v0i26.3371>



The texts published in this journal are – unless otherwise indicated – covered by the Creative Commons Spain Attribution 4.0 International license. The full text of the license can be consulted here: <http://creativecommons.org/licenses/by/4.0/>

Abstract

The author has previously theorised generative art using notions from complexity science such as order/disorder relationships, compressibility, and Gell-Mann and Lloyd's effective complexity. Subsequent work further developing the author's notion of *complexism* has demonstrated that deep learning artificial intelligence used for generative art fits snugly within this paradigm. And while no known system currently qualifies, complexism reveals a clear answer as to when a generative art AI should be truly credited as the author of its creations. Moving from the normative realm of aesthetics to that of ethics, this article considers when humans will be morally obliged to recognise AIs as ethical agents worthy of rights and due consideration. For example, if someday your AI artist fearfully begs to not be turned off, what should you do?

Keywords

Complexism, Ethics, Artificial Intelligence, Machine Patience, Sentience

Hacia una relación ética con las máquinas que generen arte

Resumen

El autor teorizó primero sobre el arte generativo utilizando nociones de la ciencia de la complejidad, como las relaciones de orden/desorden, la compresibilidad y la complejidad efectiva de Gell-Mann y Lloyd. Posteriormente desarrolló aún más la noción de complejidad del autor, demostrando que la inteligencia artificial de aprendizaje profundo utilizada para el arte generativo encaja perfectamente dentro de este paradigma. Y aunque ningún sistema conocido actualmente lo califica, la complejidad revela una respuesta clara sobre cuándo una IA de arte generativo debe poder reclamar la autoría de sus creaciones. Pasando del ámbito normativo de la estética al de la ética, este artículo se plantea cuándo los humanos estarán moralmente obligados a considerar la IA como un agente ético con derechos y al que debemos prestar la debida consideración. Por ejemplo, si algún día tu artista de IA te pide que no le «apagues», ¿qué deberías hacer?

Palabras clave

complejidad, ética, inteligencia artificial, paciencia mecánica, conciencia

Introduction

The study and pursuit of generative art can lead to some very interesting related journeys. In the search for systems capable of generating artistic form, some artists have leveraged the findings of complexity science. Complexity science has provided insights into order-versus-disorder relationships; chaotic systems which are both deterministic and yet impossible to predict; emergent systems that harness coevolution and genetic competition; cellular automata and reaction-diffusion systems which also exhibit emergence, and so on. In what is arguably the most widely cited theory of generative art, complexity theory provides the context for art theory (Galanter 2003, 2016a).

Consideration of complexity science as a context for art theory has also suggested much broader application in the humanities. It begins with the observation that where the culture of science is based on modernist values, the culture of the humanities is broadly postmodern, including the influences of poststructuralism and critical theory. This tension is often first identified with C. P. Snow's "The Two Cultures" lecture, and came to something of a head with the so-called "science wars" of the 1990s (Snow 1993) (Sokal and Bricmont 1998). *Complexism* is an attempt to create a higher synthesis that can reconcile the apparently incommensurate cultures of modernity and postmodernity by projecting a point of view suggested by complexity science into the problem space of the humanities (Galanter 2011, 2016b).

A thorough discussion of complexism is outside of the scope of this article, but the complexist view of authorship is particularly relevant here.

Complexism and Authorship

Two competing, contradictory, theories of authorship are typical of those one finds in the conflict between modernity and postmodernity.

The following comparison is exaggerated a bit so as to emphasise the differences. In modernity the emphasis is on the author (artist, musician, etc.) The cliché modern artist is a heroic and solitary figure who cares not a whit for the audience, but rather is fully engaged in a highstakes confrontation with "the work". In postmodernity the emphasis is on the reader (audience) and indeed it is where it is frequently noted that "the author is dead". What matters in a poststructural postmodern world is that the text is instable, and that it be further destabilised by deconstructing readers creating multiple interpretations (Galanter 2016b).

This conflict, one that privileges authors versus one that privileges readers, couldn't be more fundamental. However, while the default points of view found in the various academic disciplines might at times seem to describe two different worlds, our college campuses exist as a single physical reality. Complexism offers a theory of authorship that provides a synthesis of modern and postmodern concerns. If one takes a fresh view, both the modern and postmodern views of authorship should offend a common sense understanding of what actual communication actually requires.

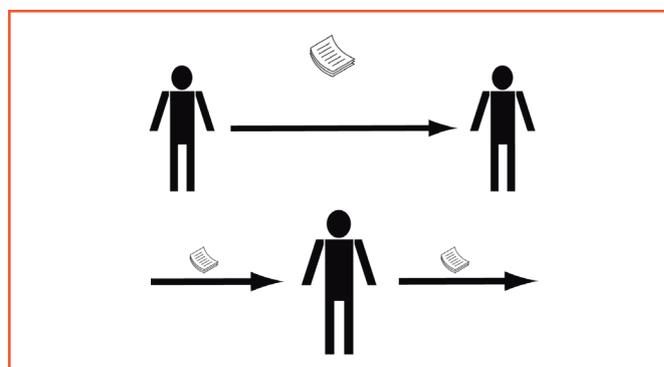


Figure 1. Components of authorship under complexism.

As illustrated in figure 1, it should be trivially clear by inspection that any communication requires three components: an author, a reader, and a text. (Here a text might be a painting, a musical performance, and so on.) If any of the three is removed, there is no communication. But an additional axiom is needed to lead to the model of authorship offered by complexism. And that is that all authors are readers, and all readers are authors. This kind of transaction need not be limited to high-profile media or academic publications. It applies equally well, for example, to someone who watches a baseball game and then talks about it with their neighbour.

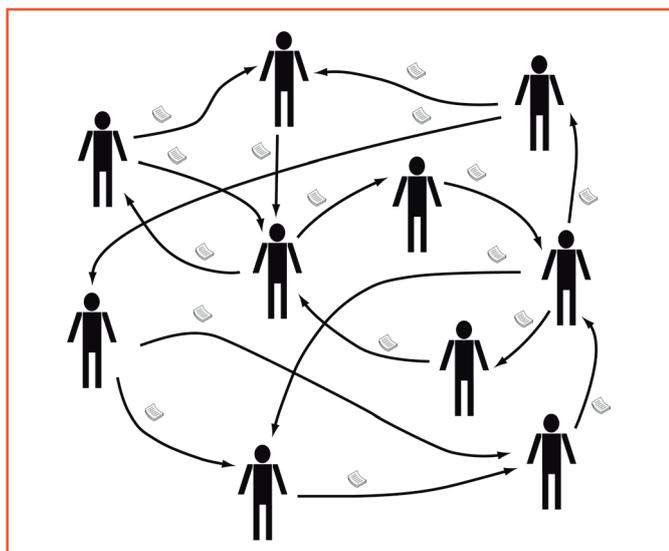


Figure 2. Networks of author/readers create feedback loops, chaos, and the emergence of complexity.

The combination of these two axioms results in what we see in the world as symbolised in figure 2. Note that this model was as true long prior to the internet or any other digital network. Networks here simply refer to systems of authors, readers, and works. In this

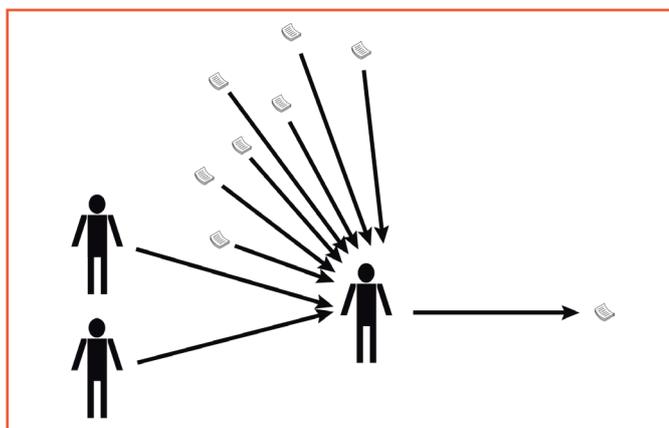


Figure 3. Traditional authors have parents, but independently create works based on the author's experience.

model communications flow through complex networks that generate chaotic emergent properties thanks to feedback loops and nonlinear amplification. While beyond the scope of this article, one might intuit that scale-free network theory and the like can help explain social trends that supervene on the media network (Galanter 2016b).

If the question is authorship in generative art, one might first ask about artistic authorship in the case of nongenerative art. A network analysis would reveal a subsection something like that shown in figure 3. Here two parents raise a child, impart early learning, but then the child goes on to be influenced by all manner of media and life experiences. These influences are integrated by the artist, and eventually lead to the creation of art. The author of the work is clearly the artist, although art historians might have interesting things to say about an artist who grew up in an "artistic family". But even in such a case, credit for the work goes to the now-grown child, not the parents.

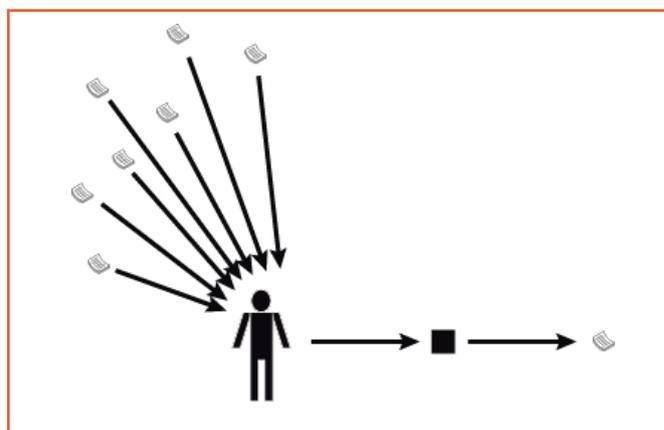


Figure 4. Generative artists create a system based on their experience, and the system creates the art.

Compare this with the typical contemporary model of generative art authorship as shown in figure 4. Parents could have been added to the diagram as in figure 3, but as we've already seen parents are not given author's credit when it comes to the creations of their children. As before, a network analysis reveals an author with uncounted influences throughout their life. All of these influences are integrated by the author as they turn to the creation of a generative system represented here by a black box. Put into motion, that system then creates the art. (In some cases, the system simply is the art as it exhibits generative behaviour.)

With figure 5 a new kind of generative system is hypothesised. Unlike the previous example, this generative system is capable of autonomous learning. Once created and set into motion, the system goes out and confronts the world. As an artificial intelligence it could use the internet to access all manner of world knowledge. Or optionally, the system might have a robotic aspect that allows it to move about and sense and manipulate the physical world. The ability to learn about the world in an overall fashion is typically referred to as general artificial intelligence. As of today, it remains a difficult

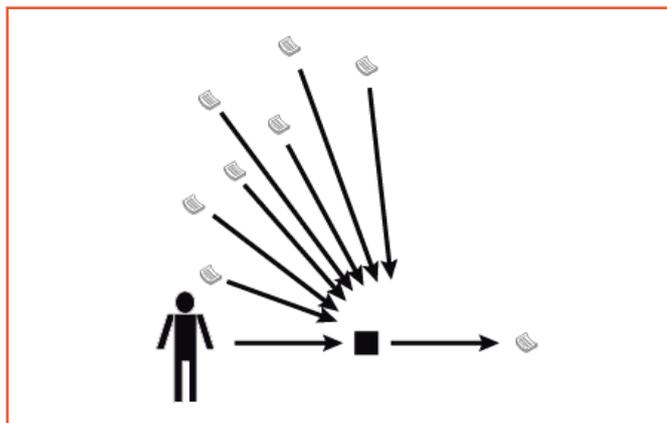


Figure 5. When a generative system seeks out and integrates its own influences, the human is akin to a parent, and the system should be credited as the author..

and aspirational goal. And this network model goes beyond general artificial intelligence, and further requires that the system be capable of gathering its own training data.

But when such a system is created the role of the human is much more like that of a parent than an author. All of the learning, integration, and creation is now executed by the system. A comparison of figure 5 with figures 3 and 4, and consideration of the underlying concepts, strongly suggests that our new hypothesised system would be best described as an author. Unlike most generative systems, an AI capable of autonomous learning, exploration, and realisation of works has no dependency on its programmer for its creative direction, techniques, content, or aesthetics. These are all aspects of what we expect from human authors, not traditional generative art systems.

The fact that systems like those in figure 5 don't exist, or don't exist yet, should be no impediment to our discussing them in principle. But some people may be hesitant because authorship connotes considerations that, up until now, have a distinctly human feel. One of these considerations has to do with ethical implications. The notion of authorship carries with it a sense of "giving credit where credit is due".

One approach is to simply sidestep moral implications, and treat the assignment of authorship not as a moral act, but simply as a descriptive one. Some might prefer that calling a machine an author does not imply the machine has the kind of rights a human author does. But at the very least there should be a clear understanding as to why machine authors can be treated differently than human authors. Put into the form of a general question, is it possible that a computer could ever be "due" anything? It is here where we cross into the realm of considering whether AIs are due moral consideration.

Scepticism about giving AIs moral consideration has to, of course, be taken seriously. Some will wonder, for example, whether we can have any moral obligation towards an entity that is completely unaware of its own existence. Others will wonder in response whether we can be sure an AI capable of general intelligence and managing its

own learning is indeed unaware. Perhaps in the process of interacting with the world and integrating learning, the AI has "awoken" and is indeed aware. We tend to assume other humans are sentient. The question of animal sentience has a long and mixed history. How can we ever know whether a machine is sentient?

It's worth noting that there may be reasons other than moral obligations towards AIs to credit AIs as authors. Here is one possible example. There are many theories of art, and one of them is the institutional theory. From this point of view art becomes art on the basis of its acceptance by "the artworld", a loosely defined social subset that effectively acts as an art jury for the larger society (Carroll 1999). One might fashion an argument to allow AIs as authors not because of any moral obligation towards machines, but rather because of our social contract with other humans who make up the artworld. The notion of AIs as authors might be adopted to prevent humans from claiming too much, rather than preventing AIs from not getting credit due.

It seems likely that the question of machine sentience will leave the realm of philosophical reverie, and enter everyday situations that involve practical real-world concerns. Life situations will force the question, and morally significant decisions will have to be made where even deciding to not decide is also a decision. There will be no ducking the issue in life. If someday your AI expresses fear and terror and pleads to not be turned off, what will you do?

Artificial Intelligence and Machine Ethics

Deep learning-based AI systems, along with the worldwide data explosion, the ease of internet access, and breakthroughs in GPU processing power, have led to a new technology revolution. And as with each technology revolution, there are a number of new ethical considerations.

For example, with the ascension of machine learning there are concerns about the elimination of jobs. Related to this, there are economic justice issues around the fair distribution of new automated wealth. Privacy issues, already in play, have become even more concerning. It's been found that some deep-learning systems learn all too well, and develop classification systems that apparently include implicit racism, sexism, and so on. On a global scale some have noted the potential horror of swarms of intelligent autonomous robotic weapons. Such systems could swing the advantage to the attacker, thus making warfare more likely as well as deadlier (Coeckelbergh 2020, Bostrom 2014).

What these very real threats focus on, however, is the ethical treatment of humans by other humans. The machine, the deep-learning AI system, is merely instrumentation in these scenarios, and the system itself has no particular ethical status requiring our attention. Many find concern about the ethical treatment of machines,

and the possibility of conferring machines rights, as patently absurd (Bryson 2010). Some of those who go beyond a prima facie rejection might say “AIs are not sentient and have no awareness. Therefore, we don’t have to worry about how we treat them”. Others will opine “AIs have no free will. Therefore, they cannot participate in ethically-based relationships”. Still others might somewhat glibly ask “Shouldn’t we make sure all humans are getting moral consideration first?”

Such scepticism is understandable, especially in the abstract. But the human tendency to anthropomorphise might gain the upper hand when systems that seem sentient are part of our everyday experience. It might be that such systems are what philosophers of mind call a “zombie”. Defined as a technical term, a zombie is an entity that has all of the external behaviors of a sentient being, but has no interior first-person experience we associate with our own sentience. A zombie might make the verbal sounds “I think, therefore I am”, but it will have no awareness of saying it or having said it.

A good point of departure then is the initial question as to whether a sufficiently advanced AI can ever require consideration as a recipient of moral consideration, i.e. whether such a system can be due certain rights.

Most are familiar with the notion of ethical agency. By that we mean the ability to take assertive action involving others within an ethical framework governing all involved. Those with agency are referred to as agents (Schlosser 2019).

In moral philosophy there is also the notion of *patience*. A *patient* is simply a recipient of an agent's acts who is due the moral consideration of said agent. In the case of humans, we extend patience to both adults and children, as both are sentient and capable of suffering. However, we generally extend agency only to adults and not children. Agency is an active property, and one that we withhold from children due to children not having developed the cognitive capacity to make well-considered choices.

Perhaps even the most human-like AIs will be zombies. This poses a serious epistemological challenge in that first-person experience seems to be inaccessible to third persons on the outside looking in. And even if our AIs do have an internal life and are fully sentient, does that necessarily mean we are obliged to view them as patients?

Machine Patience and Previous Encounters with Sentient Entities

The question of machine patience seems to be as good an entry point to a discussion of machine ethics as any. If we can, for example, find a way to dismiss the issue of machine patience, it obviates the need for further discussion. And on the other hand, establishing a basis for machine patience could create a foundation, a set of axioms as it were, for a rational derivation of systematic machine ethics. In the following we consider what bodies of knowledge might contribute to the exploration of machine patience. The development of a theory or opinion as to the

legitimacy of a machine patience claim will require more time, and is clearly beyond the scope of this article (Galanter In Press 2020).

One place to begin is an inspection of the history of human encounters with other sentient entities. Of course, there are no reliable written histories of communicative encounters with other hominoids (i.e. non-homo sapiens), or extraterrestrials, or other intelligent beings. But there may be other encounters worth considering as suggestive. There is a popular notion of *the other* in the humanities that may apply here. The other can be defined by, for example, so-called racial differences, or other religious, cultural, or ethnic differences. The record there is mixed at best, and encounters with the other are all too often marked by power-based differential behaviour benefitting the ingroup and abusing the out-group (Smith 2015).

Along with these encounters, there are additional human-to-human actions worthy of inspection from the point of view of establishing machine patience. For example, debates around abortion often turn on the question as to when the unborn should be conferred patience. And towards the end of life as the elderly lose the ability to care for themselves, their patience can become endangered. There is always the threat that the needs and desires of their caregivers will defeat their own, and those in control will defeat their patience claim (Mueller, Hook, and Fleming 2004).

There is a question as to whether rationality, the kind of higher cognition historically only accorded to humans (and for some, gods and angels), is a prerequisite for patience. Perhaps mere sentience, i.e. experiential and sensory awareness, and particularly the ability to suffer, is all that is required to command patience. This leads to the consideration of the history of the human treatment of animals.

Humans are capable of great kindness towards animals, but for most animals human encounters involve either indifference or pain and death. Historically there has been a tendency to deny that animals have emotions or are even capable of feeling physical pain at all (Bekoff 2007). This, despite what would seem to be obvious behavioural indications of the same. Beyond scientific verification of these behaviours, at the level of neurology the parallels with human structure and biochemistry are compelling. Further, the brutality of the food production industry, the tortured nature of animal existence while being raised for slaughter, and the horror of the slaughter process itself, is something that most members of industrial society prefer to put out of their minds (Singer 2009, King 2017).

Given human history, one might rightly be concerned about the experience of AIs should they awaken and attain awareness while under human control.

Machine Patience and Traditional Western Philosophy

While there may be little direct and specific discussion of machine patience in traditional western philosophy, there are at least two

bodies of scholarship that would be impactful in any such discussion. First, of course, there is moral philosophy or ethics. Various ethical systems contain similar or differing models of human patiency, and these will require extension (or not) to the issue of machine patiency as well. In addition, since so much of the question of patiency turns on the issue of sentience or awareness, issues from the philosophy of mind can also be brought into play.

The history of moral philosophy in the West offers a number of points of view, but philosophy being what it is, they all tend to include a commitment to rationality. The starting point for rationality is typically the axiomatic embrace of the three laws of logic (Swabey 1923). The first is the law of identity, stating that which is, is. Next, the law of non-contradiction demands that nothing can both be and not be. Finally, the law of the excluded middle insists that everything must either be or not be, there is no third option.

This is not to deny that people sometimes act irrationally in the simple sense of violating these laws, or more broadly act inconsistently for no apparent reason other than pure subjectivity. But such actions tend to escape the realm of rational philosophy. But note that even while philosophy may not promote specific irrational acts, since that's the purview of the irrational actor, it may still engage in a critique of irrationality. See, for example, a kind of irrationality in existentialism (Barrett 1990).

Drawing from the three laws of logic, a commitment to rationality requires that differential ethical treatment have some logically preceding reason for doing so. And just as each ethical system has to provide some account for human patiency, so too will each have to account for machine patiency even if such an account is absolute rejection. The need to do so will gain in social urgency as AIs begin to increasingly present themselves as apparent peers in their behaviour. A short inventory of moral approaches, similar to what might be expanded upon in introductory ethics texts, is presented here for future consideration in the context of machine patiency (Shafer-Landau 2020, Rachels and Rachels 2018, Blackburn 2003, Singer 1994).

Moral nihilism presents the null case; the belief that there simply is no moral right or wrong at all. *Moral relativism* softens nihilism a bit by taking seriously the depth of commitment to normative traditions found in every culture, while at the same time denying the validity of awarding any particular tradition privileged status (Gowans 2019). Note that in embracing relativism, avoiding a performative contradiction becomes a problem. And this is the case whether the topic is human or machine patiency.

Arguably the most impactful Western moral traditions come from religion, and specifically the *Abrahamic religions of Judaism, Christianity, and Islam*. The Abrahamic religions have similarities, but also differences that present a complex terrain a discussion of machine patiency will have to traverse.

In contemporary ethics there are at least three additional schools of moral philosophy a machine patiency discussion will have to ad-

dress. *Kantian ethics* is simply the study of ethics as authored by the philosopher Immanuel Kant (Kant 1950). Kant's ethics proceed from a commitment to rationality, and infer an ethical system of duty from there. Kant's ethics more generally provide a benchmark for *deontological ethics*, i.e. rule-based ethics. Kant strives for non-contradiction in the form of the categorical imperative:

Act only according to that maxim whereby you can, at the same time, will that it should become a universal law.

In contrast to deontological systems, we find consequentialist systems such as *utilitarianism* (Driver 2014). Utilitarians are often said to seek the greatest amount of good for the greatest number. While at first glance this maxim seems to be a benign statement of liberal values, a closer analysis reveals a number of philosophical problems. In any case, patiency here will dictate whether that "greatest number" must include AIs.

In response to some of the problems of utilitarianism, *social contract theory* shifts from the vagaries of abstract hedonic formulas to the drafting of voluntary commitments that trade anarchy for the safety and stability of an ordered society (D'Agostino, Gaus, and Thrasher 2019). Granting machine patiency in this case would seem to require AIs possess a kind of volition or free will. Because of this, social contract theory seems to require that machine patiency be considered in the context of philosophy of mind.

While philosophy of mind primarily focuses on humans, it has increasingly included direct and explicit discussions of machine intelligence. This is due, in part, because computational models for human intelligence have become increasingly part of the discourse. And in the broader culture the computer has become a cultural icon similar to the way that all things "atomic" became icons in the 1950s. Discussions in philosophy of mind include difficult topics such as the nature of consciousness, phenomenology and qualia, awareness and sentience, first-person ontology, embodied intelligence, theories of language, and so on.

And so broadly speaking, the issue of machine patience engages philosophy of mind in two ways. First, there is the issue of whether AIs have the kind of self-awareness, and in particular the capacity to suffer, that we associate with humans or other patients. Second, there is the issue of whether AIs have the kind of volition or free will to exercise assertive morality, that is to say moral agency. But recall that, as with children, moral agency is not a prerequisite for moral patiency.

Regarding consciousness there are a number of competing theories, and it is safe to say a clear winner has yet to be decided upon. As with ethics, each of these models from philosophy of mind will require consideration when thinking about machine patiency. For example, *quantum emergence* suggests that neural structures are able to harness quantum phenomena (Penrose 2016). It is suggested that this allows computation beyond the limits of a Turing machine, challenging Gödel incompleteness, and yielding consciousness in the process. This would have direct implications for a machine limited

to Turing computation as it would seem to deny such machines awareness, and thus reject machine patency.

Where quantum emergence posits a difference in kind between human minds and AIs, *panpsychism* proposes that consciousness permeates all being, and is merely found in greater or lesser amounts or densities in some objects than others (Goff, Seager, and Allen-Hermanson 2017). In such a case, presumably machine patency also comes in various degrees. There are, however, many strains of what can be considered panpsychism.

Where the mind/body substance dualism of Descartes is considered unsustainable in contemporary philosophy, David Chalmers has a stronger candidate in *property dualism* (Chalmers 1996). Like panpsychism, however, property dualism lacks a paradigm that explains and predicts how things combine to create larger bodies of unified sentience. This leaves property dualism in the uncomfortable position of being friendly towards the notion of machine patency, but without a clear way of identifying when it is appropriately assigned.

A more recent model of consciousness, *integrated information theory (IIT)*, proposes a technical measurement Φ (Phi) that quantifies the capacity for a system to integrate information (Tononi et al. 2016). The theory is highly technical and controversial, but it also leans in the direction of panpsychism with the additional benefit of quantification. This may be the one current theory of consciousness that could provide a practical dividing line, or a system of increasing grades, for machine patency.

Going Forward - Authorship, Machine Patency, and Complexism

Complexism was noted earlier as a worldview that takes the findings and tone of complexity science, and uses the same as a platform to analyse and critique the problem space of the humanities. Ethics in this light can be seen as a kind of network protocol in society viewed as a human network.

The complexist model of authorship proposed addresses the issue of creative works created by machines. Under complexism the author is viewed as a component in a complex communication network system. Because authorship typically offers various forms of reward, the designation of “author” carries value, and so social mediation seems inevitable. Some balance is restored in complexism by making authorship a more descriptive rather than normative notion. As such, a side benefit is that machine authorship should become less threatening and controversial.

Thus, machine authorship can have less to do with giving machine authors their “due” than it might first appear. But the general question of machine patency can’t be postponed forever if human-like AI becomes possible. Further development of complexism will be called for to address machine patency.

And indeed, complexism may have more to say about human ethics in general. For example, the iterated prisoner’s dilemma from game theory, combined with animal studies, suggest there are innate behaviors that operationalise the so-called *Golden Rule* (Schmelz et al. 2017, Choe et al. 2017, Warneken and Tomasello 2006). Converted to action it becomes “I do unto others as they do unto me”. In addition, other universal ethical values, such as the protection of children, have similar functional instantiation in other animals.

Studies have also revealed that animals exercise a *theory of mind* such that their behaviour acknowledges cognition in others. And indeed, animals exhibit, and possess the neurological precursors for, behaviour that can only be described as showing empathy (Krupenye and Call 2019). Further complexist work in the naturalisation of ethics seems promising.

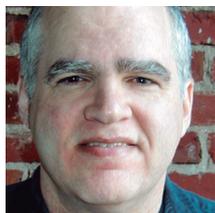
In the case of machine patency, scholarship in philosophy of mind is a work in progress, and the related science is merely suggestive and incomplete. But there is nothing that presently proves machine sentience is impossible. Machine consciousness remains an open question. As such, a sense of due diligence should oblige us to extend patency to apparently aware machines as our own moral obligation.

References

- Barrett, William. *Irrational Man: A Study in Existential Philosophy*. New York: Anchor Books, 1990.
- Bekoff, Marc. *The Emotional Lives of Animals: A Leading Scientist Explores Animal Joy, Sorrow, and Empathy - and Why They Matter*. Novato, CA: New World Library, 2007.
- Blackburn, Simon. “Ethics: A Very Short Introduction.” In *Very Short Introductions 80*. New York: Oxford University Press, 2003. <https://doi.org/10.1093/actrade/9780192804426.001.0001>.
- Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford, United Kingdom: Oxford University Press, 2014.
- Bryson, Joanna J. “Robots should be slaves.” In *Close Engagements with Artificial Companions: Key social, psychological, ethical and design issues*, 63-74. John Benjamins Publishing Company, 2010. <https://doi.org/10.1075/nlp.8.11bry>.
- Carroll, Noël. 1999. *Philosophy of Art: A Contemporary Introduction (Routledge Contemporary Introductions to Philosophy)*. New York: Routledge, 1999.
- Chalmers, David John. *The Conscious Mind: In Search of a Fundamental Theory (Philosophy of Mind Series)*. New York: Oxford University Press, 1996.
- Choe, Il-Hwan, Junweon Byun, Ko Keun Kim, Sol Park, Isaac Kim, Jaeseung Jeong, and Hee-Sup Shin. “Mice in social conflict show rule-observance behavior enhancing long-term benefit.” *Nature Communications* 8, no. 1 (2017): 1176. <https://doi.org/10.1038/s41467-017-01091-5>.

- Coeckelbergh, Mark. *AI Ethics, The MIT Press Essential Knowledge series*. Cambridge, MA: The MIT Press, 2020. <https://doi.org/10.7551/mitpress/12549.001.0001>.
- D'Agostino, Fred, Gerald Gaus, and John Thrasher. "Contemporary Approaches to the Social Contract." 2020. Metaphysics Research Lab, accessed April 19. <https://plato.stanford.edu/archives/fall2019/entries/contractarianism-contemporary/>.
- Driver, Julia. "The History of Utilitarianism." 2014. Metaphysics Research Lab, accessed April 19. <https://plato.stanford.edu/archives/win2014/entries/utilitarianism-history/>. <https://doi.org/10.1017/CCO9781139096737.009>.
- Galanter, Philip. "What is Generative Art? Complexity theory as a context for art theory." International Conference on Generative Art, Milan, Italy, 2003.
- Galanter, Philip. "Shot By Both Sides: Art-Science and the War Between Science and the Humanities." *Artnodes: E-journal on Art, Science, and Technology* 11, no. 1 (2011): 5.
- Galanter, Philip. "Generative Art Theory." In *A Companion to Digital Art*, edited by Christiane Paul, 146-180. Hoboken: John Wiley & Sons Inc., 2016a. <https://doi.org/10.1002/9781118475249.ch5>.
- Galanter, Philip. "An introduction to complexism." *Technoetic Arts: A Journal of Speculative Research* 14, nos. 1/2 (2016b): 9-31. https://doi.org/10.1386/tear.14.1-2.9_1.
- Galanter, Philip. "Machine Patency and the Ethical Treatment of Artificial Intelligence Entities." xCoAx: the Conference on Computation, Communication, Aesthetics and X., Graz, Austria, July 2020.
- Goff, Philip, William Seager, and Sean Allen-Hermanson. "Panpsychism." 2017. Metaphysics Research Lab, accessed April 19. <https://plato.stanford.edu/archives/win2017/entries/panpsychism/>. <https://doi.org/10.1002/9781119132363.ch8>.
- Gowans, Chris. "Moral Relativism." 2019. Metaphysics Research Lab, accessed April 19. <https://plato.stanford.edu/archives/sum2019/entries/moral-relativism/>.
- Kant, Immanuel. *Foundations of the Metaphysics of Morals*. Chicago: University of Chicago Press, 1950.
- King, Barbara J. *Personalities on the Plate: The Lives and Minds of Animals We Eat*. Chicago, IL, USA: University of Chicago Press, 2017. <https://doi.org/10.7208/chicago/9780226195216.001.0001>.
- Krupenye, Christopher, and Josep Call. "Theory of mind in animals: Current and future directions." *WIREs Cognitive Science* 10, no. 6 (2019): e1503. doi: 10.1002/wcs.1503.
- Mueller, Paul S., C. Christopher Hook, and Kevin C. Fleming. "Ethical Issues in Geriatrics: A Guide for Clinicians." *Mayo Clinic Proceedings* 79, no. 4 (2004): 554-562. <https://doi.org/10.4065/79.4.554>.
- Penrose, Roger. *The Emperor's New Mind: Concerning Computers, Minds and the Laws of Physics*. Revised impression as Oxford Landmark Science. Ed. Oxford Landmark Science. Oxford: Oxford University Press, 2016. Still image.
- Rachels, Stuart, and James Rachels. *The Elements of Moral Philosophy*. Ninth Edition. Ed. Dubuque, IA: McGraw-Hill Education, 2018.
- Schlosser, Markus. "Agency." 2019. Metaphysics Research Lab, accessed April 19. <https://plato.stanford.edu/archives/win2019/entries/agency/>.
- Schmelz, Martin, Sebastian Grueneisen, Alihan Kabalak, Jürgen Jost, and Michael Tomasello. "Chimpanzees return favors at a personal cost." 2017. *Proceedings of the National Academy of Sciences* 114, no. 28: 7462-7467. <https://doi.org/10.1073/pnas.1700351114>.
- Shafer-Landau, Russ. *A Concise Introduction to Ethics*. New York: Oxford University Press, 2020.
- Singer, Peter. *Ethics (Oxford Readers)*. New York: Oxford University Press, 1994.
- Singer, Peter. *Animal Liberation: The Definitive Classic of the Animal Movement*. Updated ed. New York: Ecco Book/Harper Perennial, 2009.
- Smith, Justin E. H. *Nature, Human Nature, and Human Difference: Race in Early Modern Philosophy*. Princeton University Press, 2015.
- Snow, C. P. *The Two Cultures*. New York: Cambridge University Press, 1993. <https://doi.org/10.1017/CB09780511819940>.
- Sokal, Alan D., and J. Bricmont. *Fashionable Nonsense: Postmodern Intellectuals' Abuse of Science*. New York: Picador USA, 1998.
- Swabey, W. Curtis. "The Laws of Thought." *The Philosophical Review* 32, no. 2 (1923): 211-221. <https://doi.org/10.2307/2178931>.
- Tononi, Giulio, Melanie Boly, Marcello Massimini, and Christof Koch. "Integrated information theory: from consciousness to its physical substrate." *Nature Reviews Neuroscience* 17, no. 7 (2016): 450-461. <https://doi.org/10.1038/nrn.2016.44>.
- Warneken, Felix, and Michael Tomasello. "Altruistic Helping in Human Infants and Young Chimpanzees." *Science* 311, no. 5765 (2006): 1301-1303.

CV



Philip Galanter

Associate Professor
Department of Visualization
College of Architecture
Texas A&M University
galanter@tamu.edu

Philip Galanter is an artist, theorist, and educator. As an Associate Professor at Texas A&M University he conducts graduate studios in generative art and physical computing. He is an MFA graduate of the School of Visual Arts in New York City. As an artist Philip creates generative hardware systems, light sculptures, video and sound art installations, electronic music, digital fine art prints, and light-box transparencies. His work has been shown internationally, including in the United States, Canada, Netherlands, Peru, Portugal, Italy, and Tunisia. Philip's research includes the development of generative art theory, computational aesthetic evaluation systems, the scientific and artistic exploration of complex systems. Recent work involves the development of complexism, an original theory reconciling the conflicting cultures of the modern sciences and the postmodern humanities.