

Chapter forthcoming 2018 in
A Networked Self Vol. 5: Human Augmentics, Artificial Intelligence, Sentience
Edited by Zizi Papacharissi

Tamara Shepherd
University of Calgary
tamara.shepherd@ucalgary.ca

AI, the persona, and rights

Introduction: AI and rights

The notion of rights in digital culture is a fraught one. Frameworks for inalienable human rights, themselves critiqued on the grounds of abstraction, Eurocentrism, and gender bias, have been subject to intense debate when transposed to networked spaces. The 2017 RightsCon in Belgium, for instance, featured talks on diverse aspects of internet rights, from regulatory jurisdiction to environmental impacts of technology to hate speech online to trade agreements. Germane to all of these themes, one conclusion that has emerged over the last 20-odd years of the internet's popularization has been a sense that digital culture is in some ways an inherently commercial construct, where ideals of democratization are merely rhetorical tools that shore up consolidated corporate power (e.g., Dean, 2005; Gillespie, 2010; Zuboff, 2015).

Given the overwhelmingly commercial nature of what most people experience as digital culture, and the location of that brand of commercialism predominantly in the U.S. (Jin, 2013), how might constitutional rights to personal integrity and freedom be considered? One framework for reinstating the legal viability of the person in digital culture is

William McGeeveran's (2009) formulation of "persona rights." McGeeveran sees much of digital culture as defined by the boundaries of social platforms supported by advertising that commodifies users' endorsements, in other words, a "like" economy (e.g., Gerlitz & Helmond, 2013). Since user data (privacy) and creative content (intellectual property) is appropriated by social platforms, McGeeveran argues, an individual should have recourse to their status as legal persona to exercise monopoly control over their own image.

Persona rights thus offers a novel legal framework for bolstering the integrity of the person within commercial spaces for digital culture.

Of course, persona rights reflects a liberal context and can thus be critiqued on the grounds that it re-entrenches neoliberal individuality as the basis for rights, forgoing socialist, feminist, and Indigenous versions of collective rights. However, as promising as such approaches may be for resisting the encroachment of neoliberal rationality, the current configuration of digital culture is one controlled by a handful of consolidated techno elites. This is especially the case in contemporary advancements in artificial intelligence (AI): the development of machine learning based on tightly controlled proprietary algorithms, growing in complexity from recommender systems (e.g., Netflix) to deep learning (e.g., Google). By creating algorithms with the capacity to optimize prediction, machine learning strategies are being used in diverse contexts from medical diagnosis to autonomous vehicles to weapons systems. Paying particular attention to the development of artificial general intelligence (AGI) – the capacity for computational machines to solve a variety of complex problems across domains by controlling themselves autonomously and even self-consciously (Pennachin & Goertzel, 2007) – this

chapter examines how machine-learning algorithms add fundamental complications to the already fraught terrain of digital rights. By framing such rights in terms of persona rights, I hope to add to a growing suite of means to reassert the human as a rights-bearing agent within posthumanism.

The posthuman conceit of AI

Thinking machines have been framed as “posthuman persons” (Hamilton, 2009, p. 143), in that their interactions with human interlocutors form the basis for our perception of their intelligence. AI systems both learn from humans and take their inspiration from the human nervous system, materializing Donna Haraway’s cyborg in a way that suggests a displacement of transformative agency onto machines. In this posthuman context, developments in AI exert the same pressure on the idea of the human as “special” that advances in our understandings of animal cognition do, which links the idea of posthuman rights to those of animal rights and positions humans as stewards of these other intelligences (Haraway 1991; 2008; Hayles 1999).

Posthumanism has been a salient concept within internet studies more generally, where networked communication rests on the unit of the individual user-person, whose subjectivity is configured through a “liberal strategy of individual self-worth” (Balakrishnan, 2016, p. 112). As social platforms have come to dominate most users’ experiences of the internet, personalization has become the hallmark of popular platforms driven by data-based advertising models. This intense personalization rests on the AI manifested in machine learning algorithms that benefit from huge amounts of data being

inputted by users of platforms like Netflix, Facebook, Google, and Amazon. Developed by large technology companies, such machine learning algorithms tend to be proprietary and non-transparent. The lack of transparency makes sense given the trade secrets that characterize technological development, but has significant implications when these algorithms move toward deep learning.

Deep learning is a subset of machine learning where the algorithm itself – rather than the programmer – defines features to be analyzed and how the outcomes should be optimized by modeling how the human brain works. Inputs to the algorithm are analyzed and processed. If the processing results in a correct outcome, those analytical pathways are reinforced; if the outcome is incorrect, the pathways are reconfigured (Sample, 2017). This sort of independent adjustment through learning is called backpropagation and conceptually resembles neural plasticity in human brains. Google has different divisions currently working on deep learning, including Google Brain and Google DeepMind. These divisions are working on backpropagation models but also developing an alternative that uses even less computing power, which is called the evolutionary strategies approach (Domingos, 2015). The evolutionary strategies model even more closely resembles how the brain works, where information only travels forward – not backward – across neuronal synapses (Lake et al., in press, p. 36). And the name for this process, “evolutionary,” more closely aligns AI rhetorically with the brain by using the language of a particularly human view of evolutionary adaptation: “Humans are the species that adapts the world to itself instead of adapting itself to the world. Machine learning is the newest chapter in this million-year saga” (Domingos, 2015, p. 3). Deep

learning represents a future of machines both adapting themselves to the world and in turn adapting the world to themselves, no longer subject to supervised learning under the guidance of programmers but engaged in reinforcement learning and autonomous decision making (Stone et al., 2016, p. 15).

When considering advancements in AI systems through the prism of posthumanism, perhaps the most tantalizing possibility is that of artificial general intelligence (AGI). Today's deep learning AIs are trained to learn specific tasks such as object or speech recognition (Stone et al., 2016; Weston et al., 2015). AGI presents a general purpose alternative that can handle a diverse set of tasks by drawing on a more human-like flexible memory (Pennachin & Goertzel, 2007). Such flexibility is made possible by mimicking the functional structure of the brain through neuromorphic computing. Neuromorphic computing replaces traditional computational models that separate inputs/outputs, instruction-processing, and memory, instead integrating these modules in a way similar to the brain (Stone et al., 2016, p. 9). Deep learning researchers foresee AGI arising from such attempts to get closer and closer to brain-like processing, where future neural networks are projected to be "endowed with intuitive physics, theory of mind, causal reasoning, and other capacities," including the capacity to "effectively search for and discover new mental models or intuitive theories" that can be used as the basis for subsequent learning (Lake et al., in press, p. 3)

In order to reach such a posthuman ideal, though, AGI must overcome a series of obstacles in its evolutionary progression. One significant barrier is computing power: an

average brain contains about 100 billion neurons, each connected to up to 10,000 other neurons through a total of between 100 trillion and 1,000 trillion synapses. To match that in computing power will require leaps in processing technology, perhaps only achievable through parallel or quantum computing. Nonetheless, it seems that AGI is somewhere on the horizon. And even though machines think in ways that are distinct from human thinking, deep learning still relies on homophily to the brain in terms of structural inspiration from neuroscience (Le et al., 2012), as well as psychological benchmarks that rest on comparisons with human intelligence (Kahn et al., 2007). As N. Katherine Hayles (1999) contends about AI, its trick lies in the tests (including famously, the Turing test) used to assess intelligence. If human intelligence is the benchmark for determining the effectiveness of AGI, we might miss the distinctly posthuman capacities that come from alternative models of learning (Stone et al., 2016, p. 13). Already, AI outstrips human capacities in tasks such as instantaneous prediction, and this is where posthumanism suggests the need for a politics. Given the development and deployment of AI within technology companies emblematic of late capitalism, what's needed is an extension of Winner's (1977) sense of technology as having a politics to account for AI's potential reshaping of democratic rights in a technocratic context (Damnjanović, 2015).

Rights and commerce

It is crucial to consider the posthuman context when adapting conceptions of rights, typically framed as liberal, humanist rights, to the specific situation of AI. Human rights is a relatively new concept, with that phrase only in common usage since the 1940s and the adoption of the United Nations' Universal Declaration of Human Rights (UDHR).

Yet the idea of inalienable rights of the person extends much further back in Western thought (from Ancient Greek societies through the Middle Ages), and concepts such as “natural rights” can be seen to come into their own as the cornerstone of Enlightenment philosophy. For example, in his *Second Treatise Concerning Civil Government* (1689), John Locke argued for men’s natural right to life, liberty, and property. Such arguments furnished state-level rights protections particularly in contemporary revolutionary contexts, as evidenced in the English Bill of Rights (1689), the American Declaration of Independence (1776), and the French Déclaration des droits de l’homme et du citoyen (1789).

While all of these implementations of the idea of natural rights differ according to their specific political contexts, they share a notion that people should be naturally endowed with certain basic rights, protections or freedoms that support respect for the person’s agency as a human being. These kinds of rights are framed by current human rights discourse as moral rights – rights attributable to a person simply because they exist as a human being – and should thus be distinguished from legal rights, or rights upheld by courts in a specific jurisdiction. Human rights treatises attempt to recognize moral rights as legal rights, although this is complicated by diverse understandings of both, subject to historical circumstance. For example, Kantian moral philosophy posits morality as fundamentally rational and moral responsibility as the basis for rights. Each person’s self-governing reason to act dutifully (in accordance with the categorical imperative), in Kant’s view, provided a deontological version of rights whereby that person could be viewed as possessing equal worth and deserving of equal respect as another. Yet the

universalizability of Kantian moral philosophy, while it accords with ideas about universal human rights, exists in abstraction more so than in lived reality. A more situated historical approach to rights as ethical rather than moral positions human rights more directly within contemporary political contexts by seeking to define what exactly is meant by rights under the current “human rights regime” (e.g., Rawls, 2001).

The current regime of the UDHR reflects something of a continuity between the Enlightenment moment for natural rights within emergent capitalism and the late capitalist context of the twentieth century. In both eras, ideas of personhood coincided with models for extracting surplus value from the person in the form of labor power, in Marxian terms. Legally recognized as “persona,” personhood has come to take on a commodity form, where “the persona confirms the labor and authoring capacity of the individual person” (Hamilton, 2009, p. 190). What changes with digital culture is the transition to surveillance capitalism (Zuboff, 2015), which warps the values of individualism central to liberal and neoliberal versions of the persona as a rights-bearing agent by segmenting individuals into even smaller units – “dividuals” (Deleuze, 1992; Terranova, 2004) – through data collection regimes that feed into AI systems on networked platforms.

While humans get thus reconfigured as dividuals through their contractual relationships with technology companies that run social platforms, the evolution from AI to AGI suggests a new formulation of posthuman rights. Anthropomorphizing machines has consequences for legal frameworks around agency (Hamilton, 2009, p. 175). For

instance, a 2016 proposal to the European Parliament recommends granting personhood status to thinking machines according to the EU Charter of Fundamental Rights (Delvaux, 2016). Largely, this appeal to the persona of AI rests on concerns about liability, where machines might be held legally responsible for their own actions. The proposal goes on to suggest that the autonomy of advanced AI systems “raises the question of their nature in the light of the existing legal categories – of whether they should be regarded as natural persons, legal persons, animals or objects – or whether a new category should be created” (Delvaux, 2016, p. 5). The concern here is the legal status of AI in relation to existing property regimes, including the way that machine learning relies on the collection of “personal data as a ‘currency’ with which services can be ‘bought’,” challenging other EU regulations around data protection (Delvaux, 2016, p. 8). This proposal is remarkable for both noting the confounding categorization of AI in relation to rights while also positioning contemporary legal rights frameworks squarely within surveillance capitalism’s version of data as licensed property and currency.

Persona rights in AI

Private law of contract, manifested in social platforms’ Terms of Service agreements, exerts significant pressure on contemporary versions of rights under surveillance capitalism. The whole idea of data as currency gets legally recognized through Terms of Service that most internet users do not even read (Klang, 2004), but that dictate the contractual relationship between the persona and the platform. According to the common stipulations of Terms of Service agreements, platforms commodify the content and personal information that users provide. As Lawrence Lessig (2006) explains, private law

of contract thus displaces government regulation on the internet, according to the profit motive (pp. 185-7). While users can exert some influence over the fairness of these contracts through traditional channels like the court system, or new technological channels like the development of alternative internet architectures, large technology companies still dominate. Further, while Terms of Service agreements vary by website and host country, many of the most popular sites globally are based in the U.S. (Jin, 2013). As such, state-sponsored policy around users' rights not only needs to consider the commercialization of persona, but the geopolitical imbalance of such control.

More broadly, a Western bias underpins the contractual relationship between platform and persona through the form of the Terms of Service agreement itself. As Carole Pateman (1988) argues, the way that one's identity or persona can generate proprietary interest under capital emerges from the libertarian legacy of self-ownership. The idea of owning oneself translates into alienable rights to "property in the person," which subjects the person, as private property, to a framework she calls "contractarianism."

Contractarianism contains an inherent paradox for Pateman, given that contracts binding the person – such as employment contracts – alienate the person-as-laborer's right of self-government while simultaneously demanding labor as a humanistic practice, a practice marked by self-government and autonomy (Pateman, 2002, p. 47). A parallel paradox is apparent in social media platforms' Terms of Service: platforms "retain licensing rights over user content and expression, amid celebrations of user agency and the democratization of cultural production" (Shepherd, 2012, p. 107). This disjuncture might be taken even further in the context of machine learning, where the contract users enter

into suggests full access to their private data and intellectual property as resources on which to train platforms to better personalize and meet user needs, or indeed, solve complex global problems (Zuckerberg, 2017).

Under surveillance capital, however, the idea of the individual persona is challenged by the way Terms of Service contracts facilitate the splitting up of individuals into dividuals – data derivatives that can be aggregated and reconfigured according to the pattern generating logic of big data (Amoore, 2011). In this context, individual rights within contractarianism become even more tenuous. For example, American legal frameworks for intellectual property have only strengthened corporate claims to digital culture, as in the Digital Millennium Copyright Act (1998), and Federal Communications Commission privacy laws were repealed in early 2017 by Congress and the Trump Administration. Given these kinds of legal challenges to the integrity of individual privacy and intellectual property, which only stand to increase in the context of AI, it seems important to attempt to reinstate the legal salience of the person – despite its shortcomings – as a means to reassert even basic versions of liberal rights.

As a legal concept, the term persona rights refers to an individual's control over commercial appropriation of their identity in social media platforms. McGeeveran (2009) explains how, within the context of behavioral advertising practices, persona rights law “transcends the narrower focus of other paradigms on protecting information privacy or preventing misleading advertising” (p. 1154). Persona rights opens up more familiar understandings of online privacy or intellectual property by focusing on the practice of

endorsement in social media platforms. Platforms like Facebook leverage their users' identities within algorithmic advertising models that deploy the "like" button (Gerlitz & Helmond, 2013), attempting to predict through machine learning what kinds of content or products will lead to greater user endorsement. For McGeveran, similar to celebrity endorsement, user endorsement might be subject to "two related but distinct legal claims: the tort of appropriation and the right of publicity" in U.S. law (McGeveran, 2009, p. 1149). By protecting against unauthorized commercial uses of one's identity (the tort of appropriation) and maintaining monopoly control over one's own image (the right of publicity), these two legal instruments serve the function of recognizing the integrity and dignity of personal identity, in light of commercial exploitation, as sanctioned by the state. Persona rights law thus rests on the premise that individual users should have recourse to legal and regulatory protection of their rights to control their personal identities online, even as they get reconfigured into data derivatives used as inputs for machine learning.

While the persona rights framework has its limitations, including maintaining neoliberal individualism, proving dignitary harm (McGeveran, 2009), conflating privacy with property (Samuelson, 2000), and legal jurisdiction of specific tort law (Hamilton, 2009), the concept of persona rights offers a reinstatement of personal integrity in a posthuman context where AGI threatens to displace the concept of identity even further under surveillance capitalism. So far, due to the lack of robust state protections, suggestions for how to tackle the problem of endorsement underlying commercial AI have largely come from within the industry as it seeks to avoid alienating user-consumers. For example,

Google has proposed a privacy solution that combines federated learning with differential privacy. Federated learning entails training AI systems with data that is aggregated rather than connected to an individual's identity (McMahan et al., 2017). This would be done through encrypted updates that prevent Google from seeing sensitive personal information, however, in practice, it would be possible to reverse engineer those updates to reveal the training data. Differential privacy is a strategy devised by Apple that applies complex mathematical techniques to render sensitive information unidentifiable. Yet the drawback of this approach is that it would also entail users sending more data to Apple than ever before (Greenberg, 2016). Apparently, the impetus of personhood rights to re-focus legal protection on the person will remain crucial as AI training systems develop.

From another perspective, moreover, personhood rights not only reinstates the legal salience of the person but the specificity of the human person within rights frameworks that are encroached upon by non-human or posthuman persons. As Peter Kahn and his co-authors suggest in their article "What is a human?" (2007), AI stretches the limits of humanity by being essentially determined by its ability to meet human psychological benchmarks.

AI's success, they argue, is reliant on mimicking human-like capacities through "categories of interaction that capture conceptually fundamental aspects of human life" (Kahn et al 2007, p. 366). One of these is creativity or the capacity of AI to generate what might be recognized as intellectual property. As the authors note, even if AI cannot be said to develop consciousness, the viability of any AGI would be fundamentally predicated on creativity in approaching a novel task using past learning unrelated to that specific task (Kahn et al., 2007, p. 378). In terms of privacy, AI represents accelerated,

immanent, and intimate data collection that stands to undermine the idea of contextual integrity all together (Kahn et al, 2007, p. 373; Nissenbaum, 2004). Finally, perhaps the ultimate benchmark for AI is autonomy, which speaks to the heart of posthumanism by acknowledging that autonomy for both humans and machines is subject to debate about the degree to which decision-making is conditioned by internal and external factors. In liberal formulations of rights, humans are invested with autonomy and by extension moral responsibilities that accompany basic rights of dignity and integrity (e.g., Dworkin, 1978). For Kahn et al. (2007), autonomy, or at least autonomous behavior, as a benchmark for AI thus raises the issue of rights beyond just persona rights of privacy and intellectual property. While the authors concur that such benchmarks for determining the success of AI are not meant as ontological categories but rather functional, psychological ones (Kahn et al., 2007, p. 365), the benchmarks nonetheless illuminate the constructedness of rights frameworks alongside the constructedness of human-like versions of privacy and intellectual property in AI.

Regulating AI

The question so far addressed in this chapter about the status of persona rights given the development of AI and AGI leads naturally to a consideration of legal protections for liberal rights. Especially around privacy and intellectual property, advancements in machine learning demand new frameworks for understanding the regulatory implications of applications already in use, such as assistant AI systems (Apple's Siri, Amazon Echo, Google Home) that collect ambient voice data and computational creativity systems like Adobe Sensei that automates design. Add to these the potential leaps in AGI – depicted

evocatively for example in Louisa Hall's recent novel *Speak* (2015), where sentient robots become friends for withdrawn children – and it seems that regulation has a difficult task ahead to keep pace with the diverse implications of technological developments.

Perhaps the greatest challenge in devising new regulatory frameworks is the opacity of the technology. AI systems are typically inscrutable in terms of how certain inputs result in outputs, in what has been termed the explainability problem (Heaven, 2013). This is apparent in the European Parliament proposal mentioned above, where the status of AI is itself unclear (Delvaux, 2016). So, even though the U.S. government has made attempts to leverage technical expertise (Stone et al., 2016, p. 10), for example in the task force advising on the May 2016 report of the National Science and Technology Council, “Preparing for the Future of Artificial Intelligence,” there are limits to regulating something that can't be fully understood even by developers. Nonetheless, the report suggests seven overarching mandates in AI policy, including that AI should be used for public good, bias must be eliminated from data, and global partnerships are necessary to ensure transparency.

The co-articulation of public good, bias, and transparency in this regulatory proposal suggests a link back to liberal conceptions of rights where individuals might expect fair and equal treatment under the law. And yet eliminating bias from AI is as tricky a proposition as eliminating bias from humans. Algorithms, it has been argued, are inherently ideological; they represent a “machinic subjectivity” that interpellates us

through our data (Jones 2014, p. 251). For instance, much has been said about the way AI systems enact a discriminatory reinforcement of existing social inequalities (Crawford, 2016; Leurs & Shepherd, 2016). Often these critiques rest on a supervised learning model where AI is fed data labeled by people and AI systems rest on invisible human labor (Gray & Suri, 2017). In reinforcement learning, however, unlabeled data is processed by algorithms and so it is not human but “algorithmic responsibility” that underlies discrimination (Banavar, 2016).

In order to police algorithmic responsibility without the need for restrictive government intervention, a number of industry-led initiatives have been proposed. This sort of self-regulation is seen to be potentially more effective than state policy due to non-specialist regulators’ difficulty in keeping up with the technology (Banavar, 2016, Stone et al., 2016). For example, The Elon Musk-backed OpenAI project, announced in December 2015, focuses on developing AI for public good. The Partnership on AI, a collaborative effort involving Google, Facebook, Amazon, IBM and Microsoft, was launched in September 2016. The founders of LinkedIn and eBay set up the Ethics and Governance of Artificial Intelligence Fund in January 2017 to support academic research on the safety of AI systems. These initiatives share a preoccupation with transparency around AI, matching calls from regulators and academics for third-party audits of algorithmic bias (Wachter et al., 2016). A “right to explanation” may help to ensure that industry-led initiatives live up to their promises for fair AI (Wachter et al., 2016), by providing clear explanations of what’s behind the black box of algorithmic decision-making so that people, even non-experts, can audit the process (Banavar, 2016, p. 4; Gray & Suri, 2017;

Pasquale, 2015). This suggestion could fit in with the social systems analysis that Kate Crawford and Ryan Calo (2016) support for uncovering the social, political, and cultural values embedded in AI.

While increased transparency is an important and laudable goal for regulators, transparency alone is not sufficient to address the larger regulatory implications of AI and AGI. Returning to algorithms' explainability problem, the posthuman conceit of AI suggests that these systems work like the human brain (which itself is something that neuroscientists don't yet fully understand). Yet AI does not work like the brain, even though it takes inspiration from the brain. Instead, developments in deep learning come from computational rather than neuronal processes, and moreover, these processes often have unpredictable outcomes. As legal scholar Ian Kerr (2017) has argued, if AI developers cannot fulfill their "duty to explain," then unexplainable AI should be prohibited unless individuals can take recourse in "a right not to be subject to a decision that is based solely on automated processing." This proposal pivots on explainability as the issue that underlies transparency and represents a relatively interventionist stance on AI development as socio-technical.

Such intervention may be necessary, despite attempts at industry self-regulation, given other, more longstanding problems with calls for transparency in the tech sector. To take transparency reports of the major technology companies as a prime example, the self-reporting done by industry players offer "a very particular kind of reporting, which may cater to demands for openness and disclosure about government surveillance and

“censorship, but provide a very specific response in a preformatted and selective shape” (Flyvverbom 2016, p. 8). Given that these same tech companies are behind recent advancements in AI, where machine learning is largely geared toward functions that can be monetized, this critique of transparency gains additional salience. Especially in the context of AGI, it is not difficult to imagine how it will be under attendant capitalist pressures to generate surplus value, continuing the imperatives of social media platforms to harness general intellect and flexible labor (Cohen, 2008; Jarrett, 2015). As the current limitations in processing power give way, machine learning will represent the “next frontier” for imperialist ambitions of tech companies to accelerate commodification. Perhaps the strongest limitations in this case will be philosophical. To return to the liberal political theory underpinning the idea of rights, the line between human and non-human or posthuman remains deeply invested by our own beliefs in the primacy of consciousness. Perhaps it is on these grounds that the persona can be recuperated along with its rights to self-determination.

Conclusion: individual vs. collective rights

This chapter has deployed the persona rights framework to help consider the double challenge of AI to liberal rights: 1) that AI systems appropriate individuals’ private data and intellectual property through machine learning algorithms; and 2) that the imminent development of AGI challenges the whole concept of persona through its posthuman conceit. Given the late capitalist context in which deep learning systems are emerging, where successful monetization of the technology rests on huge resources in data, persona rights offers something of a tactical response to maintaining human dignity and

sovereignty. Through control over the uses of one's endorsement – encompassing personal data and creative output – persona rights bolsters individual entitlement to control and use one's own persona (McGeveran, 2009, p. 1132).

Yet persona rights maintains a relatively liberal version of the individual as a coherent core over which sovereign control can be exercised – going against philosophies of technology that see the self as a social construct (Austin, 2010). Indeed, if the notion of a discrete self or persona might be challenged on philosophical grounds, it is also being undermined within digital culture as a society of control where individuals better describes the units through which people are made useful to machine learning as data derivatives (Amoore, 2011; Deleuze, 1992). The ways in which individual sovereignty fails as a precondition for rights suggests two paths forward: one in which the liberal individual is reconstituted, for example through persona rights law; and the other in which a wholly new paradigm for collective rights might be articulated.

The loss of the sovereign individual thus merits thinking more strategically about resisting the larger configuration of neoliberal individualism within global capital. One place to start this thinking is through a reconfiguration of the idea of private property. Letting go of the notion of control over the self as property opens up onto alternative frameworks of the commons, as suggested by free culture advocates (e.g., Boyle, 2008; Lessig, 2006), feminist political philosophers (e.g., Pateman, 1988), and Indigenous studies scholars (e.g., Stabinsky & Brush, 2007). The questions to ask from this vantage point would be how collective rights might help maintain human dignity in an AI era:

how can the colonization of general intelligence by tech companies be prevented? It seems particularly urgent to articulate posthuman rights before human and machine brains eventually merge in the envisioned singularity according to the efforts of technologists like Elon Musk. Given this urgency, what may look like anthropomorphism in the neural metaphors for AI, where machines are assigned subjectivity, can in fact portend a politics (Jones, 2014): the posthuman may be the privileged route into the postcapitalist.

References

- Amoore, L. (2011). Data derivatives: On the emergence of a security risk calculus for our times. *Theory, Culture & Society*, 28(6), 24-43.
- Austin, L. M. (2010). Control yourself, or at least your core self. *Bulletin of Science, Technology & Society*, 30(1), 26-29.
- Balakrishnan, S. (2016). Historicizing hypertext and Web 2.0: Access, governmentality and cyborgs. *Journal of Creative Communications*, 11(2), 102-118.
- Banavar, G. (2016). Learning to trust artificial intelligence systems: Accountability, compliance and ethics in the age of smart machines. White Paper prepared for IBM Corporation. Retrieved from:
https://www.research.ibm.com/software/IBMResearch/multimedia/AIEthics_Whitpaper.pdf

- Boyle, J. (2008). *The public domain: Enclosing the commons of the mind*. New Haven, CT: Yale University Press.
- Cohen, N. S. (2011). The valorization of surveillance: Towards a political economy of Facebook. *Democratic Communiqué*, 22(1), 5.
- Crawford, K. (2016). Artificial intelligence's white guy problem. *The New York Times*, 25 June. Retrieved from:
https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html?_r=0
- Crawford, K., & Calo, R. (2016). There is a blind spot in AI research. *Nature*, 538, 311-313.
- Damnjanović, I. (2015). Polity without politics? Artificial intelligence versus democracy: Lessons from Neal Asher's Polity Universe. *Bulletin of Science, Technology & Society*, 35(3-4), 76-83.
- Dean, J. (2005). Communicative capitalism: Circulation and the foreclosure of politics. *Cultural Politics*, 1(1), 51-74.
- Deleuze, G. (1992). Postscript on the Societies of Control. *October*, 59, 3-7.
- Delvaux, M. (2016). Draft Report with recommendations to the Commission on Civil Law Rules on Robotics. European Parliament Committee on Legal Affairs, 31 May. Retrieved from: <http://www.europarl.europa.eu/committees/en/juri/draft-reports.html>
- Domingos, P. (2015). *The master algorithm: How the quest for the ultimate learning machine will remake our world*. New York: Basic Books.
- Dworkin, R. (1978). *Taking rights seriously*. Cambridge: Harvard University Press.

- Gerlitz, C., & Helmond, A. (2013). The like economy: Social buttons and the data-intensive web. *New Media & Society*, 15(8), 1348-1365.
- Gillespie, T. (2010). The politics of 'platforms'. *New Media & Society*, 12(3), 347-364.
- Gray, M. L., & Suri, S. (2017). The humans working behind the AI curtain. *Harvard Business Review*, 9 January. Retrieved from: <https://hbr.org/2017/01/the-humans-working-behind-the-ai-curtain>
- Greenberg, A. (2016). Apple's 'differential privacy' is about collecting your data – but not *your* data. *Wired*, 13 June. Retrieved from: <https://www.wired.com/2016/06/apples-differential-privacy-collecting-data/>
- Hamilton, S. (2009). *Impersonations: Troubling the person in law and culture*. Toronto: University of Toronto Press.
- Haraway, D. J. (2008). *When species meet*. Minneapolis: University of Minnesota Press.
- Haraway, D. J. (1991), *Simians, cyborgs and women: The reinvention of nature*. New York: Routledge.
- Hayles, N. K. (1999). *How we became posthuman: Virtual bodies in cybernetics, literature, and informatics*. Chicago: University of Chicago Press.
- Heaven, D. (2013). Not like us: artificial minds we can't understand. *New Scientist*, 219(2929), 32-35.
- Jarrett, K. (2015). *Feminism, labour and digital media: The digital housewife*. London: Routledge.
- Jin, D. Y. (2013). The construction of platform imperialism in the globalization era. *tripleC: Communication, Capitalism & Critique*, 11(1): 145-172.

- Jones, S. (2014). People, things, memory and human-machine communication. *International Journal of Media & Cultural Politics*, 10(3), 245-258.
- Kahn Jr, P. H., Ishiguro, H., Friedman, B., Kanda, T., Freier, N. G., Severson, R. L., & Miller, J. (2007). What is a human? Toward psychological benchmarks in the field of human-robot interaction. *Interaction Studies*, 8(3), 363-390.
- Kerr, I. (2017). Speech before the House of Commons' Standing Committee on Access to Information, Privacy and Ethics. 4 April. Retrieved from: <https://techlaw.uottawa.ca/news/professor-ian-kerr-intervened-house-commons-discuss-pipeda>
- Klang, M. (2004). Spyware—the ethics of covert software. *Ethics and Information Technology*, 6(3), 193-202.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (In press). Building machines that learn and think like people. *Behavioral and Brain Sciences*.
- Le, Q. V., Ranzato, M., Monga, R., Devin, M., Chen, K., Corrado, G. S., Dean, J., & Ng, A. Y. (2012). Building high-level features using large scale unsupervised learning. *Proceedings of the 29th International Conference on Machine Learning*, Edinburgh, Scotland.
- Lessig, L. (2006). *Code: And Other Laws of Cyberspace, Version 2.0*. New York: Basic Books.
- Leurs, K., & Shepherd, T. (2016). Datafication & discrimination. In M. T. Schäfer & K. van Es (eds.), *The datafied society: Studying culture through data* (pp. 211-231). Amsterdam: Amsterdam University Press.

- McGeeveran, W. (2009). Disclosure, Endorsement, and Identity in Social Marketing. *University of Illinois Law Review*, 2009(4), 1105-1166.
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Aguera y Arcas, B. (2017). Communication-Efficient Learning of Deep Networks from Decentralized Data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, Fort Lauderdale, Florida.
- Nissenbaum, H. (2004). Privacy as Contextual Integrity. *Washington Law Review*, 79, 101-139.
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Cambridge, MA: Harvard University Press.
- Pateman, C. (2002). Self-Ownership and Property in the Person: Democratization and a Tale of Two Concepts. *The Journal of Political Philosophy*, 10(1), 20-53.
- Pateman, C. (1988). *The sexual contract*. Stanford, CA: Stanford University Press.
- Pennachin, C., & Goertzel, B. (2007). Contemporary approaches to artificial general intelligence. In *Artificial general intelligence* (pp. 1-30). Springer Berlin Heidelberg.
- Rawls, J. (2001). *Law of peoples*. Cambridge, MA: Harvard University Press.
- Sample, I. (2017). Google's DeepMind makes AI program that can learn like a human. *The Guardian*, 14 March. Retrieved from: <https://www.theguardian.com/global/2017/mar/14/googles-deepmind-makes-ai-program-that-can-learn-like-a-human>
- Samuelson, P. (2000). Privacy as intellectual property? *Stanford Law Review*, 52(5), 1125-1173.

- Shepherd, T. (2012). Persona rights for user-generated content: A normative framework for privacy and intellectual property regulation. *tripleC: Communication, Capitalism & Critique*, 10(1), 100-113.
- Stabinsky, D., & Brush, S. B. (Eds.). (2007). *Valuing local knowledge: Indigenous people and intellectual property rights*. Washington, DC: Island Press.
- Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenain, A., Shah, J., Tambe, M., & Teller, A. (2016). Artificial Intelligence and Life in 2030. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel. Stanford University. Retrieved from: <https://ai100.stanford.edu/2016-report>
- Terranova, T. (2004). *Network culture: Politics for the information age*. London: Pluto Press.
- Weston, J., Chopra, S., & Bordes, A. (2014). Memory networks. *Proceedings of the 3rd International Conference on Learning Representations*, 7-9 May, San Diego.
- Winner, L. (1977). *Autonomous technology*. Cambridge, MA: MIT Press.
- Zuboff, S. (2015). Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1), 75-89.
- Zuckerberg, M. (2017). Building global community. Facebook, 16 February. Retrieved from: <https://www.facebook.com/notes/mark-zuckerberg/building-global-community/10103508221158471/>