OPEN FORUM

AI ethics as subordinated innovation network

James Steinhoff[1]

Abstract

AI ethics is proposed, by the Big Tech companies which lead AI research and development, as the cure for diverse social problems posed by the commercialization of data-intensive technologies. It aims to reconcile capitalist AI production with ethics. However, AI ethics is itself now the subject of wide criticism; most notably, it is accused of being no more than "ethics washing" a cynical means of dissimulation for Big Tech, while it continues its business operations unchanged. This paper aims to critically assess, and go beyond the ethics washing thesis. I argue that AI ethics is indeed ethics washing, but not only that. It has a more significant economic function for Big Tech. To make this argument I draw on the theory of intellectual monopoly capital. I argue that ethics washing is better understood as a subordinated innovation network: a dispersed network of contributors beyond Big Tech's formal employment whose research is indirectly planned by Big Tech, which also appropriates its results. These results are not intended to render AI more ethical, but rather to advance the business processes of data-intensive capital. Because the parameters of AI ethics are indirectly set in advance by Big tech, the ostensible goal that AI ethics sets for itself—to resolve the contradiction between business and ethics—is in fact insoluble. I demonstrate this via an analysis of the latest trend in AI ethics: the operationalization of ethical principles.

Keywords AI ethics · Innovation · Machine learning · Industry · Political economy · Operationalization

1 Introduction

In the wake of mounting social criticism following several scandals from 2016 on, the Big Tech firms which lead artificial intelligence (AI) research and production have developed an apparent interest in AI ethics, referred to variously as responsible AI, trustworthy AI, socially-beneficial AI, democratic AI, and human-centered AI, among other terms. Regardless of how it is termed, the idea is that the scandals with which data-intensive capitalism is rife derive from an ethical deficit in AI research, production and deployment, which can be remedied by an increased focus on developing the morals and ethical behaviour of computer scientists and engineers (Green Hoffman and Stark 2019). Interest in AI ethics is now shared by governments, international organizations, NGOs and academic researchers. Yet, as it has proliferated, AI ethics has itself become the subject of criticism. Most prevalent is the claim that AI ethics is merely "ethics washing" (Metzinger 2019). The notion is that since the research and production of AI is led by profit-seeking companies, discussion of ethical matters is an act of cynical dissimulation serving, at best, a public relations function for those companies.

However, the ethics washing claim is complicated by the fact that AI ethics is not just something done by Big Tech. There are a wide variety of actors involved in AI ethics, including academics, non-profits, grassroots organizations and small companies. Not all of these actors have priorities directly aligned with big tech; indeed some are in conflict with it. A second complication comes from the fact that research on environmental "green washing" (from which ethics washing derives its name) suggests that such efforts have little efficacy, as the audiences to whom they are directed are not so easily duped (Rahman et al. 2015; de Jong et al. 2020). Taken together, these two complications raise the question of whether AI ethics is indeed merely cynical dissimulation.

This paper contends that while AI ethics is indeed ethics washing, it also serves the economic exigencies of data-intensive capital more directly. It does not do so by advancing AI ethics' proclaimed goal of rendering AI more ethical.

✉ James Steinhoff
  james.steinhoff@ucd.ie

1 School of Information and Communication Studies,
  University College Dublin, Dublin, Ireland

Instead, AI ethics functions as a subordinated innovation network (Rikap 2021). A subordinated innovation network is a dispersed social relation through which Big Tech wields indirect control over research conducted outside of its legally owned resources, directing it towards ends which will advance its commodity production, circulation and other business processes. This is the primary mechanism of intellectual monopoly capitalism (Rikap 2021), which depends on the appropriation of knowledge produced by individuals and organizations outside of a monopoly capital and its conversion into commodities. AI ethics, I contend, is truly about neither AI nor ethics, but rather the accumulation of capital. The economic function of AI ethics as a subordinated innovation network is thus at odds with AI ethics' proclaimed goal of rendering AI more ethical. Indeed, AI ethics is wracked by an internal contradiction between capital and ethics. This is a contradiction which cannot be resolved, except by evacuating the notion of ethics of any content and letting it be defined by capital. This explains why AI ethics exhibits vacuity; such that one practitioner describes the field as having undergone a "moral collapse" (Abdurahman 2020).

The paper proceeds as follows. First, I briefly contextualize the appearance of AI ethics. Next, I review the critical literature on AI ethics and consider the ubiquitous accusation of ethics washing. Then I show how ethics washing criticisms point to a contradiction immanent to AI ethics. I suggest that to fully understand this contradiction and its significance for AI ethics, we need to switch from analysing the discourse of AI ethics and adopt a political economy perspective. I review several attempts which have been made to make this switch, but find them inadequate. Then I introduce the theory of intellectual monopoly capitalism. Drawing on this theory, I argue that AI ethics is an example of a subordinated innovation network—a dispersed network of labourers whose research is indirectly planned by Big Tech, and whose outputs are appropriated by Big Tech. I substantiate my contention with a case study of the latest trend in AI ethics: the "operationalization" of ethical principles. I argue that existing attempts at operationalization provide evidence of Big Tech's subordination of AI ethics research. This is a theoretical paper and the argument advanced here will benefit from the subsequent empirical investigation. However, it provides a basis for asserting that there can be no interesting future for AI ethics unless it begins from a stance of intentional incompatibility with the capitalist production of AI.

2 Context

Research on the ethical dimensions of AI predates the current technological milieu centered on machine learning (Wallach and Allen 2008), but the contemporary phenomenon of AI ethics has its beginning in the mid-2010s when machine learning emerged as a viable technique in many application domains. In the wake of manifold scandals including that involving Cambridge Analytica, Facebook and the 2016 US presidential elections, as well as the exposure of Google's secret plans to produce military drone vision systems, the Big Tech firms at the head of the AI industry grappled with a rising social backlash from diverse sectors of society: the so-called "techlash" (Rosenberg et al. 2018; Foroohar 2018; Green 2021). From 2016 on, AI companies began issuing statements proclaiming their ethical AI principles. By 2019, nearly all the US Big Tech companies including Microsoft and Google, some Chinese Big Tech companies and organizations such as Baidu and the Artificial Intelligence Industry Alliance, several smaller but influential AI firms such as DeepMind, as well as several think tanks and industry-adjacent organizations like The Partnership on AI had some form of ethical AI principles on display (Green 2021; Arcesati 2021). As of early 2020, there were 167 AI ethics guidelines documents around the world (AlgorithmWatch 2020). The number has no doubt increased since.

3 Content and Critique

According to Big Tech, the production of AI commodities is a profoundly ethical endeavour. IBM tells us AI ethics is "a framework that guides data scientists and researchers to build AI systems in an ethical manner to benefit society as a whole" (IBM Cloud Education 2022). Tencent argues that "just as Noah's Ark preserved the fire of human civilization, the healthy development of AI needs to be guaranteed by the 'ethical ark'" (Cao 2020). Microsoft (nd.a) declares that it is "committed to the advancement of AI driven by ethical principles that put people first". Google (nd.a) says that the "vast opportunity" presented by AI "carries with it a deep responsibility to build AI that works for everyone". Sometimes Big Tech makes an additional claim, asserting that AI ethics are identical to good business sense. Eric Horvitz, Microsoft's Chief Scientific Officer, describes responsible AI as "a critical part of innovation across organizations" (Microsoft nd.b) while Google (nd.a) states plainly that "values-based AI is good for your business".

Such formulations raise immediate questions since ethics refers to a vast field with an ancient, global history, full of myriad possible positions. One might ask: which values and ethical systems are good for the AI industry? An answer to this question is not provided in the ethical AI discourse of Big Tech. Perhaps the only thing that can be said about ethics in general is that it is not universal or timeless; rather, a given ethical theory or system necessarily arises within particular social relations (Noonan 2003; Robles Carrillo 2020). Clearly, AI ethics arises in the context of the AI

industry. Capitalist industry, of course, has specific needs, and is indisputably not compatible with every possible system of ethics, as I will discuss later. But first, let us consider existing critiques of the content of AI ethics.

According to Hagendorff's (2020) analysis, AI ethics understands ethics primarily in terms of principles of accountability, privacy and fairness. This is striking because these aspects "are those for which technical fixes can be or have already been developed" and "those which may be "most easily operationalized mathematically and … implemented in terms of technical solutions" (Hagendorff 2020, p. 103). The analysis of Jobin et al. (2019) recognizes similar terms as most prevalent: transparency, justice and fairness, non-maleficence, responsibility and privacy. Again, these are possible technical fixes, with the exception of non-maleficence, which is so vague as to be meaningless in an industry context. One might wonder how many companies there are that produce openly maleficent commodities.

Ethical aspects which are less amenable to technical fixes receive little mention. Hagendorff (2020) points out that "almost no guideline talks about AI in contexts of care, nurture, help, welfare, social responsibility or ecological networks" (p. 103) while Jobin et al. (2019) note a "thematic underrepresentation of sustainability and solidarity". In sum, as Sloane (2019) puts it: AI ethics is conceived of in a way which does not require examination of "historic, systematic and complex inequalities".

No doubt taking note of such critiques, some AI producers have since updated their principles, albeit slightly. While, as of November 2022, Microsoft has eschewed any substantial changes, Google's (nd.b) AI principles now open with "Be socially beneficial" which is explained as: "we will take into account a broad range of social and economic factors, and will proceed where we believe that the overall likely benefits substantially exceed the foreseeable risks and downsides" while "we will continue to respect cultural, social, and legal norms in the countries where we operate". Beyond this, the principles remain close to those noted by Jobin et al. and Hagendorff, including: bias, safety, accountability, privacy, scientific excellence and availability.

Such minor revisions do little to address the chorus of criticism that argues that AI ethics is merely "ethics washing", a facade or a cynical gesture (Metzinger 2019). In an early and influential article, Ochigame (2019) holds that ethical AI is "aligned strategically with a Silicon Valley effort seeking to avoid legally enforceable restrictions of controversial technologies". Others deride it as a "smokescreen" (Sloane 2019, p.3), a "marketing strategy" (Hagendorff 2020, p.113) and "yet another proxy for advancing various types of interests—be they financial in the case of private companies, or political in the case of states" (Vică et al. 2021, p.91). Such critics agree that AI ethics is "largely deployed to gain competitive advantage (between firms,

industries, nations) rather than initiating a genuine push towards social justice" (Sloane 2019.

4 Contradiction

As the critiques cited above indicate, ethics washing critiques are motivated by the context of the capitalist industry in which most AI research and development occurs. This industrial context generates a contradiction within AI ethics, between capital and ethics. Some critics have addressed this contradiction directly. Green (2021) holds that "[w]hen ethical ideals are at odds with a company's bottom line, they are met with resistance" (p.214). Phan et al. (2021) argue that "attempts to reconcile a contradiction between ethics and commercial profit usually results in ethical products being shaped to consumer demand or the business needs of 'end users'" (p.11). Ebell et al. (2021) argue that AI ethics has a "fundamental conflict of interest" (p.133). Chen et al. (2022) agree, holding that "Business needs are often in conflict with ethics and transparency" and that "ultimately industrial and practical applications will be the determining factor in ethical behavior of AI" (p.4).

The contradiction between ethics and capital manifests in several ways. One report suggests that this contradiction manifests in a cognitive dissonance in people working in ethical AI who are torsioned "between external pressures to respond to ethical crises at the same time that they must be responsive to the internal logics of their companies and the industry" (Moss and Metcalf 2019). This torsion is reportedly why people working in AI ethics have severe "burnout" problems exceeding the already high norm in the tech industry at large (Heikkilä 2022). The contradictory nature of AI ethics also suggests why it is that a study which surveyed 211 software companies found that "AI ethics guidelines have not had a notable impact on practice" (Vakkuri et al. 2020, p.3).

In the course of my research into AI ethics, I conducted interviews with people working in the field in a variety of contexts, including academia, Big Tech, startup companies, grassroots organizations and the policy space. While this paper draws on that interview data primarily as background information, it is worth including excerpts from two interviewees who spoke precisely to the contradiction highlighted in the literature. One researcher who has held both academic and industry positions told me: "there's a split [within AI ethics] There's the people who work for the big tech companies. And then there are those who don't". This researcher went on to supply the following vivid imagery:

> it's a little bit like tectonic plates that … for a moment between 2010 and 2020, let's say … the interests of at least some big tech companies or some people in big

tech companies and the interests of academic researchers who were critical of this space, they were aligned and … they were kind of moving in parallel. And now, the plate has gone under. And that's causing earthquakes of all sorts.

An industry data scientist painted AI ethics in a similar light, describing an antagonistic divide between two factions: "I wouldn't even say there's an uneasy alliance ... there's this wildly unequal amount of distribution of resources … big tech has so much fucking money". According to these accounts, the contradiction between capital and ethics is manifest in the formation of opposed factions within AI ethics, in addition to its manifestation in the anodyne content of AI ethics, cognitive dissonance of workers, and ineffectiveness in the application of ethical principles. Let us examine this contradiction in greater detail.

We have noted already that ethics is a vast and varied field about which few generalizations can be made. To see how such a broad field can come into contradiction with capital, we need to understand what exactly capital is. Capital is a quantity of value invested in the production of commodities (including the purchase of labour power and materials) with the intent of selling those commodities to generate more value than was initially invested. Labour power is importantly purchased for less than the value it generates in the course of production, producing "surplus-value" which accrues to capital (Marx 1990, 326). Karl Marx schematized this process as the circuit M-C-M', or: money-commodity-more money (Marx 1990, 251). Capital is thus defined as the increase of value via commodity production which relies on the appropriation of value from labour. Capitalism is the mode of production based on this particular social relation. Unlike ethics, which is a broad field, capital and capitalism have very narrow meanings. Ethics does not factor into capital at the definitional level. Indeed, Marx held that "the immanent law" of capital: "to produce as much surplus-value as possible" was simultaneously its only "moral imperative" (Marx 1990, p.1051). In other words, capital cannot have a moral imperative.

Capitalist firms must seek to increase their value by completing the circuit of capital again and again. They must do so ruthlessly, because they must compete against rival firms on the market, and if they fail to take a sufficient share of value from those competitors, they will eventually cease to exist. This is true whether a firm sells ballistic missiles, soap or AI, and regardless of the set of values held by the particular capitalist at the helm of that firm. The structure of capitalist production thus manifests, via competition, as a suite of "coercive laws" which limit the possible actions of capitalist firms (Marx 1990, 433). Any diversion of resources to ends other than the increase of value detracts from a firm's ability to achieve its necessary goal.

It does not take a radical critic of capitalism to come to such conclusions. In fact, some of the most ardent supporters of an unfettered capitalism agree that a contradiction exists between capital and ethics. The economist Milton Friedman, one of neoliberalism's greatest champions, argued the following:

What does it mean to say that the corporate executive has a 'social responsibility' in his capacity as businessman? If this statement is not pure rhetoric, it must mean that he is to act in some way that is not in the interest of his employers (Friedman 2007 [1970], p.174)

According to this view, business has a narrow definition which does not include the diverse range of interests that might fall under the category of social responsibility—another way of saying ethics. The only possible "social responsibility" for business, he holds, is "to increase its profits" (Friedman 2007 [1970], p.173). This view that capitalism has nothing to do with ethics was a convenient theory for capitalists and helped justify the neoliberal dismantling of Keynesian economic policies which involved forms of social responsibility such as welfare. While Keynes saw capitalism as a productive yet dangerous system, the unethical excesses of which had to be restrained by government interventions (Freeman et al. 2007), for neoliberalism, ethics is replaced by the market: "the operation of a market … is seen as an ethic in itself, capable of acting as a guide for all human action, and substituting for all previously existing ethical beliefs" (Treanor 2005).

These two analyses coming from very different perspectives both maintain that the contradiction between ethics and capital is insoluble, even if ethics is a vaguely defined term. They correctly point out that capital's very specific needs mean that, by definition, ethics must be incompatible with it. However, despite the rather clear and simple nature of this proposition, it is not one endorsed by most proponents of capitalism today.

Unsurprisingly, firms do not represent their operations as incompatible with ethics. The World Economic Forum even aims to surmount the contradiction between ethics and capital via their programme of so-called "stakeholder capitalism" in which the non-ethical Friedmanian logic of shareholder capitalism will be supplanted by a capitalism in which firms "seek long-term value creation by taking into account the needs of all their stakeholders, and society at large" (Schwab and Vanham 2021). The stakeholders they refer to are no less than "all human individuals" and "the natural environment we all share" (Schwab and Vanham 2021). In other words, social responsibility or some form of ethics should and can, be incorporated into capitalist production. This view that the contradiction between capital

and ethics can be overcome has precedents. And these precedents are also connected to AI ethics.

In an excellent historical and theoretical study of business ethics, Gabriel Abend describes its central premise as the notion that:

> business and morality can be reconciled, everyone will win, capitalism is not morally bad even if there will always be a few bad apples among business people (as among any other group), and the ethics of business can and should be improved through education, incentives, organizational design, or legislation (Abend 2014, p.145).

This is the very premise behind stakeholder capitalism. Despite a century of effort, business ethics has yet to provide a convincing argument for its central premise. As Abend (2014) demonstrates, the history of business ethics is marked by "little novelty and originality … Normative prescriptions, codes of ethics, business ethics classes, speeches in the legislature, newspaper editorials, and outraged reactions to scandals repeat themselves over and over again a constant déjà vu (p.651). Indeed, drawing on Abend's analysis, Greene, Hoffman and Stark (2019) argue that AI ethics is best understood as yet another instance of that central premise of business ethics, with little to add to its repetitive history other than a new technology of interest.

To gather up the various threads discussed thus far, ethical AI is, according to its self-presentation, trying to resolve a contradiction. However, like business ethics before it, ethical AI faces an insoluble contradiction because a solution would require AI-producing capital to function sub-optimally as capital. This conclusion can be reached by following the logic of either Marx or Friedman. If the contradiction is insoluble, then AI ethics cannot serve to make the AI industry more ethical. It is therefore reasonable to level the charge of ethics washing at AI ethics. But I argue that it is more than a simple attempt to dupe consumers with an ethical facade. I contend that AI ethics also serves an economic function for Big Tech. The insolubility of the contradiction within AI ethics, is, I think, key to understanding what the economic function of AI ethics is. To grasp this function, we need to shift levels of analysis, from the discursive to the political economic, and to situate AI ethics amid broader changes in the capitalist mode of production.

5 An ethical economy?

One political economic interpretation of AI ethics comes from before the widespread commercialization of machine learning. Arvidsson (2010) holds that the rise of ethics in the tech industry is "more than just a cynical move" or a "matter of benevolence" (p.637). The rise of ethics is actually a manifestation of a new "ethical economy" characterized by "the growth of a number of strategically central, productive practices: all working according to a logic where value is related to the quality of social relations, and not to the quantity of productive time" (Arvidsson 2010, p.637). Arvidsson goes on to suggest that the "ethical economy is likely to be central to the emerging economic ecology of the information society" and that it might "even become hegemonic" such that ethics could even replace labour as the source of value in capitalism (Arvidsson 2010, p.637–8). Arvidsson is drawing on the notion, developed by *post-operaismo* thinkers such as Hardt and Negri (2001), that the proliferation of information technology will reconfigure the capitalist mode of production, away from a centralized industrial model, towards a decentralized mode of networked production which capital cannot directly command, but only parasitically appropriate the output of. The inadequacies of the *post-operaismo* approach have been demonstrated both theoretically and empirically (Pitts 2017) and particularly in the context of the AI industry (Steinhoff 2021), so there is little reason to pursue this interpretation further. While so-called "ethical consumerism" may be informing business strategies (or the ethics washing of them), its applications are limited by the necessary constraints of capitalist production (Newholm 2017) and it is safe to say that capitalist industry retains its historical mechanism of the appropriation by the capital of value produced by labour.

Less radically than Arvidsson, Phan et al. (2021) argue that AI ethics should be understood as "a marketplace of ethical skills, signals and knowledge" which they call an "economy of virtue" (p.1). In this economy, "virtue and ethics are the primary objects that are produced and circulated by groups inside Big Tech through the establishment of, for example, ethics boards and working groups and also outside, from Universities, research institutes, consultancies, and other allied industries" (Phan et al. 2021, p.1). Here the idea is not that ethics supplants labour as the source of value, but that ethics becomes an increasingly important mode of representation, by which actors signal their participation in processes of social betterment. Hu (2021) makes a similar point, arguing that "just as Big Tech needs 'ethics' on its side to maintain public goodwill, "ethics" ventures need Big Tech for their own legitimacy" (p.240–241). Hu elaborates: "ethical tech institutions are in fact parasitic on the continual moral failures and disappointments of a hegemonic tech industry. These groups and efforts survive only because Big Tech has chosen to engage the ethics discourse while it has blocked most other political movement-building" (Hu 2021, p.241).

There is much of value in the analyses of Phan et al. and Hu. Both rightly point out how AI ethics involves people and institutions outside of Big Tech, and they demonstrate how such external actors—whatever their goals—can be

brought into alignment with those of Big Tech. However, their analyses still do not go very far beyond the standard ethics washing thesis. Saying that AI ethics serves as a form of virtue signaling means that it amounts to little more than an attempted trick. It provides the benefit of an appearance of participation in ethical activity. I contend that AI ethics serves another function of more substantial benefit to Big Tech. Hu's analysis points the way to such a line of thought by contending that AI ethics is an avenue chosen by Big Tech while it blocks others. To grasp the economic function we need to gain a different perspective on AI ethics by considering it from the other side of the class divide. Rather than as solely a technique of capital, we need to approach AI ethics as also a kind of labour.

6 Intellectual monopoly capital

Whatever else AI ethics might be, it is for most people involved with it, a kind of labour, meaning they engage in it in the course of working for a living. As we have seen, AI ethics has a theoretical precedent in business ethics. However, from a labour perspective there is a longer history and wider perspective that AI ethics can be situated in as well, which pertains to the digital networking of the capitalist mode of production.

While capitalism has relied on global networks of trade since its earliest days, since the 1970s and advances in communications and transportation technologies (Martin 2016), the production of commodities has been radically fragmented across the world into "global value chains" (Johnson 2018), with each moment of production located wherever the requisite commodities, including labour-power, are cheapest (Gereffi, Korzeniewicz and Korzeniewicz 1994). For example, the production of an Apple iPhone utilizes inputs from 43 countries across six continents (Petrova 2018) and all computing firms, including the rest of Big Tech, rely on similarly dispersed processes. The World Bank (2020) estimates that "almost half of all trade" moves through global value chains.

While global value chains are spatially distributed, they are not characterized by a homogenous distribution of wealth and power. Rather, they are characterized by a funneling of resources in one direction and the imposition of command in the other. While the bulk of labour is performed in poorer regions, the largest share of value accrues to the richest regions (Suwandi et al. 2019). The operations of global value chains are directed by the centralized powers of large corporations which are capable of coordinating—and disciplining—the many participating firms along the chain (Tsing 2009). The rise of global value chains is thus far from an international democratization of production. It is better understood as the evolution of capitalist planning.

The combination of words "capitalist planning" may sound strange to some ears. Planning—the "direct allocation" of resources (Mandel 1986, p.7)—is usually associated with socialist economies, while capitalist economies are said to rely on markets to allocate resources without explicit planning. However, at least since the end of the Second World War, capitalist economies have also engaged in planning "to deal with the economic, as much as political, consequences of high employment policies" among other factors (Warren 1972). As corporations grew to unprecedented sizes in the latter half of the century, they could not depend on existing markets to absorb their burgeoning outputs and thus began developing means to manipulate the circulation of commodities to their advantage (Baran and Sweezy 1966). Capitalist production then called for "an immense amount of social coordination that was not previously required" (Braverman 1998, p.186). In the past two decades, corporations such as Walmart and Amazon have pushed capitalist planning to new heights, accelerating global value chains and subjecting markets to ever more sophisticated manipulation via data surveillance, targeted advertising and recommendation systems. As Phillips and Rozworski (2019) show, today's market economy is "rife with planning" (p.50).

The apex of capitalist planning so far, according to Cecilia Rikap, occurs with the development of intellectual monopoly capitalism (IMC). This refers to a particular form of monopoly capitalism developed by large corporations that produce intangible commodities, including pharmaceuticals and, most relevant to the purposes of this paper, software. Indeed, the Big Tech companies which launched the ethical AI phenomenon are exemplars of IMC for Rikap.

While conventional monopolies wield power within a given market, destroying rivals or making it impossible for new competitors to enter, intellectual monopolies have "power [which] extends beyond the market and takes the form of capitalist planning of production and innovation" (Rikap 2021, p.11). IMCs rely not only on the planning of production and circulation via the construction of global value chains. They also rely on planning in the phase of innovation which precedes production (Rikap and Lundvall 2021, p.46). The notion of innovation is often treated as sacred in acritical industry discourse and business scholarship; as a magical property of capitalism (Florida and Kenny 1993). Theorization of innovation is usually traced back to the neoclassical economist Joseph Schumpeter; however, Schumpter himself noted that Marx had discussed the topic long before him (Schumpter 1943, p. 82, cited in Walsh 2021). While a detailed theoretical analysis of innovation is beyond the scope of this paper, we can note that it, in fact, has a very simple definition within a capitalist mode of production, since, as discussed above, capital has a narrow definition. As Walsh (2021) succinctly puts it, within capitalism, "innovation is firstly a vehicle for the accumulation

of capital; any other concerns come second" (p.7). With this in mind, it is easy to understand how IMCs plan innovation.

While global value chains rely on the outsourcing of labour, IMCs outsource innovation via the creation of "innovation networks" consisting of companies, research organizations and universities which work on research and development in ostensible partnership with a monopoly (Rikap 2021, p.175). While organizations within an innovation network are ostensibly regarded as partners, they lack the power to influence the agenda of the IMC with which they partner, just as contributors to Apple's global value chain are unable to direct the development of the iPhone. Such organizations thus exist in a relation of "subordination" since, while they contribute to innovation processes, the outputs of these "are mostly transformed into intangible assets by the intellectual monopoly" (Rikap 2021, p.175). In other words, while subordinated organizations contribute to the production of new knowledge and thereby commodities, they tend not to retain ownership over this knowledge or its commodified forms.[1] Subordinated organizations accept this relationship "because this is their best survival strategy, but this does not mean that the relationship is equally beneficial" to them and the IMC (Rikap and Lundvall 2021, p.47).

The ability to utilize innovation networks is akin to the ability to construct global value chains in that both are possible only for large firms with sufficient resources. The creation and control of subordinated innovation networks is thus an advanced form of planning which depends on the "capacity of certain firms to organize long-term capital accumulation beyond their legally owned capital" (Rikap 2021, p.11). IMCs "plan the production and innovation processes of subordinated firms and other organizations" by "controlling management's critical parameters. They also define R&D agendas, clauses of exclusivity, commercial credit conditions, quality standards and other regulatory matters" (Rikap 2021, p.11). In other words, IMCs set the directions and priorities of the research their partner organizations engage in, such that it benefits their particular goals, and is amenable to incorporation into their existing business processes.

The creation of innovation networks is ubiquitous in Big Tech. Rikap and Lundvall (2021) demonstrate this via a comparison between the high frequency with which intellectual monopolies co-publish research with other organizations to the low frequency with which they co-patent related research. While Google authored 6447 publications up to 2019, with 3397 co-authoring organizations, only 65 (0.3%) of its 25,538 applied and granted patents are co-owned with

another organization (Rikap and Lundvall 2021, p.50). Since intellectual monopolies do not share ownership of the vast majority of the patents relating to their co-published research, they are evidently appropriating knowledge from subordinated organizations.

The subordination of an innovation network may also occur in a more diffused manner, and a form which is less easy to document, if a network consists not of subordinated firms, but of various individuals within and without a variety of organizations—some of which may not be formally employed. Here, an example is open-source software development. While such software is freely available to anyone, it accrues particular benefits to IMCs which are able to incorporate it into their commodities at scale (Rikap 2020; Rikap and Lundvall 2021) and who have the resources and technology ecosystems to reap various other benefits, from on-ramping skilled labour to locking future applications into their infrastructure (Dyer-Witheford, Kjøsen and Steinhoff 2019, p.54–56; Birkinbine 2020). Innovation networks of this diffuse sort reach out for input beyond the labour market into the commons.

A commons is usually taken to refer to a resource which is accessible to all members of society; in other words, resources not governed by the now-pervasive strictures of private property and capital. However, Bollier (2014) argues that a commons is better described as "a resource + a community + a set of social protocols" which are used to manage that resource (p.15). In other words, a resource along with a set of social relations and actors in which it is embedded. A commons may be enclosed when its resource is wrested from its existing social relations and transferred into relations congruent with commodity exchange. Allen and Potts (2016) argue that innovation truly begins, not with the valiant entrepreneur, but within "innovation commons" in which shared knowledge accrues around particular technologies and applications. They suggest that "defence against enclosure" is a vital component of encouraging innovation even if, for them, the ultimate goal of innovation is commodity production (p.1047). As an advanced form of capitalist planning, the IMC model does not attempt to enclose innovation commons in a conventional sense, instead it grants them ostensible autonomy while directing its operations and appropriating its outputs—what Rikap calls subordination. This, I suggest, is how the AI ethics phenomenon should be understood.

7 AI ethics as subordinated innovation network

My contention is that AI ethics is a subordinated innovation network. Like the innovation network constituted by open source software research, it is highly diffuse, composed of some individuals which are paid by Big Tech for their

---

[1] Also pivotal to the IMC model, though beyond the scope of this paper, is the development of intellectual property rights, which facilitate the capture and concentration of intangible commodities by IMCs (Rikap and Lundvall 2021, p.26).

work, but also of individuals who work in startup companies, universities, research labs and NGOs. The AI ethics phenomenon as a whole performs an innovation function for the Big Tech companies which are capable of productively appropriating the output of this network. The following sections sketch the rudiments of this theory.

7.1 Planning

The first aspect of this theory is that AI ethics is an instance of capitalist planning. The key is that the problem set by AI ethics, the resolution of the contradiction between capital and ethics, is insoluble. The only possible way out of the contradiction is for ethics to be made isomorphic to capital. Capital cannot cease being capital; so all possibilities for a compromise must skew towards the benefit of capital and the attenuation of ethics. Thus, "the plan" for AI ethics is not to resolve the contradiction between capital and ethics, but to maintain it, as this interminable conflict is potentially productive of new ideas which are predisposed not to conflict fundamentally with capitalist AI production, including the making of new commodities or otherwise enhancing business processes.

As Hu (2021) and Phan et al. (2021) recognize, AI ethics is for many contributors a form of work, and those workers are thus drawn into a relationship of dependence on the AI firms which they critique. Such workers will seek to keep generating research, and thus sustain the contradiction, because like all of us, they need to work to survive. However, since the contradiction is insoluble, their work—whatever its particular conclusions—must go down one of two paths. The first path is to acknowledge the insolubility of the contradiction between capital and ethics, and thus to render fruitless one's own AI ethics research (and undermine one's chances of funding from industry sources). The second path is to attenuate the ethical component and accept capital's framing of AI ethics as something compatible with the accumulation imperative. In this way, the dispersed labourers of the AI ethics innovation network are subordinated to the IMCs which lead AI research and production. While AI ethics practitioners may work outside the legal boundaries of Big Tech, the parameters of AI ethics research are already set such that we can regard this as an instance of capitalist planning.

7.2 Output

The second aspect of this theory regards the nature of AI ethics outputs, or the kinds of things that are produced by AI ethics research. I contend that AI ethics generates innovations useful to Big Tech, but these are not innovations which render Big Tech's AI operations more ethical.

Innovation takes on the significance of mythological proportions in industry and economics discourses, but as we have seen, it has a necessarily narrow meaning within a capitalist economy: the opening of new avenues for accumulation. To understand the contribution of ethical AI to Big Tech's capital accumulation we need to consider the types of commodities produced by Big Tech, and the inputs on which they rely. Outputs include: targeted advertisements, software, including AI, and software-related services, all of which rely on the collection of large quantities of data for training machine learning models. Many IMCs are thus data-driven IMCs (Rikap 2022). Since most of the data of interest to IMCs derives from the surveillance of users of applications and platforms, these firms may accurately be called surveillance capitals. Firms employing a surveillance capitalist model tend to follow a cyclical business process which Shoshana Zuboff (2019) calls the dispossession cycle.[2]

The dispossession cycle begins with *incursion* in which a data collection/surveillance function is deployed in a new context, which may not be legal or ethically or socially palatable to most people. This is followed by *habituation* in which acceptance of the new incursion is inculcated such that it becomes the "new normal". Third, comes *adaptation*, which refers to how after introducing an invasive business practice and receiving backlash, surveillance capitals deploy "superficial but tactically effective adaptations that satisfy the immediate demands of government authorities, court rulings, and public opinion" (Zuboff 2019, p.170). This leads to the final stage of *redirection* in which projects are reconfigured to operate in ways not apparently subject to the criticisms already offered, while continuing their underlying operations unabated. This cycle drives the expansion of the "perpetual-motion machine" of data collection to ever new sectors (Zuboff 2019, p.170). Apologies will abound after a new incursion, but no substantial changes occur, indeed they cannot, insofar as the harvest of data must continue or these companies cannot continue to be profitable. As Zuboff emphasizes, the surveillance practices to which people object cannot be reformed—they can only be hidden behind a facade of concern.

My contention is that AI ethics outputs contribute primarily to the adaptation and redirection phases of the dispossession cycle, as they generate means of tweaking commodities and business processes in ways which may address immediate concerns, without inhibiting essential mechanisms. This is not mere ethics washing as a facade, as it collects actual

---

[2] Zuboff's analysis of surveillance capitalism should be a component of any analysis of data-intensive capitalism; however, it suffers from an unjustified insistence that surveillance capitalism is somehow a pathological mutation of capitalism which differs in essence from its traditional, benevolent form. Morozov (2019) provides an effective in-depth critique of this aspect of her work.

innovations which are applicable to the advancement of data-intensive capital valorization. To test the theory of AI ethics as a subordinated innovation network we can consider the latest development in ethical AI: operationalization.

8 Case Studies: The Operationalization of AI Ethics

In an influential critique, Mittelstadt (2019) argues that "[p]rinciples alone cannot guarantee ethical AI" and suggests the "real ethical challenges" will come in figuring out how to "translate and implement" principles. Several other critics have made the same essential point (Dignum 2019; Rességuier and Rodrigues 2020; Hagendorff 2020). The idea is that ethical AI principles need to be transformed into concrete methods which can be put to use during the design and deployment of AI systems—in other words, operationalized.

Discussion of operationalizing ethical principles now appears in the ethical AI discourse of Big Tech. IBM has published a report on "AI ethics in action" (IBM Corporation 2022) and Microsoft has established no less than three internal groups tasked with AI ethics operationalization: the AETHER Committee, the Office of Responsible AI (ORA), Responsible AI Strategy in Engineering (RAISE) as well as other groups devoted to ethics more broadly, including the Ethics and Society team. Business-oriented publications such as *Harvard Business Review* and *Forbes* now publish articles such as "A Practical Guide to Building Ethical AI" (Blackman 2020) and "Operationalizing AI Ethics, No Longer An Option But An Imperative" (Dhinakaran 2021). Let us consider some examples.

First, two examples from Big Tech. Here I draw on a white paper published by the UC Berkeley Center for Long-Term Cybersecurity. In this document, Cussins Newman (2020) analyses three attempts at operationalizing ethical AI principles, two of which are relevant to this paper. The first concerns OpenAI. Cussins Newman (2020) praises OpenAI for operationalizing AI ethics in the staggered release of their large language model GPT-2. In the wake of early concerns expressed about the malicious uses of large language models, OpenAI decided to release GPT-2 incrementally. Only certain functionalities would be available to begin with as part of a risk-mitigation strategy. Cussins Newman (2020) holds that this approach effectively operationalized AI ethics, bucked industry and computer science trends and allowed OpenAI to conscientiously "monitor uses, engage with partner organizations on particular research questions, and promote awareness of impacts" (29).

The second example concerns Microsoft's aforementioned AETHER Committee. AETHER was established with the expressed intention of facilitating internal deliberation about AI ethics by establishing seven working groups devoted to topics like Bias and Fairness. Through these working groups, workers can research and bring topics of concern to management via a supposedly transparent process. Cussins Newman (2020) glowingly assesses AETHER because:

> it provides a clear signal to employees, users, clients, and partners that Microsoft intends to hold its technology to a higher standard. AETHER shows one pathway by which companies can empower employees to voice concerns and work toward new company practices and policies supporting the responsible development and use of AI (20).

Both the establishment of AETHER and the staggered release of GPT-2 are taken as evidence of successful operationalization because they "represent shifts in practices and policies that were made across entire companies and organizations, with evidence of spillover effects to other parts of the AI ecosystem already present" (Cussins Newman 2020, 12). The broad idea seems to be that these gestures were not guided purely by ruthless capitalist calculation, but by a genuine desire to render AI more ethical. However, at the time of writing three years later, we are in a better position to assess how effective these operationalizations have been. In both cases, spillover effects pertaining to ethical AI are not obvious. Indeed, without exaggeration, the AI industry has moved in the opposite direction.

The staggered release of GPT-2 has not become an industry standard. On the contrary, spurred by the release of OpenAI's shockingly capable ChatGPT in 2022, competitors in the AI industry are currently scrambling to release their own large language models, even as employees publicly complain that these models are not ready for launch (Elias 2023). Google's Bard model gave an incorrect answer during its first public use, causing a 7% ($100 billion) drop in Alphabet shares (Sherman 2023) while Baidu's Ernie model underwhelmed viewers on its public unveiling. Ernie was demonstrated only in prerecorded video, presumably due to performance anxieties, and its presentation was followed by a 10% drop in Baidu shares (Liao 2023). A race to the market, rather than an ethically staggered release, seems to accurately describe the AI industry today.

But what about AETHER? Interestingly, this does not seem to have achieved more ethical AI either. While AETHER still exists, in early 2023, Microsoft laid off its entire Ethics and Society team, which was reportedly tasked with ensuring that the abstract ethical principles developed by groups such as AETHER were "actually reflected in the design of the products that ship"-in other words, operationalization (Schiffer and Newton 2023). The Ethics and Society team generated ethical output, including exercises, games and frameworks which operationalized ethical principles (Lane 2020). Its dissolution followed the 2020 firing of AI

ethicist Timnit Gebru from Google. And like Gebru, the Ethics and Society team was critical of their employer's new interest in generative machine learning models, pointing out concerns around the DALL-E image generator developed by OpenAI in collaboration with Microsoft. According to the recording of a meeting obtained by *Platformer*, Microsoft Vice President of AI John Montgomery told the Ethics and Society team that the priority was to "move them [OpenAI's models] into customers hands at a very high speed" (Schiffer and Newton 2023). The market takes priority and ethics is relegated to a secondary status. Both of these examples show how when the "coercive laws" of the market manifest, AI ethics and whatever incidental innovations it might produce quickly become a luxury.

Operationalization also appears outside of Big Tech. Canca (2020), founder of consulting company AI Ethics Lab, has published one operationalization procedure in the *Communications of the Association for Computing Machinery*. Canca (2020) holds that "operationalized AI principles for ethical practice will also help organizations confront unavoidable value trade-offs and consciously set their priorities". Canca's method for operationalization is to identify core versus instrumental ethical principles. He sees as core those principles which have intrinsic value, such as human autonomy, while instrumental principles, such as privacy, derive from core values. This dichotomy can be deployed to judge whether an ethical situation involves core or instrumental values, and whether the core and instrumental values at hand are related or not (e.g. whether the instrumental value derives from the relevant core value or another). However, Canca (2020) notes, when "core principles point in opposite directions, we face a real ethical dilemma". In such a case, his solemn advice is that an "ethics expert should be brought in to apply ethical theories". This framework acknowledges "unavoidable value trade-offs" but does not recognize that values, of any sort, might come into conflict with the exigencies of capital accumulation. Core values might come into conflict with one another, but never with the capitalist organization of the economy. Operationalization proceeds here by completely ignoring the industrial context of AI.

One more example comes also from academia/industry. Publishing in academic journals, Morley et al. (2020; 2021) offer the notion of ethics-as-a-service. This approach to operationalization happens to be the same one sold by the AI ethics consulting company Digital Catapult, which funded the research on which the papers are based. The ethics-as-service approach consists of three elements: ethical principles, "a reflective development process" (Morley et al. 2021, p.246) and a distributed system of responsibility shared between internal actors (employees) and external actors (an ethics board). By asking developers to reflect on their design decisions, and by distributing responsibility

among many actors, the goal is to transcend the rigidity of AI principles and achieve a flexible, context-sensitive ethical perspective on AI.

Reflection is surely a good thing for AI development, but is it sufficient to overcome the insoluble contradiction within AI ethics? Morley et al. (2020) admit that it is "hard" to encourage the adoption of ethical AI tools to "practically-minded ML developers, especially when the competitive advantage of more-ethically aligned AI is not yet clear" (p.2161). They go on to elaborate:

> Taking the time to complete any of the 'exercises' … and investing in the development of new tools or methods that 'complete the pipeline', add additional work and costs to the research and development process. Such overheads may directly conflict with short-term, commercial incentives … Unless a longer-term and sector-wide perspective in terms of return on investment can be encouraged (Morley et al. 2020, p.2161).

In other words, doing AI ethics presents an obstacle to doing AI business, unless somehow capital can become other than capital. This, as both Marx and Friedman recognized, is impossible. Thus while offering operationalization as a solution, Morley et al. admit that it is no solution at all. Yet, Morley et al. (2021) argue that AI ethics is "not futile" because "the experience of other applied ethics fields (for example, medical ethics and research ethics) shows that it is possible to operationalise abstract ethical principles successfully" (p.244). However, the comparisons made here are inapt, as Mittelstadt (2019) explains, in that most research occurs, to a large degree, in public institutions and medicine is a unique field in which the "interests of patients and medical practitioners remain aligned at some fundamental level which encourages solidarity and trust ... Comparable solidarity cannot be taken for granted in AI development" since it is "largely developed by the private sector".

In sum, the possibilities for the operationalization of AI ethics are predetermined by the requirement that they acknowledge the priority of capital accumulation as a given. Operationalization implements principles which are amenable to capital, so it is to be expected that operationalization does not modify any underlying operations in AI production or deployment. When operationalization comes into direct conflict with business operations, it is easily discarded.

9 Conclusion

In a discussion of Boltanski and Chiapello's *New Spirit of Capitalism*, Jarrett (2022) describes their contention that "the 'amorality' of capitalism requires that it have enemies" (p.117). They hold that it is only by responding to the criticisms of its enemies that capitalism can generate "the moral

foundations that it lacks" (Boltanski and Chiapello 2005, p.163). Yet, capitalism can only accept certain moral foundations which do not contradict the imperative to increase capital. Adaptation has hard limits. AI ethics recruits capital's enemies to contribute to a subordinated innovation network, but a peculiar one that dares not speak its true purpose, and necessarily fails in its expressed purpose. We might ruefully call it an immanently stymied innovation network.

AI ethics as it stands is a dead-end enterprise. Any interesting AI ethics needs to begin from a perspective which does not prioritize the needs of capital or accept them as given. This would be an AI ethics which acknowledges the fundamental and insoluble contradiction between the accumulation of capital and human flourishing in its myriad forms, one of which is the ethical sphere. The insoluble contradiction within AI ethics should be brought to the fore and explicitly addressed.

Data availability The datasets analysed in the current study are not publicly available due to ethics guidelines on privacy but are available from the corresponding author on reasonable request.

Declarations

Conflict of interest The corresponding author states that there is no conflict of interest.

References

Abdurahman, JK (2020) On the Moral Collapse of AI Ethics. Medium. https://upfromthecracks.medium.com/on-the-moral-collapse-of-ai-ethics-791cbc7df872 Accessed 20 Jan 2022.

Abend G (2014) The moral background: an inquiry into the history of business ethics. Princ Univer Press 11:332

AlgorithmWatch (2020) AI Ethics Guidelines Global Inventory, AI Ethics Guidelines Global Inventory. https://inventory.algorithmwatch.org Accessed 23 Nov 2022.

Allen DWE, Potts J (2016) How innovation commons contribute to discovering and developing new technologies. Int J Comm. https://doi.org/10.18352/ijc.644

Arcesati R (2021) Lofty principles, conflicting incentives: AI ethics and governance in China. Mercator Institute for China Studies. https://merics.org/en/report/lofty-principles-conflicting-incentives-ai-ethics-and-governance-china

Arvidsson A (2010) Speaking out: The ethical economy: new forms of value in the information society? Organization 17(5):637–644. https://doi.org/10.1177/1350508410372512

Baran PA, Sweezy PM (1966) Monopoly Capital: An Essay on the American Economic and Social Order. NYU Press, New York

Birkinbine B (2020) Incorporating the Digital Commons: Corporate Involvement in Free and Open Source Software. University of Westminster Press, London

Blackman R (2020) A Practical Guide to Building Ethical AI. Harvard Business Review. 15 October. https://hbr.org/2020/10/a-practical-guide-to-building-ethical-ai. Accessed 9 Nov 2022.

Bollier D (2014) Think Like a Commoner: A Short Introduction to the Life of the Commons. New Society Publishers, Gabriola Island

Boltanski L, Chiapello E (2005) The New Spirit of Capitalism. Verso, London

Braverman H (1998) Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century. NYU Press, New York

Canca C (2020) Operationalizing AI ethics principles. Commun ACM 63(12):18–21. https://doi.org/10.1145/3430368

Chen J, Storchan V, Kurshan E (2021) Beyond Fairness Metrics: Roadblocks and Challenges for Ethical AI in Practice. arXiv:2108.06217 [cs] [Preprint]. http://arxiv.org/abs/2108.06217 Accessed 14 Oct 2021.

de Jong MDT, Huluba G, Beldad AD (2020) Different shades of greenwashing: consumers' reactions to environmental lies, half-lies, and organizations taking credit for following legal obligations. J Bus Tech Commun 34(1):38–76. https://doi.org/10.1177/1050651919874105

Dignum V (2019) Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way. Springer Int Publish Cham. https://doi.org/10.1007/978-3-030-30371-6

Dyer-Witheford N, Kjøsen AM, Steinhoff J (2019) Inhuman power: artificial intelligence and the future of capitalism. Pluto Press.

Ebell C, Baeza-Yates R, Benjamins R, Cai H, Coeckelbergh M, Duarte T, Hickok M, Jacquet A, Kim A, Krijger J, MacIntyre J, Madhamshettiwar P, Maffeo L, Matthews J, Medsker L, Smith P, Thais S (2021) Towards intellectual freedom in an AI Ethics Global Community. AI and Ethics 1(2):131–138. https://doi.org/10.1007/s43681-021-00052-5

Elias J (2023) Google employees criticize CEO Sundar Pichai for 'rushed, botched' announcement of GPT competitor Bard. CNBC. 10 February. https://www.cnbc.com/2023/02/10/google-employees-slam-ceo-sundar-pichai-for-rushed-bard-announcement.html

Foroohar R (2018) Year in a Word: Techlash. Financial Times. 16 December. https://www.ft.com/content/76578fba-fca1-11e8-ac00-57a2a826423e

Freeman RE, Martin K, Parmar B (2007) Stakeholder Capitalism. J Bus Ethics 74(4):303–314. https://doi.org/10.1007/s10551-007-9517-y

Friedman M (2007 [1970]) The Social Responsibility of Business Is To Increase Its Profits. In Ch.Zimmerli W, Richter K, Holzinger M (eds) Corporate Ethics and Corporate Governance Springer, Berlin.

Google (nd.a) Responsible AI. https://cloud.google.com/responsible-ai Accessed 5 Jan 2023.

Google (nd.b) Our Principles. https://ai.google/principles/. Accessed 9 Nov 2022.

Green B (2021) The contestation of tech ethics: a sociotechnical approach to technology ethics in practice. J Social Comp. 2(3):209–225. https://doi.org/10.23919/JSC.2021.0018

Greene D, Hoffmann AL, Stark L (2019) Better, Nicer, Clearer, Fairer: A Critical Assessment of the Movement for Ethical Artificial

Intelligence and Machine Learning. Proceedings of the 52nd Hawaii International Conference on System Sciences. 2122–2131.

Hagendorff T (2020) The Ethics of AI Ethics: An Evaluation of Guidelines. Mind Mach 30(1):99–120. https://doi.org/10.1007/s11023-020-09517-8

Hardt M, Negri A (2001) Empire. Harvard University Press, Cambridge

Heikkilä M (2022) Responsible AI has a burnout problem, MIT Technology Review. https://www.technologyreview.com/2022/10/28/1062332/responsible-ai-has-a-burnout-problem/. Accessed 28 Oct 2022.

Hu L (2021) Tech ethics: speaking ethics to power, or power speaking ethics? J Social Comp 2(3):238–248. https://doi.org/10.23919/JSC.2021.0033

IBM Cloud Education (2022) AI Ethics. https://www.ibm.com/cloud/learn/ai-ethics. Accessed: 23 Nov 2022.

Jarrett K (2022) Digital Labor. Polity, Cambridge

Jobin A, Ienca M, Vayena E (2019) The global landscape of AI ethics guidelines. Nature Mach Intel 1(9):389–399. https://doi.org/10.1038/s42256-019-0088-2

Johnson RC (2018) Measuring global value chains. Ann Rev Eco 10(1):207–236. https://doi.org/10.1146/annurev-economics-080217-053600

Lane, M (2020) Responsible Innovation: The Next Wave of Design Thinking. Microsoft Design. 19 May. https://medium.com/microsoft-design/responsible-innovation-the-next-wave-of-design-thinking-86bc9e9a8ae8

Liao R (2023) China's ChatGPT rival Baidu Ernie is off to a rough start. TechCrunch. 16 March. https://techcrunch.com/2023/03/16/baidu-chatgpt-launch-china/

Mandel E (1986) In defence of socialist planning. New Left Rev I(159):5–37

Martin C (2016) Shipping Container. Bloomsbury, London

Marx K (1990) Capital, vol I. Penguin, London

Metzinger T (2019) Ethics washing made in Europe. Der Tagesspiegel Online. 8 April. https://www.tagesspiegel.de/politik/eu-guidelines-ethics-washing-made-in-europe/24195496.html. Accessed 19 Jan 2022.

Microsoft (nd.a) Responsible AI principles from Microsoft. https://www.microsoft.com/en-us/ai/responsible-ai. Accessed 9 Nov 2022.

Microsoft (nd.b) Responsible AI Resources. https://www.microsoft.com/en-us/ai/responsible-ai-resources. Accessed 23 Nov 2022.

Mittelstadt B (2019) Principles alone cannot guarantee ethical AI. Nature Mach Intel 1:501–507

Morley J, Floridi L, Kinsey L, Elhalal A (2020) From what to how: an initial review of publicly available ai ethics tools, methods and research to translate principles into practices. Sci Eng Ethics 26(4):2141–2168. https://doi.org/10.1007/s11948-019-00165-5

Morley J, Elhalal A, Garcia F, Kinsey L, Mökander J, Floridi L (2021) Ethics as a service: a pragmatic operationalisation of AI ethics. Mind Mach 31(2):239–256. https://doi.org/10.1007/s11023-021-09563-w

Morozov E (2019) Capitalism's New Clothes. The Baffler. https://thebaffler.com/latest/capitalisms-new-clothes-morozov. Accessed: 23 Mar 2023.

Moss E, Metcalf J (2019) The Ethical Dilemma at the Heart of Big Tech Companies. *Harvard Business Review*. 14 November. https://hbr.org/2019/11/the-ethical-dilemma-at-the-heart-of-big-tech-companies. Accessed 20 Apr 2022.

Newholm T (2017) Relocating the 'Ethical Consumer.' In: Norman R (ed) Ethics and the Market. Routledge, London and New York, pp 162–184

Noonan J (2003) Can There be Applied Philosophy Without Philosophy? Interchange 34(1):35–49

Ochigame R (2019) The Invention of 'Ethical AI': How Big Tech Manipulates Academia to Avoid Regulation. The Intercept. https://theintercept.com/2019/12/20/mit-ethical-ai-artificial-intelligence/. Accessed 23 Nov 2022.

Petrova M (2018) We traced what it takes to make an iPhone, from its initial design to the components and raw materials needed to make it a reality. CNBC. 14 December. https://www.cnbc.com/2018/12/13/inside-apple-iphone-where-parts-and-materials-come-from.html. Accessed 4 Jan 2023.

Phan T, Goldenfein J, Mann M, Kuch D (2021) Economies of Virtue: The Circulation of 'Ethics' in Big Tech. Sci Cult. https://doi.org/10.1080/09505431.2021.1990875

Pitts, FH (2017) Critiquing capitalism today: New ways to read Marx. Springer.

Rahman I, Park J, Chi CG (2015) Consequences of 'greenwashing': Consumers' reactions to hotels' green initiatives. Int J Contemp Hosp Manag 27(6):1054–1081. https://doi.org/10.1108/IJCHM-04-2014-0202

Rikap C (2021) Capitalism, Power and Innovation: Intellectual Monopoly Capitalism Uncovered. Routledge, London and New York

Rikap C (2022) Amazon: a story of accumulation through intellectual rentiership and predation. Compet Chang 26(3–4):436–466. https://doi.org/10.1177/1024529420932418

Rikap C, Lundvall BÅ (2021) Tech Giants' Corporate Innovation Systems. In: Rikap C, Lundvall BÅ (eds) The Digital Innovation Race: Conceptualizing the Emerging New World Order. Springer International Publishing, Cham

Robles Carrillo M (2020) Artificial intelligence: From ethics to law. Telecommunications Policy. https://doi.org/10.1016/j.telpol.2020.101937

Rosenberg M, Confessore N, Cadwalladr C (2018) How Trump Consultants Exploited the Facebook Data of Millions. The New York Times. 17 March. https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html. Accessed 30 Nov 2022.

Schiffer Z and Newton C (2023) Microsoft just laid off one of its responsible AI teams. Platformer. 14 March. https://www.platformer.news/p/microsoft-just-laid-off-one-of-its

Schumpeter J (1943) Capitalism. Socialism and Democracy, Routledge, London

Sherman N (2023) Google's Bard AI bot mistake wipes $100bn off shares. BBC. 8 February. https://www.bbc.com/news/business-64576225

Sloane M (2019) Inequality Is The Name Of The Game. Thoughts On The Emerging Field Of Technology, Ethics And Social Justice. In Proceedings of the Weizenbaum Conference 2019 Challenges of Digital Inequality - Digital Education, Digital Work, Digital Life. 1–9. Berlin

Steinhoff, J. (2021) *Automation and Autonomy: Labour, Capital and Machines in the Artificial Intelligence Industry*. Springer Nature

Suwandi I, Jonna RJ, and Global, Bellamy Foster J (2019) Global Commodity Chains and the New Imperialism. Monthly Review. 1 March. https://monthlyreview.org/2019/03/01/global-commodity-chains-and-the-new-imperialism/ Accessed 4 Jan 2023.

Treanor P (2005) Neoliberalism: origins, theory, definition. http://web.inter.nl/users/Paul.Treanor/neoliberalism.html

Tsing A (2009) Supply chains and the human condition. Rethink Marx 21(2):148–176. https://doi.org/10.1080/08935690902743088

Vakkuri V, Kemell K, Kai-Kristian K, Kultanen J, Abrahamsson P (2020) The current state of industrial practice in artificial intelligence ethics. IEEE Softw 37(4):50–57. https://doi.org/10.1109/MS.2020.2985621

Vică C, Voinea C, Uszkai R (2021) The emperor is naked. Információs Társadalom. 21(2):83–96. https://doi.org/10.22503/inftars.XXI.2021.2.6

Wallach W, Allen C (2008) Moral machines: teaching robots right from wrong. Oxford University Press, New York

Walsh S (2021) Marx, subsumption and the critique of innovation. Organization. https://doi.org/10.1177/13505084211015377

Warren B (1972) Capitalist planning and the state. New Left Rev I(72):3–29

World Bank (2020) World development report 2020: trading for development in the age of global value chains. World Bank, Washington, DC. https://doi.org/10.1596/978-1-4648-1457-0

Zuboff S (2019) the age of surveillance capitalism: the fight for a human future at the new frontier of power. Profile Books, London