

Deep Learning's Governmentality

The Other Black Box

Jonathan Roberge/Kevin Morin/Marius Senneville

Introduction

Frank Pasquale's 2016 book *The Black Box Society* is now considered a landmark study in law-related disciplines, in the social sciences and beyond. The topic in itself, digitalization and the invasiveness of the internet, is of the utmost importance. The book reveals the inversion of operational secrecy by digital platforms and the extensive access to users' private data. Facebook, Google and the like collect and aggregate bits and bytes of information to create massive profiling schemes, the *modus operandi* of which little is known. Problematically, as private actors, they acquire massive data and "knowledge" about the society and our individual behaviors that we don't. To that complex and most certainly unpleasant reality, Pasquale's analytical rigor and finesse contribute valuable insights on potential regulations and on the possibility of developing a smarter citizenry. However, we want to argue that there is another, broader and maybe more cultural reason why *The Black Box Society* draw so much attention—thus further explaining the success of the book. The image of a concealed, networked entity is evocative of some common fears. It captures a sense of "loss of control" vis-à-vis the latest automation processes. Such an algorithmic black box, in other words, taps into a diffused anxiety regarding what is to be called a "known-unknown", i.e. something we recognize to be a mostly hidden form of knowledge production. The image of a black box is a disenchanted one; here lies its strength *as well as* its weakness.

While mostly in line with Pasquale's effort to decipher the opaqueness of our data-driven world, it also appears significant to question the limits of the black box as a heuristic if not holistic image. Scholars such as Geiger (2017), Sudmann (2018), Burrell (2016) or Bucher (2016) have explored this territory. The latter, for instance, has argued that "the widespread notion of algorithms as black boxes may prevent research more than encouraging it", noting that the notion is "too readily used" (84). She then calls for critical scrutiny of algorithms and algorithmic systems using a three-step method: i) "do not fear the black box"; ii) "do not expect the solution to be inside" and iii) "consider the boxing of the box". Whereas

the first step could be understood as encompassing the entire process, the others could be conceived as forming a complementary pair, together examining the inside and outside of the box. Moreover, such an approach is particularly suited to analyzing the recent shift towards deep learning algorithms, but also to machine learning techniques of different sorts, and everything nowadays labelled artificial intelligence. How and to what extent these algorithms are practically and symbolically different from the previous so-called ‘generations’. What would they entail in terms of opacity, ambiguity and vagueness? Where, what, and whom should we look at to develop critical insight and robust interpretation? These questions, once docked to Bucher’s steps, could serve as guide to this chapter.

Examined closely, the idea of “not expecting the solution to be inside” matches perfectly with the ingrained logic and historical development of deep learning algorithms. “Open sesame!” is a task that cannot be programmed, or “learned” for that matter. Despite its biological inspiration and the romantic-teleological accounts of the field’s historical development (Rosenblatt, 1958; Hinton & al. 2006), the fact remains that what stands for “learning” is in fact adaptation and self-tweaking. The mathematical structure modifies itself while interacting and coping with the data stream coming from the outside world (Litvinski 2018; L’heureux & al. 2017). Backpropagation, recursive loops and other subtleties thus not only reflect but also enact a reality in flux. Another way of looking at such uncertainty is with the discrepancy between the more classical symbolic approach to AI and today’s connectionist or neo-connectionist shift (Cardon & al. 2018). Whereas the first relied on deduction, explicit modeling, abstract rules and programmable languages to create a logical and formal mode of reasoning, the second is based on induction, whereby connected hypotheses and approximations produce “optimized” perceptions and predictions about what is going on in the data, inasmuch as data translates into improved rates of predictability (Mackenzie 2017, Sudmann 2018). Layers of non-linear calculus thus inform something of a “deep” but shallow architecture which does not necessarily form an inexplicable AI, but which pushes the limits of its explicability further away. If not fully black, the box of current AI is very grey, to say the least. This can also be seen in the problems scholars are now facing concerning the reproducibility of small-scale theories, where current practices of publication generally prevent them from sharing both source code and training database, or to address the hazards related to the arbitrary setting of hyperparameters, or even the unavoidable randomness inherent in the process of generating training values (Hutson 2018). Managing and massaging that much data is never an easy task, especially not in an experimental environment, and even less in the real world where the saying “garbage in, garbage out” remains thoroughly valid. The multiple problems nowadays with bias fall under the same

category, and could serve as one last example here, namely that the box cannot by its very definition be the solution to something bigger than itself.¹

The internal problems sketched above might not even compare to what is at stake with Bucher's insight probing to "consider the boxing of the box". In fact, there is a long history of social sciences in general and in STS in particular to devote a great deal of attention to everything surrounding a given piece of technology (Bijker & al. 2012). In the specific case of deep learning algorithms, it is all the more important to remember that "they are embedded in larger, far more complex assemblages" that never cease to influence their shape and content (Gillespie 2014: 3). The question, then, is how to make sense of such molding pressures and the kind of opacity they produce. It is about the complexity of a given context or a given "ecology", yet we want to argue that the best way to consider such boxing is through a *networked* approach. As stated elsewhere, "[...] there is not one box, but multiple boxes. The opacity of algorithms is more precisely expressed in different forms of opacity, all of which, in specific ways, are contingent on the *in-betweenness* of a plethora of actors, both human and non-human" (Roberge & Seyfert 2018: 2; Latour 1987). First, it is difficult not to acknowledge an intense division of labor within this domain of innovation—a situation that often translates into developers working on a dataset without fully knowing for whom, to which end and why. Here agency is divided amongst many little hands. Second, a networked approach would consider the actual implementation of deep learning tools and techniques as more in flock than in a row, adjusting to one another more than collaborating. This has been well documented in the literature on algorithmic finance, for instance, where competing stakeholders deploy "algotrading" tools to bolster attack or defense maneuvers (Seyfert 2018; Knorr-Cetina & Preda 2011; Castelle & al. 2016). Whether social sciences will be more attentive in the future to the combined, butterfly-like effects of all these actors' efforts into "boxing the box" remains to be seen. However, considering the understudied state of this phenomenon we urgently need to give more attention to the management, ordering and decision processes that shape what deep learning algorithms come to be about in the real world. More straightforwardly, the political economy of AI is one of our biggest and most opaque boxes today. In this chapter, we intend to contribute to the ongoing debate by analyzing what is at stake in this new form of socio-technical governmentality, i.e. what are the tensions, struggles, efforts at coopting knowledge, power, etc. Taking the Montreal AI hub as a case study, and following a 2016-2018

¹ Of late, IBM has announced that it would allow access to a library comprising over two million images for facial recognition training with the hope that enhanced accuracy would help curb bias. The position of NGOs such as the *American Civil Liberty Union* in that case and in other similar ones is that better facial recognition is still bad news for minorities facing discrimination across a variety of social settings. See Browne 2019.

ethnographical investigation², we will focus on how stakeholders deploy multiple strategies and resources, including the building of legitimacy through symbolically-laden media operations. Our is thus empirical, while also network approach being theoretically informed; in that sense we hope to answer calls by preeminent scholars to develop critical thinking through studies which are *in situ* (Kitchin 2014; Mackenzie 2018).

I. Governmentality—What about it, and what Does it Change to the Study of Deep Learning?

Debates surrounding the increasing complexification of today's political economy and how it should translate into new understandings of power gain a great deal of intelligibility once one adheres to Michel Foucault's concept of governmentality (2004a, 2004b). This practice-oriented analysis of the 'problem of government' has allowed scholars to re-orient their focus on "the ceaseless transactions which, variably, modify, displace, upset, or insidiously shift the funding, the investment modalities, the centers of decision, the forms and types of control, the relations between local and central authorities, etc." (Foucault, quoted in Lascoumes 2004: 3; our translation.) Indeed, the French philosopher has had a pivotal role in identifying—at least—these three logics: i) how power and knowledge are inseparable, ii) how these introduce mobile and networked dynamics, and iii) how all of this allows to think about authority as enacted *by* and *as a* set of technologies. That said, the difficulty with Foucault is that he never properly wrote about the digital. Of late, it is Mackenzie who has endeavored to apply the concept to the study of what he calls 'machine learners', i.e. naive Bayes classifiers, decision trees, neural networks, and a range of others that fall under the broad category of AI (2013; 2017; 2018). All of this, according to him, corresponds to a "data practice that re-configures local centers of power and knowledge by redrawing human-machine relations" (2017: 9). How research, development and implementation is organized; by whom, for what purposes, and through which means and discourses, is different from London to the Silicon Valley, or China and Canada for that matter. Likewise, how power relations and the distribution of authority are shaped specifically by the structure of its organizations and institutions varies from subfield to subfield—finance, military, transport, etc. Mackenzie is thus very helpful by providing such ecosystemic, if not ecological views. At the same time, in his book, he runs the risk of over-emphasizing an internal examination of the technology,

2 We totalized 12 interviews with machine learning specialists, 4 additional ones with scientific journalists, and over 400 articles from local francophone and anglophone newspapers and monthly publications.

and thus is only partially able to 'consider the boxing of the box'—to refer to Burcher once again.

What would it take to be able to provide insights that would be both local and 'architectural'—that is, able to demonstrate how particular constructions and transformations occur from the outside in? One such way is by looking at the distance between governmentality and what is deemed today as 'governance', and how, in fact, the latter is the topic of the former. In Montreal and probably elsewhere, the discourse related to governance enjoys great momentum, as the idea itself serves as a sort of empty-signifier that can tactically be given meaning. Governance, like progress or innovation, readily means "*good governance*" and many stakeholders involved in the construction of the Montreal hub conflate the two in order to bolster the institutional-public support for market-oriented developments in deep learning, the details of which will be presented shortly. For now, suffice to say that the very idea of a deep learning governance in Quebec's metropolis seeks to implicate pretty much everyone as "partners" in a game of collective self-management and purposive social change. In play is what scholars such as Walters have identified as "[an] emphasis on self-governing networks" drawing heavily on "the imagery of cybernetics and complexity theory" (2004: 29-30; see also Simard 1979 for a similar theoretical approach applied to Québec). Power and authority here are conceived as enablers: they allow for the circulation of resources, not for their constraint or restriction. As will be made clear below, everything related to ethics—the industry-backed Partnership on AI or the Montreal Declaration for a Responsible Development of AI—is tainted by an idea of self-regulation and its distinctive way of translating into a loose, sickly effort to *not* legislate. Power and politics have not disappeared for that matter; while governance might present itself in the best light, as lightweight *government at a distance*, the point is that it represents itself as an efficient, if understudied form of governmentality.

What is it about the Montreal deep learning hub that makes it worthy of scientific analysis? Part of the answer relates to the fact that Quebec is a rather small society, well developed but still marked by the concentration of its elites—social, political, economic, cultural, etc. As for the historical context in which the province has addressed the most recent "AI awakening", it is important to recall the role of Canada's CIFAR in subsidizing deep learning research, even when the technique was highly unfashionable (Hernandez 2014; Cardon & al. 2018; Engemann and Sudmann 2018). Star scientists such as Hinton (University of Toronto), former students Bengio (Université de Montréal) and to a lesser extent³, LeCun (NYU and now Facebook) are both the inheritors and the best promoters of what is now a C\$125 million pan-Canadian AI strategy. When, for instance, the talk of

3 LeCun has worked less in Canada and more in France and USA in recent years, although he still enjoys important media coverage.

an ongoing “AI revolution” emerged in Quebec’s francophone mediascape, Bengio himself came to be introduced as one of the leaders of Montreal’s AI ecosystem, itself presented as one of the most world-renowned hubs of cutting-edge AI innovation (See Bourgault 2017). This peculiar dynamic can be further demonstrated by the fact that its name appears in 126 of 161 articles focusing on AI developments published by Montreal newspaper *La Presse* between May of 2016 and July of 2017 (Bourgault 2017). The point here is that, when considering the Quebec’s AI field, two correlations appear clearly: firstly, between the emergent rhetoric of a revolution and the rise of a charismatic leader; and secondly, between the accentuated hype surrounding deep learning and AI and the capacity of local actors to rapidly set in motion the relevant institutions. “Hype is low on informative content,” scholar Guice rightly observes, “but directly states the relevance of the information to a social context” (1999: 85). In order to bolster the Montreal hub, former Quebec’s Economy and Innovation Minister Anglade noted that her government would not “sprinkle” public investment (Rettino-Parazelli 2017). That led first to the creation of an advisory committee and, subsequently, of an AI Cluster initially equipped with a budget of C\$100 million. The two most interesting facts about the cluster is that it devoted 80% of its funds to Bengio-directed, Université de Montréal-led MILA (Montreal Institute for Learning Algorithms) all while being officiated by Breton—Université de Montréal’s dean—and well-known businessman Boivin, who several months later also became the head of MILA’s board (see IA.Québec 2018). This suggests that in this particular context, and in this rather short period of time, what good governance meant was delivering efficiency; whereas a broader, more reflexive and critical perspective would instead have interrogated what it means in terms of circulating elites, and why the effort to maximize efficiency still needs to justify and legitimate itself through at least the appearance of duly-conducted administrative processes.

Another way of considering ‘the boxing of the box’ in the Montreal case is to have a look at the conjunction between efforts geared towards the launch of the aforementioned Montreal Declaration for a Responsible Development of AI and the creation, in late 2018, of the International Observatory on the Societal Impacts of AI. Fully endorsed by the government and its main scientific institutions, both make claims to an epistemological posture of “knowledge co-construction” with the public, the different stake-holders, etc., that in practice serves as a malleable, if not shallow, signifier. The Declaration, for instance, proposes a list of ten principles that are all more general and abstract than the other, with some overly naive or in contradiction with the current economic reality of deep learning—Principle 6.2 for instance states that “AI development must help eliminate relationships of domination between groups and people based on differences of power, wealth, or knowledge” (IA responsible 2017). For its part, the Observatory is still nascent, but in its very constitution already signals a poor understanding of social sci-

ences' role in studying social impacts, with for instance more members coming from computer sciences than sociology and communication studies altogether. Importantly, what the Observatory and the Declaration have in common is the cybernetic view of governance introduced above. On the one hand, the management of knowledge production obliterates any notion of checks and balances or arms-length regulatory principles, notions central to the very idea of modernity. On the other hand, it appears that all current virtue signaling efforts, including the Declaration, the Partnership for AI and the like, emerge as what Wagner calls "an escape from regulation" (2018).⁴ All in all, the Quebec government's involvement in the development of its Montreal hub is one not of creating barriers and obstacles, but rather one to usher and foster the circulation of whatever is deemed 'positive', namely any twists and turns that exhibit a form of action from the government or the stakeholders, knowing that the legitimacy of who gives reflects on who receives and vice versa.

II. The 'Triple Helix' Remix and the Role of Open Science

At this point, it would be tempting to declare that, in spite of the initiatives of numerous actors, including significant gestures by the government of Quebec, it still is "business as usual". This, however, would be misleading in at least two separate ways. First, while it is accurate to say that the Cluster, the Declaration and the Observatory all participate in building a certain public perception of everything AI, it is not possible to adequately fit it to an ideology that would hide any sort of naked truth.⁵ In other words, to be critical is to question how the box is made, not to put it on fire. Second, the expression 'business as usual' undercuts how much the advent of deep learning and related AI techniques is changing the power-knowledge topography of the province, notably the pivotal role universities are called upon to play. A governmentality approach must therefore be attentive to the structuration as well as the tensions involved here—which is also to say the historical and geographical subtleties that make higher education in Quebec something both North American but also profoundly influenced by the French universalistic approach,

4 Many have indeed noticed how increasingly frequent calls for "ethical AI" from industry figures often correlate with ongoing campaigns against "overly coercive" government regulation; see both Wagner, 2018 and Greene & al., 2019. A recent variation on this theme seems to be industry-backed regulations (Simonite, 2019); already, accusations of "regulatory capture" have been expressed (Biddle, 2019).

5 Such a tradition finds an emblematic figure in the early Habermas while he was for instance saying that "[...] [ideologies] replace traditional legitimations of power by appearing in the mantle of modern science and by deriving their justification from the critique of ideology. Ideologies are coeval with the critique of ideology" (Habermas 1971, 99).

and how, starting in the 60s and 70s, it made a substantial push towards a democratization of access.⁶ Interestingly, Université de Montréal and McGill University, both at the forefront of MILA, are historically considered more 'elite' schools while still enjoying a great deal of public support. The MILA itself is important not only because it attracted most of the government-backed Cluster's money, but also as it comes to embody the displacement and, really, the refinement of what the literature calls the triple helix—a schematic model of innovation where corporate actors come to mesh with university and government ones (Etzkowitz & Leidessdorff 2000)—to now a “quadruple helix” where start-ups, too, are considered key strategic partners. Confusing small and large, it is not rare in Montreal to see international corporations such as Microsoft being equated with the local Maluuba, Facebook being considered as the emerging FAIR-MTL or Google as an embryonic DeepMind—with media celebrating even the smallest of investments.⁷ All of this participates in what we argue is an ecological mentality that blurs the symbolic—a hub is positive by its very nature—and the practical, by the virtue of the latest trend in what Hoffman and others have called “academic capitalism” (2017; see also Slaughter and Rhodes 2010). In turn, the reality corresponds less to the early French influence on Quebec's higher education system than to a mode of “Silicon Valley-isation” or “Stanford-isation”, terms borrowed from Salter (2018).

Common to all AI developments is the fact that they are guided by and inseparable from a specific ethos or model of “open science” (Leonelli 2013; Mirowski 2018). Researchers see the sharing of information as, *prima facie*, progress in and of itself; discovery and innovation are meant to be picked up by and benefit the entire “community” in what is thus an ecological as well as cybernetic mentality which, again, has roots in a certain Californian “rebelliousness”.⁸ Today, these norms prove to be very efficient, especially with regards to the following three dimensions. First, the obligation to choose to pursue either an academic career or a career in private R&D becomes less of an issue when one can publish freely, which is now allowed, if not encouraged in most basic research-inclined industrial labs—in fact, it is not rare to see papers co-authored by scientist at Facebook or DeepMind along with university-affiliated researchers. Knowing the shortage of qualified personnel in AI, such open science practices are thus instrumentally adopted

6 The ten institutions networked across the territory under the umbrella of Université du Québec is emblematic in that regard.

7 See for instance the summary of investments made to the local ecosystem in 2017 in Mathys 2017.

8 Here, we want to refer to what Saxenian (1994) and others have described as the characteristically innovative way Silicon Valley academic and industry actors had to produce new organizational forms at an impressive rate. On many accounts, this distinctive way of establishing collaborative ties between actors pertaining to different professional categories but to a common cultural background has spearheaded the privileged understanding of how to lead technological innovation these days—see also Storper & al. 2015.

as part of the repertoire necessary to navigate the “war to attract talent” (Hernandez & King 2016; Metz 2016a; 2017). Second, unrestricted circulation of people and ideas should allow for companies to track the best of university research. Internships, grants, and philanthropic donations large or small, contribute to secure access to computer science labs and to reach researchers where they are. For a city such as Montreal, this has proven very helpful, even if, from this decentralization and openness, it is impossible to conclude that its hub is a “plaque-tournante”—after all, others like Paris, Singapore, Pittsburg, etc., have benefited too. Thirdly, this openness is not only geographical, but temporal, as the adaptation between the different helixes, companies and university labs in particular, is intended to happen more or less in real-time. The pace of research here is as important as the commercial turnover rate that transforms an algorithmic architecture into an API, an innovation in a recommendation system, etc. While openness translates into windows of opportunity and good timing into fierce competition between companies, it is especially important to understand that the logic sustaining the entire model really is one of “strategic openness” (Ananny & Crawford 2018). The knowledge being produced in universities turns out to be “open for business” in a new and understudied sense, especially with regards to its wider implications in contexts such as Quebec.

The fact that deep learning and associated AI technologies signal a substantial displacement of wealth, prestige and power finds numerous and all the more empirical examples to which we will come in a few moments. For now, however, it appears that a necessary transition implies to question the broader significance of the “exploitable epistemology” (Levy & Johns 2016) set in motion through the quadruple-helix and open science nexus. As part of this research, a series of interviews with individuals involved in AI in Montreal were conducted, with most expressing largely consensual views, except for two or three more critical figures. The first one came from a computer scientist working in healthcare. Her critique points to structural elements in the transformation of research financing in Quebec and Canada:

It's a concentration of millions of dollars, it's as if you're betting on a single number at the casino roulette. There's a variety of different types of research done, not all from the deep learning or big data strain [in AI] but that are also innovative—but you're not betting on them, you're only betting on deep learning. You're pushing everyone in the same direction and you forget that innovation is not necessarily of all going in the same direction. You also need to leave some to be sure that research in its totality is somewhat diversified. That, I see as a threat. It's going to siphon everything in the same direction [...]. In fact, everyone is rushing into it.

This comment could serve as a proxy or hint as to how and why even scientific institutions partake in the kind of self-fulfilling prophecy that makes deep learning a reality. Again, institutions, hype and the pressure towards “Stanfordized” research go hand in hand. Another key example would be the attribution of 29 CIFAR Research Chairs late in 2018, some to prominent Facebook associates such as Pineault (McGill) or Vincent (Université de Montréal) (CIFAR 2018). For less trendy research streams, of course, this draws a path in which difficult access to funding would blend with its equivalent in terms of strenuous access to students—at least two other computer scientists in public universities from our samples talked about how they barely have any grad students nowadays. Looking at the longer term, chances are that the situation will become only more cyclical and detrimental.

Another related issue emerging with respect to the meaning of the “exploitable”, even weaponized, epistemology implicated here, concerns the handling of databases: who owns them, how are they released, and for what purposes. For Big Tech companies as Google, who just open-sourced GPipe, or Microsoft, who acquired and now runs Github, gigantic libraries of data are acting in both performative and legitimating ways. Their flaws and limitations are scarcely if ever exposed—the fact for instance that such companies still pursue patents aggressively (Simonite 2018)—especially when compared to the ecological and cybernetic benefits attributed to these platforms and widely praised in the media. It is then at a more mezzo or local level that things get more challenging. The problem is that open data for training is not exactly the same as “real-deal” data or value-added data. For instance, our interviewee who works with deep learning applications in healthcare insisted that a dangerous dynamic is developing, where start-up businesses search for any sizeable bases to access in exchange for deep learning services, or at least, make a contract allowing them to share data with a third party. In places such as Montreal, to make a profit means finding clients—insurance, banks, clinics, biotech, etc.—not yet accustomed to deep learning techniques, in a legal environment still unsure about the best way to defend privacy or to regulate any potential wrongdoing.⁹ Yet, it is probably at the micro level of the different university labs that the difference between the data “haves” and “have nots” is the most striking. Star researchers such as Bengio in Montreal—or, for that matter, Hinton in Toronto—attract funding and students because of their close connec-

9 An important parallel should be established with the way failed unicorn Theranos capitalized on the biotech industry’s important regulatory leeway to position itself as one of the biggest (if short-lived) success stories of this emergent field. Its ability to rely on the reputability of early-backers such as Gen. James Mattis, Oracle founder Larry Ellison, media mogul Rupert Murdoch or present-day Secretary of Education Betsy DeVos to sustain increasing investment rounds should be understood precisely as the result of the field’s relative newness and its lack of proper regulations (O’Brien 2018).

tion with Google and the like. But what about the other, lesser-known researchers in the field? The conundrum is that they almost never gain access to the data actually prone to broad commercial applications. To give one final example in this section, our team met with another scientist in the summer of 2017 and talked about the general sense of community, and what it meant that open science was a way for private and public actors in the field to communicate. His answer was laconic: "It's just fake. It's just fake. They share the algorithm but not the data, you can do nothing with this. [As for the meaning of "open"], it's just a word because I cannot use it".

III. Deep Learning is Redefining the Private-Public Partnership

To say that today's developments in the AI field's political economy blur the preexisting distinctions between what is deemed private and public—or, for that matter, that it amounts to a "Stanfordisation" of higher education in places such as Quebec—is not to succumb to any nostalgia for a utopian past. A descriptive and agnostic approach is indeed needed to account for, as Hoffman stated, "the complicated, subtle, and sometimes contradictory ways that commercial logics have diffused across academic culture" (2017: 727). The point is that, in Montreal and most probably elsewhere, ambiguity is in itself a form of governmentality. Weakened institutional autonomy is translated into more collaboration; buzzwords in the semantic region of "hub", "clustering", "ecosystem" and the likes are repeated and celebrated in what is then hard to decipher from public relations endeavours (see Turkina 2018, for instance). A turning point of this development was the January 2019 relocation of the MILA to Mile-Ex, a post-industrial inner-city in Montreal. The relocation of the lab occurred as it got elevated to the status of "Quebec Artificial Intelligence Institute" and came to be positioned at the forefront of the Mile-Ex's Cité de l'IA, with multiple small and big companies establishing their new facilities either in the same building or in its immediate surroundings. O Mile-Ex, the converted textile-manufacture the lab moved in, already accommodated the offices of up-and-coming startup Element AI, Royal Bank of Canada's AI branch Borealis, French military contractor Thales' AI research division and the para-public Institute for Data Valorization (IVADO), with Microsoft's Maluuba also a close neighbour (Bachand 2018; Dubuc 2018). Importantly, the idea to create the Cité represents the fourth pillar of the government's strategy in everything AI—along with the Industrial Cluster, the Declaration and the Observatory—yet, because of its weight in terms of jobs, investments, square feet of office space and the like, it is possible to argue that it is the most important. The people in charge there understand rather well the leverage associated with their interstitial position. Indeed, in interviews with media about the relocation, they were keen to ask

for additional public funding: “Attracting researchers to Montreal by telling them we only have two years of funding left, that won’t work. We need to be part of a broader, longer-term vision. We’re in the order of tens and hundreds of millions” (Pisano quoted in Rettino-Parazelli 2019).

In terms of practical, yet non-official public-private blending, there might be no equivalent in Quebec to Element AI, the fast-growing company co-founded by MILA’s director Bengio. As rightly expressed by one media commentator, “the business model is not easy to understand” (quoted in Mercure 2016); not only did it attract historic amounts of venture capital without a proper product on the market, but it continually operates under an ethics-oriented discourse of public good and social benefits while also positioning itself as an active player in the rather traditional and profit-savvy fields of logistics, insurance and banking (The Economist 2017). Bengio himself appears willing to play on both levels as he dedicates genuine efforts to promote an ethical and socially-minded development of his field while lionizing the commercial success of his company, one apparently set to become one of the first Canadian AI Unicorn (George-Cosh 2018; Vara 2018). In addition, he sometimes confuses his own numerous public and private affiliations in talks, Power Points and elsewhere, in what is now emblematic of a bigger issue, namely how the value and wealth created in the public domain tends to move away from it. The very nature of Element AI—and part of the reason for its initial valuation—is to capitalize on its access to star academics to develop ‘business solutions’ for its private-sector clients. As acclaimed in the *Journal of Small Business & Entrepreneur*, the company has “a faculty fellow network composed of over 20 world-renowned AI scientists from the top academic labs across Canada. These professors not only do research-related work for the company but [...] provide valuable advice [...]. This unique arrangement gives Element AI access to cutting-edge research” (Turkina 2018: 2). Again, what there is in this quotation relates to everything cybernetic about the new model being implemented discussed above. Helixes rotate, openness signals access, pace equals circulation and innovation, etc., in a movement that is certainly difficult, yet not impossible, to track.

While Element AI is cybernetic by essence, it is as well ecological in a very practical way. Proximity to the MILA shapes the urban space around it, and could be measured in meters. Of course, such proximity is not something to be found only in the Cité de l’IA; numerous incubators in North America, Europe and elsewhere use the model with the justification that it contributes to the cross-pollination of ideas and resources. The problem, however, is slightly different when it comes to the blurring of public and private assets. Emblematic in that regard is a Facebook post by Element AI saluting the arrival of MILA “to the neighborhood”, which showcased a picture of Bengio while emphatically adding, “see you in the stairwell” (Element AI 2018). Such metaphor usually refers to a more or less licit space; one with more or less fuzzy codes and boundaries. Who goes up, what goes

down, when, and how? In the case particularly of students-becoming-interns-becoming-students, the lack of explicit limitations never cease to be problematic both on an individual and on a cohort basis. Once aggregated, these public and private part-time or twofold affiliations reinforce a model that is poorly checked and balanced, especially in the face of its long-lasting socio-political and economic impacts and how these could be discussed and amended in the public sphere.

At current, it is such public-private intermingling that comes to colonize the many layers of the AI Montreal hub—despite certain pleas from Bengio against a reality he actually contributes to.¹⁰ The recent wave of investments made by foreign corporations in the Montreal hub has been going hand-in-hand with the increasing adoption by the newly 'partnered' scientists of this new organizational arrangement, namely, the dual affiliation model. This university-to-industry collaborative form, imported from the fields of law, management and medicine—and probably at its strongest in the biotechnology industry; see Mirowski 2012—allows scholars to keep their university professorship appointment while adding to it a commitment, on at least a part-time basis, to their new corporate employer (Serebrin 2017a; 2017b). To the list parsed throughout this chapter, we should still be adding the many names of MILA-affiliated scientists such as Larochelle at Université de Montréal and Google Brain, Precup at McGill and DeepMind or Pal at Element AI and Université de Montréal. Dual affiliation is justified by actors of the field as a novel solution where scientists are able to continue teaching and conduct basic research while also participating in industrial R&D, whereas previously, such participation would entail a complete retreat from their university teaching and basic research activities (Plamondon Emond 2017). The growing dissemination of the model thus operates at the junction of two distinct but concomitant dynamics. On the one hand, corporate actors are increasingly aware of the necessity, for their business model, to achieve an all-essential balancing-act between the preservation of the "ecosystem sustainability"—i.e., to ensure the continued formation of future generations of AI researchers and the further advancement of basic research endeavours (LeCun 2018)—and, as described in section II, the conflicting urge of immediate appropriation of specialized human resources (Metz 2016a; 2017). On the other hand, scientists are responding to constraints which are mostly presented as incentives: besides the alluring possibility of alternative sources of private funding, researchers also have to deal with the fact that access to state-of-the-art corporate computational infrastructures and some of the widest proprietary databases are indeed technological means increasingly needed for the pursuit of cutting-edge deep learning research. In turn, this new public-pri-

¹⁰ See Shead 2018. On his criticizing the increasing concentration by major tech corporations of both technological means and specialized human resources while scarcely rejecting opportunities to collaborate with them—see Mathys 2017 and Vara 2018.

vate assemblage and its peculiar way of providing many solutions at once is presented, justified and legitimized as a form of necessity, and not as a cascade of contingent choices and principles. In that sense, however, it is nonetheless highly political and a form of governmentality.

Conclusion

This chapter started by acknowledging how the black box is a powerful, yet disenchanted figure to reflect on technologies *in the making* such as deep learning, and AI more broadly. While there is an urgency that should spur social sciences inquiry, it is nonetheless important to do things right, with a certain dose of agnostic and critical reflexivity. In this light, we attempted to follow Bucher's triple advice to 'not fear the black box'; 'not expect the solution to be inside' and 'consider the boxing of the box'. So how did it go? How did the theoretical concepts apply to the practical reality and how, in turn, can better understanding of a case such as the Montreal hub inform broader and more critical reflection? Parts of the answer came in section I, where the argument was made that what is mostly at stake is the present and future political economy of AI, i.e. how the automation of knowledge production transforms power relations and how the different actors involved in deep learning are engaged in what Crandall names a particular form of "cooperative struggles" (2010). Substantial resources including money, state support, media coverage, etc. are flowing and aggregating, the details of which are precisely what must be understood about this dense and tense regime of governmentality. The new normal brought about by deep learning and AI-related technologies will be messy and ambivalent, if this is not already the case. We insisted throughout the chapter that power is more than ever a transaction and that what "control" means in these circumstances relates to a new sense of cybernetics and ecology that shall account for all types of mutualism and parasitism. In section II, we described this by digging into the Montreal example, especially how it exhibits a peculiar form of rotary motion between the helixes that are the governmental, university, established and upcoming corporate actors. Whereas actors repeatedly proclaim there is a "community" and that the Cluster, the Observatory, the Declaration and the *Cité de l'IA* make for an integrated whole, we propose a somewhat less optimistic, more realistic analysis. Open science is a case in point, as "open" translates into aggregation and as it signals an important shift in the educational model in vogue. There is such thing as a privatisation of higher education in place like Quebec, in which deep learning and AI related technologies are instrumental. This was the principal conclusion of section III. Whether you call it the "double affiliation" or the "see you in the staircase" model, what is clear is that the benefits are not equally redistributed at current and have very poor chance—at least the-

oretically—of being so in the future. In the end, it might then be such unfolding, in its many twists and turns, that constitutes the proper object of another, still-in-the-making ecology, one which could be called Critical AI Studies.

References

Ananny, Mike, and Crawford, Kate (2018 [2016]): "Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability." *New Media & Society* 20/3, pp. 973-89. <https://doi.org/10.1177/1461444816676645>.

Bachand, Olivier (2018): "Une Cité de l'intelligence Artificielle Dans Le Mile-Ex à Montréal." Radio-Canada, retrieved from <https://ici.radio-canada.ca/nouvelle/1081709/technologie-montreal-saint-urbain-intelligence-artificielle> (accessed March 11, 2019).

Biddle, Sam (2019): "Should We Trust Artificial Intelligence Regulation by Congress If Facebook Supports It?" The Intercept, retrieved from <https://theintercept.com/2019/03/07/artificial-intelligence-facebook-ibm-ro-khanna/> (accessed March 9, 2019).

Bijker, Wiebe E., Hugues, Thomas Parke, and Pinch, Trevor (eds.) (2012 [1987]): *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. Cambridge, Mass: MIT Press

Bourgault, Mathieu (2017): "Synthesis on the Quebec's AI Mediascape". *Unpublished Manuscript*, pp. 1-30.

Browne, Ryan (2019): "IBM Hopes 1 Million Faces Will Help Fight Bias in Facial Recognition." CNBC, retrieved from <https://www.cnbc.com/2019/01/29/ibm-releases-diverse-dataset-to-fight-facial-recognition-bias.html> (accessed March 9, 2019).

Bucher, Taina (2016): "Neither Black Nor Box: Ways of Knowing Algorithms." In *Innovative Methods in Media and Communication Research*, edited by Sebastian Kubitschko and Anne Kaun, pp. 81-98. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-40700-5_5

Burrell, Jenna (2016): "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms." *Big Data & Society*, 3/1, pp. 1-12. <https://doi.org/10.1177/2053951715622512>.

Cardon, Dominique, Cointet, Jean-Philippe, and Mazières, Antoine (2018): "La revanche des neurones: L'invention des machines inductives et la controverse de l'intelligence artificielle." *Réseaux* 211/5, pp. 173-220. <https://doi.org/10.3917/res.211.0173>

Castelle, Michael, Millo, Yuval, Beunza, Daniel, and Lubin, David C. (2016): "Where Do Electronic Markets Come from? Regulation and the Transfor-

tion of Financial Exchanges.” *Economy and Society* 45/2, pp. 166-200. <https://doi.org/10.1080/03085147.2016.1213985>.

CIFAR (2018): “29 Researchers Named to First Cohort of Canada CIFAR Artificial Intelligence Chairs.” CIFAR, retrieved from <https://www.cifar.ca/cifarnews/2018/12/03/29-researchers-named-to-first-cohort-of-canada-cifar-artificial-intelligence-chairs> (accessed March 9, 2019).

Crandall, Jordan (2010): “The Geospatialization of Calculative Operations: Tracking, Sensing and Megacities.” *Theory, Culture & Society*, 27/6, pp. 68-90. <https://doi.org/10.1177/0263276410382027>.

Dubuc, André (2018): “Microsoft Déménage Ses 200 Employés Dans Le Mile-Ex.” La Presse, retrieved from http://plus.lapresse.ca/screens/60fd222-94fb-49ee-ac0a-05a933c166ao__7C__o.html?utm_medium=Twitter&utm_campaign=Internal+Share&utm_content=Screen (accessed March 11, 2019).

Element AI (2018): “See you in the stairwell”. *Facebook Post*, January 29, retrieved from <https://www.facebook.com> (accessed March 9, 2019).

Engemann, Christoph, and Sudmann, Andreas, eds. (2018): *Machine Learning. Medien, Infrastrukturen und Technologien der Künstlichen Intelligenz*. Bielefeld: transcript.

Etzkowitz, Henry, and Leydesdorff, Loet (2000): “The Dynamics of Innovation: From National Systems and ‘Mode 2’ to a Triple Helix of University–Industry–Government Relations.” *Research Policy* 29/2, pp. 109-23. [https://doi.org/10.1016/S0048-7333\(99\)00055-4](https://doi.org/10.1016/S0048-7333(99)00055-4).

Foucault, Michel (2004a): “Naissance de La Biopolitique.” In *Cours Au Collège de France 1978-1979*. Paris: Gallimard-Seuil.

Geiger, R Stuart (2017): “Beyond Opening up the Black Box: Investigating the Role of Algorithmic Systems in Wikipedian Organizational Culture.” *Big Data & Society* 4/2, pp. 1-14. <https://doi.org/10.1177/2053951717730735>.

George-Cosh, David (2018): “Element AI Aims for Unicorn Status with Record Canadian Financing: Sources.” BNN Bloomberg, retrieved from <https://www.bnnbloomberg.ca/element-ai-aims-for-unicorn-status-with-record-canadian-financing-sources-1.1101206> (accessed March 9, 2019).

Gillespie, Tarleton (2014): “Algorithm [Draft] [#digitalkeywords].” *Culture Digitally* (blog), retrieved from <http://culturedigitally.org/2014/06/algorithm-draft-digitalkeyword/> (accessed March 9, 2019).

—(2017): “Regulation of and by Platforms.” In *SAGE Handbook of Social Media*, edited by Jean Burgess, Thomas Poell and Alice Marwick, pp. 254-78. London: Sage.

Greene, Daniel, Hoffman, Anna Lauren and Stark, Luke (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*.

Guice, Jon (1999): "Designing the Future : The Culture of New Trends in Science and Technology." *Research Policy*, 28, pp. 81-98.

Habermas, Jürgen (1971): *Toward a Rational Society: Student Protest, Science and Politics*. London: Heinemann.

Hernandez, Daniela (2014): "Meet the Man Google Hired to Make AI a Reality." *Wired*, retrieved from <https://www.wired.com/2014/01/geoffrey-hinton-deep-learning/> (accessed March 9, 2019).

Hernandez, Daniela, and King, Rachael (2016): "Universities' AI Talent Poached by Tech Giants." *Wall Street Journal*, retrieved from <http://www.wsj.com/articles/universities-ai-talent-poached-by-tech-giants-1479999601> (accessed March 9, 2019).

Hinton, Geoffrey, Osindero, Simon and The, Yee-Whye (2006): "A Fast Learning Algorithm for Deep Belief Nets." *Neural Computation* 18/7, pp. 1527-1554.

Hoffman, Steve G. (2017): "Managing Ambiguities at the Edge of Knowledge: Research Strategy and Artificial Intelligence Labs in an Era of Academic Capitalism." *Science, Technology, & Human Values* 42/4, pp. 703-40. <https://doi.org/10.1177/0162243916687038>.

Hutson, Matthew (2018): "Artificial Intelligence Faces Reproducibility Crisis." *Science* 359/6377, pp. 725-726. <https://doi.org/10.1126/science.359.6377.725>.

IA.Québec (2018): "Effervescence de l'intelligence Artificielle." Comité d'orientation de La Grappe En Intelligence Artificielle, retrieved from <http://ia.quebec/> (accessed March 9, 2019).

IA responsable (2017) : "Déclaration de Montréal Pour Un Développement Responsable de l'IA." Déclaration de Montréal IA Responsable, retrieved from <https://www.declarationmontreal-iaresponsable.com/> (accessed March 9, 2019).

Kitchin, Rob (2014): "Big Data, New Epistemologies and Paradigm Shifts." *Big Data & Society* 1/1, pp. 1-12. <https://doi.org/10.1177/2053951714528481>

Knorr-Cetina, Karin and Preda, Alex (eds.) (2011 [2006]): *The Sociology of Financial Markets*. Oxford: Oxford University Press.

Lascoumes, Pierre (2004): "La Gouvernementalité : De La Critique de l'État Aux Technologies Du Pouvoir." *Le Portique*, no. 13-14, pp. 1-16.

Latour, Bruno (1987) : "Science in Action: How to Follow Scientists and Engineers through Society" Cambridge, Mass: Harvard University Press.

LeCun, Yann (2018): "Facebook's Chief AI Scientist Says That Silicon Valley Needs to Work More Closely with Academia to Build the Future of Artificial Intelligence." *Business Insider*, retrieved from <https://www.businessinsider.com/facebook-yann-lecun-dual-affiliation-model-ai-experts-2018-8> (accessed March 9, 2019).

Leonelli, Sabina (2013): "Why the Current Insistence on Open Access to Scientific Data? Big Data, Knowledge Production and the Political Economy of Contemporary Biology." *Bulletin of Science and Technology Studies* 33/1-2, pp. 6-11. <https://doi.org/10.1177/0270467613496768>.

Levy, Karen EC, and Johns, David Merritt (2016): "When Open Data Is a Trojan Horse: The Weaponization of Transparency in Science and Governance." *Big Data & Society* 3/1, pp. 1-6. <https://doi.org/10.1177/2053951715621568>

Lheureux, Alexandra, Grolinger, Katarina, Elyamany, Hany F. and Capretz, Miriam A. M. (2017): "Machine Learning with Big Data: Challenges and Approaches." *IEEE Access* 5, pp. 7776-7797.

Litvinski, Oleg (2018): "On Social Mechanisms of Algorithmic Opacity." *Unpublished Manuscript*, pp. 1-32.

Mackenzie, Adrian (2013): "Programming Subjects in the Regime of Anticipation: Software Studies and Subjectivity." *Subjectivity* 6/4, pp. 391-405.

— (2017): *Machine Learners: Archaeology of a Data Practice*. Cambridge, MA: The MIT Press.

— (2018): "From API to AI: Platforms and Their Opacities." *Information, Communication & Society*, June, pp. 1-18. <https://doi.org/10.1080/1369118X.2018.1476569>.

Mathys, Catherine (2017): "Le Rebelle de l'intelligence Artificielle." L'actualité, retrieved on <http://lactualite.com/techno/2017/11/10/le-rebelle-de-lintelligence-artificielle/> (accessed March 9, 2019).

Mercure, Philippe (2016): "Montréal, Future Plaque Tournante de l'intelligence Artificielle?" La Presse, retrieved from <https://www.lapresse.ca/affaires/economie/201610/26/01-5034411-montreal-future-plaque-tournante-de-lintelligence-artificielle.php> (accessed March 9, 2019).

Metz, Cade (2016a): "Giant Corporations Are Hoarding the World's AI Talent." Wired, retrieved from <https://www.wired.com/2016/11/giant-corporations-hoarding-worlds-ai-talent/> (accessed March 9, 2019).

— (2016b): "GOOGLE OPENS MONTREAL AI LAB TO SNAG SCARCE GLOBAL TALENT." Wired, retrieved from <https://www.wired.com/2016/11/google-opens-montreal-ai-lab-snag-scarce-global-talent/> (accessed March 9, 2019).

— (2017): "FOR GOOGLE, THE AI TALENT RACE LEADS STRAIGHT TO CANADA." Wired, retrieved from <https://www.wired.com/2017/03/google-ai-talent-race-leads-straight-canada/> (accessed March 9, 2019).

Mirowski, Philip (2012): "The Modern Commercialization of Science Is a Passel of Ponzi Schemes." *Social Epistemology* 26/3-4, pp. 285-310.

— (2018): "The Future(s) of Open Science." *Social Studies of Science* 48/2, pp. 171-203. <https://doi.org/10.1177/0306312718772086>.

O'Brien, Sara Ashley (2018): "Elizabeth Holmes Surrounded Theranos with Powerful People." CNN Business, retrieved from <https://money.cnn.com/2018/03/15/technology/elizabeth-holmes-theranos/index.html> (accessed March 9, 2019).

Pasquale, Frank (2015): *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, Massachusetts: Harvard University Press.

Plamondon Emond, Étienne (2017): "Intelligence Artificielle: Un Pied à McGill et l'autre Chez Un Géant Des Technos." Le Devoir, retrieved from <https://www.>

ledevoir.com/societe/education/511873/intelligence-artificielle-un-pied-a-mc-gill-et-l-autre-chez-un-geant-des-technos (accessed March 9, 2019).

Rettino-Parazelli, Karl (2017): "L'intelligence Artificielle, Moteur Économique." Le Devoir, retrieved from <https://www.ledevoir.com/economie/500340/dominique-anglade-au-devoir> (accessed March 9, 2019).

—(2019): "Nouveaux Bureaux, Nouveaux Besoins Pour Mila." Le Devoir, retrieved from https://www.ledevoir.com/economie/546567/nouveaux-bureaux-nouveaux-besoins?fbclid=IwAR3XDUBbIUrIqnGIt5oXUfZTdp8GqYtY6kbVQsGl8pp-slKzpPQFgv_s2IQ (accessed March 9, 2019).

Roberge, Jonathan and Seyfert, Robert (2018 [2016]). "What are algorithmic cultures?" in *Algorithmic Cultures: Essays on Meaning, Performance and New Technologies*, Robert Seyfert and Jonathan Roberge (eds.), pp. 1-25, New York: Routledge.

Rosenblatt, Frank (1958): "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain." *Psychological Review* 65/6, pp. 386-408.

Salter, Christopher (2018). Personal communication, "AI Talks: The New Silicon Valley of the North, Really?", December 4, Montreal, Concordia University.

Saxenian, AnnaLee (1994): *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Cambridge, Mass: Harvard University Press.

Serebrin, Jacob (2017a): "Facebook to Open AI Lab in Montreal Headed by McGill Professor." *Montreal Gazette*, retrieved from <https://montrealgazette.com/news/local-news/facebook-to-open-ai-lab-in-montreal-headed-by-mcgill-professor> (accessed March 9, 2019).

— (2017b): "Google-Linked AI Company to Open Research Lab in Montreal." *Montreal Gazette*, retrieved from <https://montrealgazette.com/business/local-business/google-affiliated-ai-company-deepmind-to-open-research-lab-in-montreal> (accessed March 9, 2019).

Seyfert, Robert (2018): "Automation and Affect: A Study of Algorithmic Trading." In *Affect in Relation—Families, Places, Technologies. Essays on Affectivity and Subject Formation in the 21th Century*, Birgitt Röttger-Rössler, Jan Slaby (eds.), pp. 197-218. New York: Routledge.

Sudmann, Andreas (2018): "On the Media-political Dimension of Artificial Intelligence: Deep Learning as a Black Box and OpenAI.", in: Rethinking AI. Neural Networks, Biopolitics and the New Artificial Intelligence, *Digital Culture & Society* 4/1, pp. 181-200. (available on [academia.edu](http://www.academia.edu))

Shead, Sam (2018): "AI Pioneer Yoshua Bengio Says Universities Deserve More Credit." *Forbes*, retrieved in <https://www.forbes.com/sites/samshead/2018/11/13/ai-pioneer-yoshua-bengio-says-universities-deserve-more-credit-for-their-ai-research/#5cb7d54273ab> (accessed March 9, 2019).

Simard, Jean-Jacques (1979): "Québec et frères, inc. La cybernétisation du pouvoir." *Recherches sociographiques*, 20/2, pp. 239-261. <https://doi.org/10.7202/055840ar>.

Simonite, Tom (2018): “Despite Pledging Openness, Companies Rush To Patent AI Tech.” *Wired*, retrieved from <https://www.wired.com/story/despite-pledging-openness-companies-rush-to-patent-ai-tech/> (accessed March 9, 2019).

— (2019): “AMAZON JOINS MICROSOFT’S CALL FOR RULES ON FACIAL RECOGNITION.” *Wired*, retrieved from <https://www.wired.com/story/amazon-joins-microsofts-call-rules-facial-recognition/> (accessed March 9, 2019).

Slaughter, Sheila, and Rhoades, Gary (2010): *Academic Capitalism and the New Economy: Markets, State, and Higher Education*. Baltimore: Johns Hopkins University Press.

Storper, Michael, Kemeny, Thomas, Makarem, Naji Philip, and Osman, Taner (2015): *The Rise and Fall of Urban Economies: Lessons from San Francisco and Los Angeles*. Innovation and Technology in the World Economy. Stanford, California: Stanford Business Books.

The Economist (2017): “A Hybrid Startup Offers AI Services to Business.” *The Economist*, retrieved from <https://www.economist.com/business/2017/06/22/a-hybrid-startup-offers-ai-services-to-business> (accessed March 9, 2019).

Turkina, E. (2018): “The Importance of Networking to Entrepreneurship: Montreal’s Artificial Intelligence Cluster and Its Born-Global Firm Element AI.” *Journal of Small Business & Entrepreneurship* 30/1, pp. 1-8. <https://doi.org/10.1080/08276331.2017.1402154>.

Vara, Vauhini (2018): “Can This Startup Break Big Tech’s Hold on A.I.?” *Fortune*, retrieved from <http://fortune.com/longform/element-ai-startup-tech/> (accessed March 9, 2019).

Wagner, Ben (2018): “Ethics as an Escape from Regulation: From Ethics-Washing to Ethics-Shopping?” In *Being Profiling. Cogitas Ergo Sum.*, Amsterdam University Press, pp. 1-7. M. Hildebrandt.

Walters, William (2004): “Some Critical Notes on ‘Governance.’” *Studies in Political Economy* 73/1, pp. 27-46. <https://doi.org/10.1080/19187033.2004.11675150>.