


IMAGINING an AI Commons



REPORT ASSEMBLED BY JANNA FRENZEL AND FENWICK MCKELVEY

Imagining an AI Commons

Report written by Janna Frenzel with Fenwick McKelvey and Bart Simon

Table of Contents

Acknowledgements	3
Participants	4
Foreword by Bart Simon and Fenwick McKelvey	5
1. Introduction	7
2. Common-pool resources	8
3. Commoning not commons	12
4. AI is already built on commons	17
4.1 Data and public knowledge	17
4.2 Networks	19
4.3 Labour	21
4.4 Energy	22
5. Conclusion	26

This report is published under the following license:

Attribution-NonCommercial-ShareAlike 3.0 Unported (CC BY-NC-SA 3.0)

Acknowledgements

The AI Commons Workshop took place at Concordia University located on unceded Indigenous lands. The Kanien'kehá:ka Nation is recognized as the custodians of the lands and waters on which we gathered. Tiohtià:ke/Montréal is historically known as a gathering place for many First Nations. Today, it is home to a diverse population of Indigenous and other peoples. We respect the continued connections with the past, present and future in our ongoing relationships with Indigenous and other peoples within the Montreal community.

We wish to thank our neighbours in Montreal at the Indigenous Futures Research Centre whose Indigenous Protocol and Artificial Intelligence Position Paper¹ offers another way into the study of the commons. Their work on Indigenous AI epistemologies inspires us to see this report as a point of contact between Western approaches to the commons discussed here and Indigenous approaches that we are still learning. The commons, as we have also learned in Montreal, is a practical concern as much about where we live as how we treat our data or our ideas. Another direction is to imagine AI commons as spaces where we may live and be, a different kind of smart city.

The workshop coincided with the National Day of Remembrance and Action on Violence against Women. The day commemorates the deaths of Geneviève Bergeron, Hélène Colgan, Nathalie Croteau, Barbara Daigneault, Anne-Marie Edward, Maud Haviernick, Maryse Laganière, Maryse Leclair, Anne-Marie Lemay, Sonia Pelletier, Michèle Richard, Annie St-Arneault, Annie Turcotte, and Barbara Klucznik-Widajewicz in the 1989 École Polytechnique massacre, an antifeminist mass shooting. The Government of Canada states that the day “is about remembering those who have experienced gender-based violence and those who we have lost to it. It is also a time to take action. Achieving a Canada free from gender-based violence requires everyone living in this country to educate themselves and their families and communities on gender-based violence, centre the voices of survivors in our actions and speak up against harmful behaviours.” We recognize that the commons is a feminist idea, and we hope our report celebrates the feminist scholarship behind its being in this world.

Our report draws on research supported by the Social Sciences and Humanities Research Council and the Centre for the Study of Democratic Citizenship.

¹ <https://www.indigenous-ai.net/position-paper/>

Participants

This report is a collective output of the Machine Agencies Working Group at the Milieux Institute at Concordia University. We wish to thank:

- Natalia Balska for designing this report;
- Janna Frenzel for writing and managing this report;
- Nick Gertler for shooting, editing, and producing our explainer video; and,
- Robert Hunt for editorial assistance on this report.

We also wish to thank all the participants of the Imagining an AI Commons Workshop, held on December 6, 2019. Unbeknownst to us, the workshop was the last time we would be able to gather safely for a long time. We are deeply thankful for our participants' time and engagement with our workshop.

Alexandra Bahary

Marc Böhlen

Ana Brandusescu

Katharine Dempsey

Carlos Denner

Stefanie Duguay

Janna Frenzel

Karine Gentelet

Robert Hunt

Aleksandra Kaminska

Alex Ketchum

Nisa Malli

Fenwick McKelvey

Jeremiah Miller

Lorna Roth

Ushnish Sengupta

Bart Simon

Luke Stark

Joe Zeph Thibodeau

Neal Thomas

Matthew Toews

Tricia Toso

Julian von Barga

Cody Walker

Ceyda Yolgörmez

Finally, we wish to thank our animators and keynote speakers Sophie Bishop, Brett Frischmann, Eda Kranakis, and, especially, Sandra Braman, who also offered very helpful edits on this final report.

For the Word

by Bart Simon and Fenwick McKelvey

Our workshop did not aim to build an AI commons. There already was one of those. The AI for Good Movement publicly launched the AI Commons organization in 2019. Intrigued by the organization's promise to "maximize the dissemination of our common knowledge of AI," we took its name as an opportunity to trouble the conjuncture of artificial intelligence (AI) or machine learning (ML) and the commons. The workshop brought together scholars interested in the commons as a start, to imagine what that term's various meanings and histories implied for artificial intelligence.

We began not out of cynicism but out of a desire to use the commons to shift the terms of engagement. That AI/ML will have and is having an impact on human and nonhuman forms of life cannot be denied. It makes sense that most of us would want that impact to be for the "common good." Unfortunately wanting to do good tells us little about what to do or how to do it. There are many "goods" and many perceptions of the common good. There is no direct route to a discussion of different or progressive technological politics with this rubric, and thus there is a danger of "ethics washing."

To imagine an AI commons was (and is) to imagine a different form of technological politics than what currently appears before us. In an era of platform capitalism and huge multinational monopolies in technology, we have to take this dream of an AI commons seriously, especially because it is shared by scientists and engineers and, yes, even entrepreneurs. Imagining an AI commons might be compared with the thinking behind the formation of Computer Professionals for Social Responsibility, which started in the early 1980s in part in resistance to the US's Strategic Defense Initiative ("Star Wars") program. Can our introduction to the commons here inspire a common purpose, a common good for AI?

A commons, we learned, is something ambiguous. Institutionally, commons can easily become cartels where a few players manipulate their shared ownership for private gain. For an AI commons to work, we think there must be more than the desire of individuals to be good—the commons is not a contract but a way of being with one another.

Rather than starting from notions of the privatized individual with rights and privileges as such and then figuring out what our common interests are (or, even worse, what is in the common interest), we might start

as commoners: those who live in/with a commons. One path we see might chart a convergence between AI/ML practices, the ethics of those practices, and an alternative ground on which those ethics and practices might be constructed. How might an AI commons be built and how might it support shared resources—data, models, and practices—through commons-based governance?

The one-day workshop was an experiment in thinking and doing, structured around provocations from experts on the commons followed by an afternoon dedicated to writing, which took the form of a sort of writing sprint around the idea of an AI commons. The outcome was meant to be this report. As all good experiments go, the workshop did not entirely work. Practically speaking, we began our workshop by defining what our own commons would be. What were the expectations of shared writing? How did participants contribute and benefit from our own intellectual commons? Perhaps we assumed a kopimist² attitude of the workshop. More tangibly, we assumed that a few hours of writing could result in anything more than inspiration. The product of the workshop lingered, then the pandemic hit. Our report then became something other than a product of the workshop.

Two years is a long time in the life of artificial intelligence; yet, the commons still matters, still stays with us. The commons also became a problem much more felt by us in Montreal, living with a housing crisis³ and a city that remains entangled in an economic fantasy of artificial intelligence. And as we watched empires of ethical AI crumble, unicorns die, and controversial applications of AI spread, our workshop on AI and the commons stayed with us. The concept matters.

Our report is a start, our attempt to add something to the commons. There is still so much work to be done. We offer the report as a beginning. An AI commons remains provisional, an idea in progress, and as such, the best idea to share now.

² Kopimism was one facet of the European piracy movement. Originating with the Pirate Bay, kopimists advocate for the abolition of copyright and other proprietary information, a stance guided by the overt movement against the concentration of intellectual property in the hands of a few owners.

³ Our work on the AI Commons happened simultaneously and overlapped with research into how the AI industry is transforming our home city, Tiohtià:ke/Montréal. Research here builds in part on the report by the Digital Divides Project, which analyzes the impacts of Montreal's AI "ecosystem" on the neighbourhood of Parc Extension, particularly in reference to the housing crisis. [The report can be accessed here.](#)

1. Introduction

There are a lot of commons. There is the common good and common knowledge. There are common concerns. There are also data commons: repositories of data sets that anyone can use. The idea of a commons initially seems to oppose private or proprietary ownership, competition, protectionism, secrecy, enclosure, and monopoly. An AI commons invokes the early free and open-source software movement that celebrated the commons from often contradictory positions. These popular usages of the commons might be seen as a refraction of the canonical definition of the commons.

This report builds on the workshop to capture the many meanings of the commons. At the workshop, we learned about institutional definitions of commons and information commons. Our first section focuses on those institutional arrangements of AI that might become a policy agenda to address growing concerns about ownership and corporate power.

We have added our own interest in commoning, i.e., thinking about being in common. As we discuss in the report, we find a multitude of commons thinking. Data, networks, labour, and the raw materials used in hardware are already entangled with AI. The second part of our report brings forward what commons might be found or made through artificial intelligence.

We emphasize the friction between the existing commons that enable AI today and the institutional frameworks that might legitimize them to bring about better future commons by acknowledging the commons that already are.

2. Common-Pool resources

The commons is an underappreciated institutional approach to shared resources. Commons have conventionally been understood as something physically tangible: a natural resource shared among a group of people (Hess and Ostrom 2007). The term originated as a legal arrangement in feudal England and described an area of uncultivated land that could be used by “commoners” for purposes such as grazing livestock, fishing and hunting, collecting firewood and building materials, etc.⁴ Starting in the late Middle Ages and continuing into the eighteenth century, landlords enclosed these commons, thereby expropriating small farmers and destroying these arrangements of collectivity (Dyer-Witthford 2002:130). These commons have been lost, yet this approach to the commons endures as a way of thinking against narrow definitions of private property.

Unfortunately, contemporary debates about the commons are all too often associated with the anti-immigrant and white nationalist ecologist Garrett Hardin and his discredited “Tragedy of the Commons.” Hardin posited privatization and white separatism as solutions to alleged overpopulation and resource scarcity due to collective overuse (Mildenberger 2019; Southern Poverty Law Center n.d.). The commons has been lost both historically and in popular thinking, where all too often Hardin’s racist conceptualization overshadows the wealth of research and theorization of the commons.

Nobel Prize for Economics winner Elinor Ostrom debunked the “tragedy of the commons” claim. While acknowledging that there is a risk of overuse of common-pool resources, Ostrom (2015) demonstrated that collectively sharing a resource does not inevitably lead to its depletion—indeed, appropriate collective governance strategies can successfully maintain a commons. She developed eight principles of common-pool resources that demonstrated the effectiveness of commons as an institutional arrangement.

In this report, we recover some of the many meanings of the commons.

⁴ <https://www.britannica.com/topic/commons>,
see also <https://thelandmagazine.org.uk/articles/short-history-enclosure-britain> and
https://www.sup.org/books/title/?id=31468&local_ref=new

The Eight Principles of Common-Pool Resources

Through her empirical work around the world, Ostrom distilled eight principles for the management of what she called “common-pool resources” (CPRs) that would “account for the success of these institutions in sustaining the CPRs and gaining the compliance of generation after generation of appropriators to the rules in use” (2015:260):

1. Define clear group boundaries.
2. Match rules governing use of common goods to local needs and conditions.
3. Ensure that those affected by the rules can participate in modifying the rules.
4. Make sure the rule-making rights of community members are respected by outside authorities.
5. Develop a system, carried out by community members, for monitoring members’ behavior.
6. Use graduated sanctions for rule violators.
7. Provide accessible, low-cost means for dispute resolution.
8. Build responsibility for governing the common resource in nested tiers from the lowest level up to the entire interconnected system. (Walljasper 2011; based on Ostrom 2015)⁵

Since Ostrom first formulated her principles, the idea of the commons has inspired many to research and write about and to practice and envision better collective futures that centre around ethics of sharing, care, and mutual responsibility. Given the vastness of engagement with the commons, the works and examples we draw on here can only provide small glimpses into these debates. We hope, however, to lay out some key lines of thought we find particularly inspiring for imagining what a commons could be in reference to AI.

⁵ For the original list of Ostrom’s eight principles with more detail, see Ostrom 2015:260-262.

The commons is central to the operation of many sociotechnical systems today from airspace to radio frequency, even if regulation of the latter largely relies on auctions and private allocations. At our workshop, participant Eda Kranakis defined a commons as:

Resources always exist in certain kinds of spaces that have to be taken into consideration. A commons is essentially the combination of a resource, its particular spatialities, and a governance system. Any AI commons is going to involve a certain geographical space. This space could be diffused but connected through infrastructure systems. The question is: is this going to be a nationally limited space or a transnational space? And what objective does this commons have? I don't think it is possible to develop a generic commons because the issues are so different depending on what kind of a problem you are trying to solve.

This insight points to the complications that arise when trying to define a commons. Since commons are multiple and have noncontiguous boundaries, there can never be one clear-cut path to how a governance system should be arranged in relation to a specific commons, as illustrated in the quote in reference to infrastructure space versus geographical space. Neither should the commons be analyzed as goods or thought of as an archaic, outdated form of social organization that belongs to the past. Instead, commons evolve with society; with “new” commons decidedly shaped around technology, networking, and information, just like industrial and postindustrial society itself (Disco and Kranakis 2013). By drawing on examples such as airspace, weather-forecasting systems, the genetic diversity of crops, and the sea, Disco and Kranakis diagnose an amplification in scale from the local and a transnationalization of commons that involves “not only large-scale collective actors but also more abstract, multifaceted terms of access and use” (2013:2). They suggest calling these shared infrastructures and resource-spaces “cosmopolitan commons,” which represent “building blocks of regional integration and globalization” but are simultaneously always threatened by subversion or the clashing interests of the actors involved (2013:3).

Commons and the related concept of enclosure have figured prominently in making sense of the history of the internet and debates about such issues as regulation, platformization, and copyright and intellectual property. In its early days, the internet was often imagined as a collaborative medium that allowed for new forms of resource and knowledge sharing, communication, and relationship building, unregulated by any form of state authority or private entity—in other words, a type of commons. With the advent of Web 2.0, the platform business model that relies on user data sold to third parties for marketing purposes put an end to this idealization. In 2003, Dan Hunter described the shift as an enclosure via “undesirable private control of the previously commons-like Internet” that would result in a situation where “multiple parties have an effective right to prevent others from using a given resource, and as a result no one has an effective right of use” (2003:441). Hunter calls this an “anticommons,” bemoaning the decline of a democratic internet that serves as a public intellectual space accessible to all.⁶

The internet today is far from the knowledge commons that inspired some of its early adopters and advocates. We should, however, not forget the many successful knowledge commons that do exist online. Wikipedia can be complex and bureaucratic, but the output is a global encyclopedia shared as a public good. Though institutionally fraught, GNU General Public Licenses and other free software movements have built a public infrastructure that runs much of the internet. However, we have far to go. We finalized this report as Elon Musk announced plans to purchase Twitter, leaving some users feeling trapped and reluctant to move to alternatives like Mastodon—a reminder of the powerful lock-in effects of social media platforms.

In any case, first you have to think about what the relevant resources are, then identify what the relevant communities are. And you can't talk about "AI commons" in some kind of universal, abstract way. In general, one should beware of thinking about commons as a panacea, some sort of generalizable solution. The management of different resources within an AI system as a commons is going to look different depending on whether it is employed in a hospital, a school, a transportation system, and so on. You can go through a whole range of contexts and the possibilities and opportunities for governing the resources involved will be completely different: different sets of social dilemmas, different sets of objectives and goals, different social norms within communities. Plus there's also the issue of scale; some AI commons might involve a bigger role of government, while others might be more reliant on actors from within the communities.—Brett Frischmann

The commons then is a starting point for investigating regulatory frameworks for AI—a move that brings us to practical matters of data trusts, community trusts, and creative commons enabling the data mining behind today's machine learning and deep learning.⁷ Sanda Braman reminds us to also consider what types of commons might meaningfully improve the conditions of production of AI. Could training data commons be created to incentivise better, more just models? How could these data commons ensure obligations like the General Public License to give back and contribute to common-pool resources?

⁶ It is important to note here that some of the internet-as-commons enthusiasts, including Hunter, as well as some actors who are driving the more recent libertarian hype around cryptocurrencies are deeply entangled in a colonial mindset. This is particularly obvious in their usage of descriptions and metaphors that delineate cyberspace as a frontier, a sort of virtual "Wild West" where pioneers are allegedly free to roam, "explore," and forge new modes of sociality outside of centralized control and regulation. Such rhetoric conceals the extremely violent nature of historical and contemporary frontiers, their implication in the dispossession and genocide of Indigenous peoples, and their role in claiming and preparing both physical and digital spaces for capital accumulation. The (Western) concept of the commons does not necessarily align with a politics of decolonization, even if it is directed against corporate enclosures as, e.g., in the case of Occupy Wall Street (see Barker 2012; Fortier 2017). For critical engagements with data colonialism—and visions for data sovereignty—see, for example, Couldry and Mejias (2019), Kukutai and Taylor (2016), Milan and Treré (2019), and Thatcher, O'Sullivan, and Mahmoudi (2016).

⁷ Since the workshop, the Ada Lovelace Institute and the AI Council have published a report on data governance institutions for AI. See: AI Council and Ada Lovelace Institute. 2021. Exploring Legal Mechanisms for Data Stewardship. Ada Lovelace Institute. <https://www.adalovelaceinstitute.org/report/legal-mechanisms-data-stewardship/>

3. Commoning not commons

While a definition of the commons as a shared resource-space governed by institutional arrangements already moves away from commons as goods, other conceptualizations have brought collective practice, or acts of commoning, into focus. Rather than understanding the commons as a resource to be collectively exploited, the verb commoning describes a way of being: of living on and in the commons.

AI (Artificial Intelligence) is recycled human intelligence. Current methods depend exclusively on examples provided by humans, human knowledge, and do not extrapolate well to unseen circumstances.—Matthew Toews

Other approaches to the commons suggest an alternative perspective on it: not as an institutional arrangement but a material entanglement, living in close connection with others and the land without a presupposition of ownership. Would there be different things to say if we spoke of commoning AI rather than an/the AI commons?

We might start by asking who, or what, is a commoner. The condition is certain to historian Peter Linebaugh, who has proposed understanding commoning as a practice and an activity that is as much about sharing as reciprocal obligations. Linebaugh describes the commons as follows:

The opposite of the commons is the commodity. In the commodity, the social relations of creation and the social relations of subsistence are hidden. The commodity is about production. The commons, by contrast, are about reproduction. . . . Its principles are not those of the commodity, not those of accumulation, but those of subsistence and health. (Volont 2018; emphasis added)

Commoning puts relations between humans as well as relations between humans and nature front and centre. The key issue here is not one of ownership (as, for example, in reference to land) but of mutual subsistence (Volont 2018).

I think the idea of an AI commons is useful as a lens to think about the possibilities for people who make up AI training data. Machine learning draws lines around content creators and creates connections between them. How could these AI systems be repurposed as a form of commons, for example by influencers? For instance, algorithms that are used for marketing map out the connections between content creators on social media. So if this work is already being done by commercial entities, how could influencers use it for organizing and creating communities of care in order to counter the exploitative and individualizing employment of such tools?—Sophie Bishop

Commoning redefines labour as a human mutuality, as opposed to an exploitation or an exchange. Discommoning, then, describes the removal of a source of subsistence from a community that collectively manages it—and disrupts or destroys the relations that sustained this practice of commoning. As such, privatization is a form of discommoning. The struggle for the commons, or for the ability to practice commoning, essentially denotes the struggle against domination and exploitation (Volont 2018). Building on this potential of a common for all, Nick Dyer-Witheford, Atle Mikkola Kjosen, and James Steinhoff suggest that AI could be “a post-revolutionary ‘communal utility’” involving “working-class steering of AI development, for the involvement of workers and communities in determining what sorts of work should or should not be automated, and thus for a genuine determination by the ‘general intellect’ as to the design of AI” (2019:154). How these conditions come about is precisely a question raised by an AI commons.

What is more “common” right now is that we are the commons, as data producers or sources of data that all the different AI programs are using. And that puts us into a very different position when we try to negotiate with AI. I do think that there’s a lot about digital technologies that is qualitatively new. Informational meta-technologies—digital technologies and biotechnology—are different from industrial technologies. Our policy tools, our techniques of governance, are industrial in the sense that we understand things to flow in a causal stream in which you can intervene and which you can divert to make it do something else. But in the digital environment, the flows constantly recreate new causal chains. I think we are still learning about the affordances that digital technologies provide, including autonomously evolving software where even the producers do not fully understand what is happening. That’s why I use the phrase “negotiating with AI” rather than just focusing on efforts to control it. We should be looking for ways to ensure that communities of software programs have cognitive complexity and diversity and not be monolithic in their “thinking.” Only complex adaptive systems will allow us as humans and the rest of the natural world to survive and thrive.—Sandra Braman

Commoning is central to certain strands of feminist thought. Silvia Federici has observed that “‘commoning’ of the material means of reproduction is the primary mechanism by which a collective interest and mutual bonds are created” (n.d.). Since women have historically had to take on the vast majority of reproductive work and are still doing so today, they are especially dependent on communal natural resources and practices of commoning—and as a group often take the frontlines in struggles for retaining or reclaiming the commons. Understood through a feminist lens, commoning the material means of reproduction is the key mechanism to creating a collective interest and mutual bonds. At the heart of the idea, there is a challenge to create global bonds of solidarity:

No commons is possible unless we refuse to base our life and our reproduction on the suffering of others, unless we refuse to see ourselves as separate from them. Indeed, if commoning has any meaning, it must be the production of ourselves as a common subject. This is how we must understand the slogan “no commons without community.” But “community” has to be intended not as a gated reality, a grouping of people joined by exclusive interests separating them from others, as with communities formed on the basis of religion or ethnicity, but rather as a quality of relations, a principle of cooperation and of responsibility to each other and to the earth, the forests, the seas, the animals. (Federici n.d.)

Essentially, commoning then means both collectivizing the means of reproduction as well as building and cultivating relations of care. Following Maria Puig de la Bellacasa, “care is everything that is done . . . to maintain, continue, and repair ‘the world’ so that all . . . can live in it as well as possible” (2017:3; based on Joan Tronto). At the same time, Johanna Dahlin and Martin Fredriksson have proposed to think of resistance to enclosure and extraction as “working in common,” a form of commoning (2017:267). In this understanding, commoning does not necessarily need to include “productive” acts of relation building and the provision of care; it can also simply denote a collective refusal that pushes back against “the fundamental impulse of capitalism to enclose the commons” (Veltmeyer and Bowles 2014:66).

Commoning Care

“We live in a world where captains get arrested for saving people’s lives on the sea; where a person downloading scientific articles faces 35 years in jail; where people risk charges for bringing contraceptives to those who otherwise couldn’t get them. Folks are getting in trouble for giving food to the poor, medicine to the sick, water to the thirsty, shelter to the homeless. And yet our heroines care and disobey. They are pirates.”

<https://syllabus.pirate.care/#care-a-political-notion>

The creators of the Pirate Care Syllabus bring attention to how the organization of social reproduction has become a battleground of privatization and deregulation in the last few decades. Two trends have emerged: a market-oriented response that pushes digital technologies as individualized “fixes” for declining public health care provision (e.g., health and disease prevention apps) and a turn towards commoning care in the form of cooperativism, redistribution of material and immaterial resources, and other practices that “[position] care within specific forms of situated, embodied practices tinkering with technologies . . . expressing a transformative vision through commoning wealth and health” (Pirate Care n.d.). The latter response has often been restricted, scrutinized, and criminalized.

The Pirate Care Syllabus continues the tradition of crowdsourced online syllabi generated within social justice movements. For more information and inspirations for commoning care, see <https://syllabus.pirate.care/topic/piratecareintroduction/> and <https://syllabus.pirate.care/topic/commoningcare/>.

Bringing a perspective drawing on Black studies to bear on the two primary modalities of machine learning, Beth Coleman invokes Fred Moten and Stefano Harney’s undercommons to question all that surrounds AI. The undercommons, as best we can summarize here, invokes the wilds at odds with institutions and logics of governance. Coleman uses the concept to question AI’s potentialities because the undercommons belongs to “a long tradition of disruptive positionalities that abandon binaries such as master/slave, subject/object, and society/nature” (2021:10). The commons, or rather undercommons, looks for what is and what could be. Wilding AI is, Coleman concludes, an “opportunity to look at artificial intelligence—the machinic making sense of—as a process of ongoing relations, as phenomena as opposed to ‘knowledge’ represented in a database” (2021:15). The undercommons and its concept of the surround helps to recognize what might be ignored, unexpected, and colonized in AI as both a relation of power and an opportunity for radical new directions.

In a similar vein, Anna Tsing draws the connection between the undercommons and value creation or asset making, arguing that an exploitation of the former is indispensable to the latter:

Private assets most always grow out of unacknowledged commons. . . . Privatization is never complete; it needs shared space to create any value. That is the secret of property’s continuing theft—but also its vulnerability. . . . Even as entrepreneurs concentrate their private wealth through building alienation into commodities, they continue to draw from unrecognized entanglements. The thrill of private ownership is the fruit of an underground commons. (2015:271, 274)

Tsing's points resonate with how AI development and deployment is the latest iteration of the monopoly creation and enclosure by "Big Tech" that drives the largely unregulated growth of the industry. Pieter VergeDEM (2022) has suggested conceptualizing AI as a General Purpose Technology—that is, a technology that has the capacity to alter society's economic and social structures—and which therefore requires alternative regulations of ownership and governance to the ones that currently exist.

AI is an act of commoning, a political technology entangling together data worlds. Imagining an AI commons invites investigating the power dynamics of entanglement as well as possible alternative relations. The CoPilot Initiative by Microsoft GitHub and OpenAI is a good example of the problematic structure of these entanglements. Taking the world's largest repository of code—much of it free software—the two large AI firms trained a code auto-complete program. Meant as a convenience, CoPilot steered what might be seen as a commons into a new software service. The service might be a step in de-skilling coding, further positioning AI as dead labour meant to replace living labour (Belton 2021). GitHub demonstrates what is at stake in AI and commoning; its success comes from becoming a networked infrastructure built around public libraries and repositories. Here we take inspiration from Dimitris Paradopoulos's idea of insurgent posthumanism when he asks: "to read worlding not only as an opening to other social ways of being, but to other material processes as well. What does it mean to world justice today if not to enact openings, to build associations, to craft common, alternative forms of life?" (2010:148).

4. AI is already built on commons

Communication always involves something common, but commonness is also always a site for potential commodification and enclosure. Even though we don't share the idealized narrative of pre-Web 2.0 internet history as Hunter lays it out, we do acknowledge that his analysis is helpful for understanding the new waves of "enclosures" happening in the field of AI today. Just like corporate social media platforms monetize online socialities and modes of interaction, proprietary machine learning algorithms and other types of AI build on manifold "commons" without which they would not be able to exist.

In the following sections, we address some of these "underground" commons (Tsing 2015): data, labour, networks, and energy.

Commons have to be constructed. After you have discerned a potential resource space or a problem that involves a resource space that needs to be solved, you have to create this commons. And that's always an economic, political, and social problem. And the most important question is, who gets a seat at the table when it is decided what commons is created?—Eda Kranakis

4.1 Data and public knowledge

Artificial intelligence builds on various commons in a way that is consistent with Anna Tsing's (2015) observation that any proprietary commercial good relies in some way on exploitation of a commons. At its core, AI (or more specifically machine learning) is based on the availability of large sets of data, which are then used to train an algorithm to perform specific calculations. Since the early days of the internet, a vast amount of open data has been produced that AI development now profits from. As Matteo Pasquinelli and Vladan Joler note,

Mass digitalisation, which expanded with the Internet in the 1990s and escalated with datacentres in the 2000s, has made available vast resources of data that, for the first time in history, are free and unregulated. A regime of knowledge extractivism (then known as Big Data) gradually employed efficient algorithms

to extract ‘intelligence’ from these open sources of data, mainly for the purpose of predicting consumer behaviours and selling ads. (2020)

Put differently, AI becomes profitable by using data, the “first source of value and intelligence” (Pasquinelli and Joler 2020), which comes into existence through the labour and social cooperation of countless humans producing content and interacting with each other online as seen in the GitHub example above. Algorithms, in turn, “compute such value and intelligence into a model” (Pasquinelli and Joler 2020). But unlike what the dominant narratives around AI may suggest, this process is not just a technical one but a cultural and social one, since a specific dataset can only ever be a selection with specific biases built into it (see, e.g., Eubanks 2018; Gitelman 2013; Onuoha 2018).

Corporate AI appears as a new wave of enclosures, where human sociality, cooperative production, and knowledge sharing are extracted for profit. As Kate Crawford puts it, “the new AI gold rush consists of enclosing different fields of human knowing, feeling, and action—every type of available data—all caught in the expansionist logic of never-ending collection. It has become a pillaging of public space” (2021:120–21). This stands in stark contrast to visions and practices of a “digital information commons,” which once seemed to herald a sweeping democratization for the production and sharing of “useful knowledge” due to its capacity for breaking down the barriers between closed professional spaces and allowing for new citizen-led practices of online crowdsourcing to emerge (Mansell 2012b). The idea of an information commons presumes that “no one uses exclusive rights to organize effort or capture its value, and . . . cooperation is achieved through social mechanisms other than price signals or managerial directions” (Benkler 2004:1110). Such was the case, for example, with the database of digital images compiled under Creative Commons licenses that was later appropriated for developing and selling surveillance technologies based on facial recognition (Pasquinelli and Joler 2020).

The corporate AI enclosure of digital commons does not stop at extracting previously communal knowledge and sociality; it also removes the models it builds from the people whose data it relies on and whose lives are impacted by the model’s deployment. Unlike other statistical models, such as the ones produced to understand climate change, AI models are supposedly “blackboxed”: “A dataset may still be publicly available, but the metavalue of that data—the model created by it—is privately held” (Crawford 2021:120). Reading corporate AI as an enclosure thus also reveals how practices of commoning are being actively foreclosed. As proprietary algorithms hide extracted data in a protected machinic mechanism, they become inaccessible to public scrutiny and debate: “Given the degree of myth-making and social bias around its mathematical constructs, AI has indeed inaugurated the age of statistical science fiction” (Pasquinelli and Joler 2020; emphasis in original).

AI as One of Us or: It Takes a Village to Raise a Child

Torpignattara is a neighbourhood in Rome, Italy. Recently, the community welcomed a new member into its fold: [IAQOS](#), an artificial intelligence baby conceived by artists Salvatore Iaconesi and Oriana Persico.

“The message was clear from the beginning: AI was not to be understood as yet another technical infrastructure or service in the neighbourhood, something that others beyond us would maintain and offer to us under some form of service or utility. . . . What we had in mind was a new form of human/not-human community. Or, in other words, the beginning of a new possible queer-extended-neighbourhood family. . . .

We brought the notion of the newborn AI in the neighbourhood to the streets and to bars, laundries, grocery stores, through ‘IAQOS boxes’: boxes in which people could put notes to express what, in their opinion, an AI coming to the Torpignattara neighbourhood should absolutely know, as a knowledge-gift for the artificial kid. We received hundreds of notes which contained anything from requests to help humanity with climate change; advice about trans-escorts living in certain streets in the neighbourhood; poems; expressions of joyfulness for the upcoming birth; gossip; coordinates for remarkable or problematic places in the neighbourhood, and much more, in dozens of languages. We fed them to the linguistic AI, and it progressively formed its knowledge base.

It was useless, from a utilitarian point of view. You could have never used this sparse information to create a smart-city service. But it was exactly what was needed to promote the idea of a different type of actor in the city, that was growing with you, reachable, negotiable, that you can agree or disagree with, and with whom you can establish multiple types of negotiable social contracts based on relation and co-existence.” (Iaconesi and Persico 2021)

4.2 Networks

The popularity of today’s internet relies on a confusion, and certainly an aspiration, about the information commons. A decade ago, Robin Mansell noted that,

The evolution of the communication system is being guided by the economic power of the largest companies, by the political power of nation states and their institutions, by a host of other actors organized formally

into non-governmental organizations, and by individuals in their capacities as consumers and citizens in the information society. Developing a better understanding of how the power relationships guiding the technological innovation process come to be aligned with a vision of an information society that emphasizes private ownership of information and commercial markets is a goal of my analysis of the paradoxes of information scarcity and complexity. Critical examination of the prevailing vision serves as a means to understand possibilities for reconciling the goal of economic growth with the goals of social justice and the equitable distribution of resources. The evolution of the communication system is being guided also by a vision that is increasingly prevalent among members of online communities (of scientists, engineers, computer scientists, activists, and everyday users of the Internet) who favour social relationships that privilege the values of online cooperation, non-market information sharing, and common ownership of information. (2012a:25)

The internet, and AI, then are a point of struggle between two approaches to networking and communication. Sometimes, maybe as an unfortunate legacy of Hardin's "tragedy," the term commons is still employed as if it designates a free-for-all resource and a complete absence of regulation. As Dahlin and Fredriksson remind us, this is not at all what the commons is about. A commons is not an open access resource but "a resource shared by a group of people, often subjected to particular social norms that regulate how it can be used" (Dahlin and Fredriksson 2017:253). However, libertarian thinking reverberates in some open access, free software, or libre content movements.⁸ In these circles, open networks like the idealized internet have long been treated as the central pillar of an "information commons" and understood as the antithesis to hierarchically organized social spaces and information systems. But as Geert Lovink (2006) points out, elevating the idea of open networks to the gold standard of "good" online practice also fulfills an ideological function in that it conceals existing power structures and exclusion mechanisms, often leaving unclear what "open" really means. Instead of thinking in an open=good / closed=bad dichotomy, Lovink suggests, open networks should be regarded as distributed systems where it is possible to practice a "protological openness," which we interpret here as a contextualized, collective, and carefully worked out structure of common norms and regulations that evolves together with the social group that constitutes its user base.

At an infrastructural level, the internet has also been a site of constant experimentation with commons-based infrastructure, including early attempts at building wireless commons. In Canada, these efforts have been less successful, though Montreal's *Réseau Libre*, an informal grassroots group made up of volunteers, community organizations, university staff, and local businesses, attempted to build a city-wide free mesh network.

⁸ For a comprehensive analysis of free software culture and its related approaches, see Kely (2008).

Growing infrastructural intelligence—or what McKelvey (2018) calls the internet’s possession—reconfigures the traditional presumed ignorance of networks over what they transmit. Internet infrastructures now learn from the data they transmit, principally through behavioural profiling in cybersecurity, but more directly as a source of power itself. The British SCL Group—predecessor to Cambridge Analytica—together with Canada’s AggregateIQ tried to buy bulk data of citizens’ browsing histories from a major internet service provider in Trinidad and Tobago. The alleged goal was to build psychological profiles of voters so that a client, a Trinidadian political party, could improve their election results (Freeze and MacKinnon 2018). Whether the scheme did or could have worked is less important here than illuminating the deep infrastructural pressures to commodify collective communications as a means to train AI. The deeper issue is whether present intellectual property regimes will permit the collective production of networked cultures to be left unprotected or, at worst, allowed to become a proprietary AI.

In tandem, online sociality—especially in regard to content on social media platforms—has also become a resource to build recommendation engines via machine learning algorithms, sentiment analysis applications via natural language processing, or image recognition applications. Unlike, e.g., crowdmapping or other forms of grassroots networking for specific causes, the extraction of sociality through (corporate) AI arguably does not contribute to a collective distributive intelligence. To practise the aforementioned “protological openness,” user-producers would determine together what data is collected, how it is collected and stored, and for what purpose it would be used—i.e., they would engage in commoning. Instead, AI applications that “mine” online sociality are bound up in the for-profit business model of the respective platforms, turning it into an enclosed resource that is not a commons anymore.

4.3 Labour

As already mentioned in reference to data production, the AI industry would not be possible without human labour, despite prevailing narratives by corporations and other AI enthusiasts about automation and machine “intelligence.” Over recent years, different authors have pointed to the centrality of “ghost work” within AI systems, i.e., the invisibilized labour performed by humans (e.g., Gray and Suri 2019). Whether cleaning datasets, listening to audio snippets recorded by smart speakers to improve their responses, or completing one task after the next on platforms like CrowdFlower, Clickworker, Toluna, or Mechanical Turk, humans are involved in all steps of the AI “assembly line” (Pasquinelli and Joler 2020). Companies who employ machine learning in retail, marketing, and customer service simultaneously rely on a new economy of human contract labour. AI-assisted services such as Amazon’s logistical operations are implementing such new forms of hybrid human-robotic labour practices:

AI is built on the very human efforts of . . . crowdwork, the privatization of time, and the seemingly never-ending reaching, lifting, and toiling of putting boxes into order. From the lineage of the mechanized factory, a model emerges that values increased conformity, standardization, and interoperability—for products, processes, and humans alike. (Crawford 2021:57)

As AI implementation becomes more ubiquitous, old markets for human labour are destroyed while new ones for “humans-in-the-loop” emerge. Gray and Suri call this the “ever-moving frontier of AI’s development, the paradox of automation’s last mile” and argue that “AI will eliminate some work as it opens up opportunities for redefining what work humans do best” (2017; emphasis in original). This observation calls into question AI’s purported goal of full automation. As AI-assisted systems of extraction and surveillance are transforming workplaces, these systems are increasingly becoming the focal point of labour organizing (Crawford 2021:86).

Commoning Molecules


Why is hormone pollution from farming biocides, industrial meat production, or plastics largely thought of as normal and acceptable, while access to gender-affirming hormones for trans people is often highly contested, prohibitively regulated, and thought of as “unnatural”? This question is at the centre of Maddalena Fragnito’s article “[Commoning Molecules: Decolonising Biological Patents by Gender Hacking Protocols](#).” Fragnito looks at DIWO (Do-It-With-Others) workshops that embrace biohacking as a way of commoning scientific and chemical knowledge, thereby offering an alternative to the patent-based, “bio-capital market system” (2020).

“Methods [like DIWO biohacking workshops] hold the possibility of making sure that potential scientific benefits are available to everyone instead of letting science production become a tool of oppression. Moreover, they promote awareness of the overall toxicity related to the hormone molecules invasion on our bodies and territories.” (Fragnito 2020:167)

4.4 Energy

Training an AI model requires large amounts of computing power and, hence, electricity. Only recently have computer scientists, social scientists, and humanities scholars begun to explore the implications of this part of the “AI assembly line” in detail. In one of the first attempts to account for the climate impact of

AI, Emma Strubell, Ananya Ganesh, and Andrew McCallum studied common natural language processing models. They estimate that training one transformer (a type of deep learning model) with neural architecture search on a general processing unit produces roughly the same amount of CO2 emissions as building and driving five cars over their entire lifetime (Strubell, Ganesh, and McCallum 2019).



The total amount of compute in petaflop/s-days (Source: Amodei and Hernandez 2018, <https://openai.com/blog/ai-and-compute/#fn2>)

As AI models grow continuously larger, they become more computationally intensive. As more parameters, or weights, are added to neural networks, electricity and other computing-related resource consumption increases (Saenko 2020). It is estimated that the computation capacity required for AI training runs increased 300,000 times between 2012 and 2018 (Amodei and Hernandez 2018). Since larger models are more accurate than smaller ones, AI developers have largely adopted model accuracy as the sole measurement of success, which leads to a situation where “the focus on this single metric ignores the economic, environmental, or social cost of reaching the reported accuracy” (Schwartz et al. 2019). Schwartz and colleagues refer to this approach as “Red AI” and propose a “Green AI” approach that would elevate efficiency as an at least equally important criterion for AI development. Considering the lack of data and reporting mechanisms that would be necessary to systematically assess the energy consumption of machine learning, Henderson and colleagues (2020) have suggested a framework they call “experiment-impact-tracker.” This framework is supposed to facilitate “consistent, easy, and more accurate reporting of energy, compute, and carbon impacts of ML systems” (Henderson et al. 2020).

Computer scientist Kate Saenko concludes,

Unless we switch to 100% renewable energy sources, AI progress may stand at odds with the goals of cutting greenhouse emissions and slowing down climate change. The financial cost of development is also becoming so high that only a few select labs can afford to do it, and they will be the ones to set the agenda for what kinds of AI models get developed. (Saenko 2020)

Of course, the resources that AI development and implementation consumes are not limited to energy alone. As has been well documented in works such as Jennifer Gabrys's *Digital Rubbish* (2011), the sourcing of raw materials and manufacturing of computing hardware is implicated in cultural imaginings of abundance that conceal their wastefulness, afterlives, and precarious conditions of production and disposal, mainly but not exclusively taking place in regions of the Global South. The sites of technological innovation are sometimes also the most toxic. For example, Santa Clara County in California—where Silicon Valley is located—has more Superfund sites⁹ than any other county in the US, 23 in total (Nieves 2018; Schlossberg 2019). Each individual setting along the computing supply chain has its own set of enclosures and exploitation but is also a site of contestations and reclamations of commons—whether it is a drive to unionize the labour force at a tech company or documenting and fighting environmental racism, as exemplified for instance by the work of the [Silicon Valley Toxics Coalition](#).

Recent experimental and artistic works have called into question the current configuration of internet infrastructures in general as well as the way we understand “intelligence” in relation to technology and energy. The [Solar Protocol](#) group is one such instance of denaturalizing the status quo of digital connectivity. Conceived as a network of low-tech servers in different places around the globe, the group raises intriguing questions of what it means to care and maintain shared digital infrastructure (or stewardship, as they call it) and how the protocols that currently govern the internet could be reimaged and rewritten. At the heart of the network configuration is solar energy, harvested via photovoltaic panels at each server site. By routing data traffic according to the logic of the sun—page requests are sent to the server that is enjoying the most sunlight at the time—Solar Protocol taps into what they refer to as a “natural intelligence” as opposed to an “artificial” one. To make sure that most of the computation is really powered by renewable energy, the visualizations for the website happen on the server side instead of in the client's browser.

Thinking about AIs as embodied, material entities as well as energy commons more generally is thus another important emerging theme in both the scholarly literature and citizen-led debates and artistic explorations.

⁹ The term Superfund describes “contaminated sites . . . due to hazardous waste being dumped, left out in the open, or otherwise improperly managed. These sites include manufacturing facilities, processing plants, landfills and mining sites.” (<https://www.epa.gov/superfund/what-superfund>)

At the Milieux Institute, the recently formed [Solar Media Project](#) is engaging with some of these questions and conducting their own experimental and speculative work related to computation and energy.

Reparations and the Black Commons

Inspired among others by civil rights activist Fannie Lou Hamer and her collective Freedom Farms, [urban planner Julian Agyeman](#) and [architect Kofi Boone](#) have laid out ideas for a “Black commons” in the US. These ideas mostly center around redistribution and commoning of land and wealth, but also include practices of crowdsourcing in digital spheres as well as community-run healthcare.

Underlying the recent unrest sweeping U.S. cities over police brutality is a [fundamental inequity](#) in wealth, land and power that has circumscribed black lives since the end of slavery in the U.S. The “[40 acres and a mule](#)” promised to formerly enslaved Africans never came to pass. There was no redistribution of land, no reparations for the wealth extracted from stolen land by stolen labor. . . . An expanded concept of the “black commons”—based on shared economic, cultural and digital resources as well as land—could . . . be a part of undoing the racist legacy of chattel slavery by encouraging economic development and creating communal wealth. (Agyeman and Boone 2020)

Others such as [Susan Witt of the Schumacher Center for a New Economics](#) have put forward proposals for a Black Commons¹⁰ that focus on community land trusts:

Our proposal is to adopt the community land trust structure to serve as a national vehicle to amass purchased and gifted lands in a Black Commons with the specific purpose of facilitating low cost access for Black Americans hitherto without such access. In short creating one piece of a Black Reparations Movement.

The community land trust is a tested and known entity for holding working lands in a commons while at the same time facilitating leaseholders’ ability to build equity in homes and other improvements on the land. Donors would be assured that their one-time donation of land would not again enter the market but would remain a permanent part of a Black Commons. Individual leaseholders could change, and buildings sold, but the land would continue to be held in the nonprofit structure dedicated to serving those disenfranchised by a history of discriminatory practices. (Witt 2018)

¹⁰ For more literature on Black commons in North America/Turtle Island, see also Roane (2018) and Zellars (2020).

5. Conclusion

We cannot offer an answer to the question of what an AI commons could or should be. We instead can acknowledge that commons both form and inform digital assemblages. Commons then are as much a practical set of rules and regulations as a perspective that decentres the privatized and proprietary responses to AI.

We still need to learn how to imagine an AI commons. What we have for now is being attentive to the ways we already exist in common, or what Anna Tsing (2015) calls latent commons. Latent commons provide a space to make (unlikely) alliances and to forge collaboration for a common cause. But unlike understanding commoning as an activity suggests, latent commons cannot simply be turned into blatant commons through deliberate action. They emerge from collaboration where it is unexpected. In a world where institutionalized alienation predominates much of our experience, “fugitive moments” allow for entanglements between humans and nonhumans that are full of unrealized possibilities (Tsing 2015:134, 255). Tsing defines latent commons through a series of negative statements: latent commons are not exclusive to human enclaves, they are not good for everybody, they don’t institutionalize well, and they do not offer redemption.

We like this ambiguity as a final note. We already have an AI commons and yet we do not. We are the commons, and we are not. The AI commons then is always a future—one we hope might seem a little more in play through this event.

Recommended resources for further reading:

A New AI Lexicon by the AI Now Institute

<https://medium.com/a-new-ai-lexicon>

Critical Algorithm Studies: A Reading List, edited and compiled by the Social Media Collective

<https://socialmediacollective.org/reading-lists/critical-algorithm-studies/>

Critical Race & Digital Studies Syllabus, edited and compiled by Lori Kido Lopez and Jackie Land

<https://criticalracedigitalstudies.com/syllabus/>

Pirate Care Syllabus

<https://syllabus.pirate.care/>

Bibliography

- Agyeman, Julian, and Kofi Boone. 2020. "Land Loss Has Plagued Black America since Emancipation—Is It Time to Look Again at 'black Commons' and Collective Ownership?" *The Conversation*, June 18, 2020. <http://theconversation.com/land-loss-has-plagued-black-america-since-emancipation-is-it-time-to-look-again-at-black-commons-and-collective-ownership-140514>.
- Amodei, Dario, and Danny Hernandez. 2018. "AI and Compute," OpenAI Blog, May 16, 2018. <https://openai.com/blog/ai-and-compute/#fn2>.
- Barker, Adam J. 2012. "Already Occupied: Indigenous Peoples, Settler Colonialism and the Occupy Movements in North America." *Social Movement Studies* 11(3–4):327–34. DOI: 10.1080/14742837.2012.708922.
- Belton, Padraig. 2021. "Why Coders Love the AI That Could Put Them out of a Job." *BBC News*, September 6, 2021, sec. Business. <https://www.bbc.com/news/business-57914432>.
- Benkler, Yochai. 2004. "Commons-Based Strategies and the Problems of Patents." *Science* 305(5687):1110–111.
- Coleman, Beth. 2021. "Technology of the Surround." *Catalyst: Feminism, Theory, Technoscience* 7(2):1–21. <https://doi.org/10.28968/cftt.v7i2.35973>.
- Couldry, Nick, and Ulises A. Mejias. 2019. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA: Stanford University Press.
- Crawford, Kate. 2021. *Atlas of AI*. New Haven, CT: Yale University Press.
- Dahlin, Johanna, and Martin Fredriksson. 2017. "Extracting the Commons." *Cultural Studies* 31(2–3):253–76.
- Disco, Nil, and Eda Kranakis. 2013. *Cosmopolitan Commons: Sharing Resources and Risks Across Borders*. Cambridge, MA: MIT Press.
- Dyer-Witheford, Nick. 2002. "E-Capital and the Many-Headed Hydra." Pp. 129–64 in *Critical Perspectives on the Internet*, edited by Greg Elmer. Rowman & Littlefield Publishers.

- Dyer-Witheford, Nick, Atle Mikkola Kjösen, and James Steinhoff. 2019. *Inhuman Power: Artificial Intelligence and the Future of Capitalism*. Pluto Press.
- Eubanks, Virginia. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press.
- Federici, Silvia. N.d. "Feminism and the Politics of the Commons." Accessed May 13, 2021 (<http://wealthofthecommons.org/essay/feminism-and-politics-commons>).
- Fortier, Craig. 2017. *Unsettling the Commons. Social Movements Against, Within, and Beyond Settler Colonialism*. ARP Books.
- Fragnito, Maddalena. 2020. "Commoning Molecules: Decolonising Biological Patents by Gender Hacking Protocols." *Journal of International Women's Studies* 21(7):153–69.
- Freeze, Colin, and Mark MacKinnon. 2018. "Records Reveal AggregateIQ and SCL Group's Plan to Influence Politics in Trinidad and Tobago." *The Globe and Mail*, March 28. <https://www.theglobeandmail.com/canada/article-consulting-firms-in-data-scandal-first-partnered-on-project-in/>.
- Gabrys, Jennifer. 2011. *Digital Rubbish: A Natural History of Electronics*. Ann Arbor: University of Michigan Press.
- Gitelman, Lisa, ed. 2013. *"Raw Data" Is an Oxymoron*. Cambridge, MA: MIT Press.
- Gray, Mary L. 2017. "The Humans Working behind the AI Curtain," *Harvard Business Review*, January 9. <https://hbr.org/2017/01/the-humans-working-behind-the-ai-curtain>.
- Gray, Mary L., and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.
- Henderson, Peter, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. 2020. "Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning," January. <https://arxiv.org/abs/2002.05651v1>.

- Hess, Charlotte, and Elinor Ostrom. 2007. "Introduction: An Overview of the Knowledge Commons." *In Understanding Knowledge as a Commons: From Theory to Practice*, edited by Charlotte Hess and Elinor Ostrom. Cambridge, MA: MIT Press.
- Hunter, Dan. 2003. "Cyberspace as Place and the Tragedy of the Digital Anticommons." *California Law Review* 91(2):439–519. <https://doi.org/10.2307/3481336>.
- Iaconesi, Salvatore, and Oriana Persico. 2021. "When My Child Is AI. Learning and Experiencing through AI Outside the School: The Experiences of a Community AI." *QTimes Journal of Education Technology and Social Studies* XIII(1).
<https://www.qtimes.it/?p=when-my-child-is-ai-learning-and-experiencing-through-ai-outside-the-school-the-experiences-of-a-community-ai>.
- Kelty, Christopher M. 2008. *Two Bits: The Cultural Significance of Free Software*. Durham, NC: Duke University Press.
- Kukutai, Tahu, & Taylor, John, eds. 2016. *Indigenous Data Sovereignty: Toward an Agenda*. Australian National University Press. <http://www.jstor.org/stable/j.ctt1q1crgf>
- Lovink, Geert. 2006. "Introduction to Open Networks Power and Politics of Good Intentions." Pp. 123–24 in *In the Shade of the Commons: Towards a Culture of Open Networks*, edited by Lipika Bansal, Paul Keller, and Geert Lovink. Amsterdam: Waag Society.
- Mansell, Robin. 2012a. *Imagining the Internet: Communication, Innovation, and Governance*. New York: Oxford University Press.
- Mansell, Robin. 2012b. "Promoting Access to Digital Knowledge Resources: Managing in the Commons." *International Journal of the Commons*:1–20.
- McKelvey, Fenwick. 2018. *Internet Daemons: Digital Communications Possessed*. Minneapolis: University of Minnesota Press.
- Milan, Stefania, and Emiliano Treré. 2019. "Big Data from the South(s): Beyond Data Universalism." *Television & New Media* 20(4):319–35. <https://doi.org/10.1177/1527476419837739>.

- Mildenberger, Matto. 2019. "The Tragedy of the Tragedy of the Commons," *Scientific American* Blog Network, April 23.
<https://blogs.scientificamerican.com/voices/the-tragedy-of-the-tragedy-of-the-commons/>.
- Nieves, Evelyn. 2018. "The Superfund Sites of Silicon Valley," *The New York Times*, March 26.
<https://www.nytimes.com/2018/03/26/lens/the-superfund-sites-of-silicon-valley.html>.
- Onuoha, Mimi. 2018. "Notes on Algorithmic Violence," July 2.
<https://github.com/MimiOnuoha/On-Algorithmic-Violence>.
- Ostrom, Elinor. 2015. Governing the Commons. *The Evolution of Institutions for Collective Action*. New York: Cambridge University Press.
- Papadopolous, Dimitris. 2010. "Insurgent Posthumanism." *Ephemera: Theory and Politics in Organisation* 10(2):135–51.
- Pasquinelli, Matteo, and Vladan Joler. 2020. "The Nooscope Manifested: Artificial Intelligence as Instrument of Knowledge Extractivism." KIM research group (Karlsruhe University of Arts and Design) and Share Lab (Novi Sad). <https://nooscope.ai>.
- Pirate Care. N.d. "Commoning Care. Introduction." <https://syllabus.pirate.care/topic/commoningcare/>.
- Puig de La Bellacasa, Maria. 2017. *Matters of Care: Speculative Ethics in More Than Human Worlds*. Minneapolis: University of Minnesota Press.
- Roane, J. T. 2018. "Plotting the Black Commons." *Souls* 20(3):239–66.
<https://doi.org/10.1080/10999949.2018.1532757>.
- Saenko, Kate. 2020. "It Takes a Lot of Energy for Machines to Learn—Here's Why AI Is so Power-Hungry." *The Conversation*.
<http://theconversation.com/it-takes-a-lot-of-energy-for-machines-to-learn-heres-why-ai-is-so-power-hungry-151825>.

- Schlossberg, Tatiana. 2019. "Silicon Valley Is One of the Most Polluted Places in the Country," *The Atlantic*, September 22. <https://www.theatlantic.com/technology/archive/2019/09/silicon-valley-full-superfund-sites/598531/>.
- Schwartz, Roy, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2019. "Green AI." *ArXiv:1907.10597 [Cs, Stat]*. <http://arxiv.org/abs/1907.10597>.
- Southern Poverty Law Center. N.d. "Garrett Hardin." Southern Poverty Law Center. Accessed May 12, 2021 (<https://www.splcenter.org/fighting-hate/extremist-files/individual/garrett-hardin>).
- Strubell, Emma, Ananya Ganesh, and Andrew McCallum. 2019. "Energy and Policy Considerations for Deep Learning in NLP." *ArXiv:1906.02243 [Cs]*. <http://arxiv.org/abs/1906.02243>.
- Thatcher, Jim, David O'Sullivan, and Dillon Mahmoudi. 2016. "Data Colonialism through Accumulation by Dispossession: New Metaphors for Daily Data." *Environment and Planning D: Society and Space*, 34(6):990–1006. <https://doi.org/10.1177/0263775816633195>.
- Tsing, Anna. 2015. *The Mushroom at the End of the World: On the Possibility of Life in Capitalist Ruins*. Princeton, NJ: Princeton University Press.
- Veltmeyer, Henry, and Paul Bowles. 2014. "Extractivist Resistance: The Case of the Enbridge Oil Pipeline Project in Northern British Columbia." *The Extractive Industries and Society* 1(1):59–68. <https://doi.org/10.1016/j.exis.2014.02.002>.
- Verdegem, Pieter. 2022. "Dismantling AI Capitalism: The Commons as an Alternative to the Power Concentration of Big Tech." *AI & Society*. <https://doi.org/10.1007/s00146-022-01437-8>.
- Volont, Louis. 2018. "Who Steals the Goose from off the Common? An Interview with Peter Linebaugh," *Open! Platform for Art, Culture & the Public Domain*, December 22. <https://onlineopen.org/who-steals-the-goose-from-off-the-common>.
- Walljasper, Jay. 2011. "Elinor Ostrom's 8 Principles for Managing a Commons," *On the Commons* (blog), October 2. <https://www.onthecommons.org/magazine/elinor-ostroms-8-principles-managing-commmons>.

Witt, Susan. 2018. "Proposal for a 'Black Commons.'" *Schumacher Center for New Economics* (blog).
<https://centerforneweconomics.org/publications/proposal-for-a-black-commons/>.

Zellars, Rachel. 2020. "Dreams of a Black Commons on Turtle Island." *Studies in Social Justice* 14(2):454–73.