

Labour, Justice and the Mechanization of Interpretation

LARRY LOHMANN

ABSTRACT: The biggest frontier of mechanization of the past ten years has been the automation, broadly speaking, of that particular type of human labour known as *interpretation*. This includes the mechanization of *recognition* (for example, image recognition technologies used by security services), of *translation* (Google Translate), of *searching for information* (Bing or Google search engines), of *understanding* ("predictive algorithms" that know what books or movies you will like or what kind of propaganda will appeal to you, as used by Amazon, Netflix, or the Donald Trump campaign), of *trust* (blockchain technologies such as Bitcoin), and of *negotiation* ("smart contracts" as pioneered by firms such as Ethereum). How do these technologies benefit business and why have they come to prominence now? In what ways are they a continuation of older business strategies for extracting as much value as possible from labour via mechanization, such as steam-powered manufacture? Why are they inextricably entangled with the degradation and exhaustion of the work of both humans and nonhumans? Why are they adding to threats to the earth's climate? And what are the implications for social movement strategy? This article suggests a few directions for research.

Introduction: new mechanizations, new injustices

New waves of mechanization can be both bewildering and worrisome. Few are more so than the flood of computing innovations that has washed over the world in the last decade or so.

The list is impressive. Fabulously accurate machine translation among dozens of languages is now available free on demand. Artificially-intelligent personal assistants populate homes and offices as robots invade more and more factories. Machines routinely shuffle through images to identify your face out of millions of others. A new kind of money called Bitcoin operates by automating, via computer cryptography, the kind of trust you were used to placing in the state. An Internet of Things can now tell the World Wide Web whenever you turn on an appliance. Machines are able to send out individually-targeted advertisements based on unique data sets unceasingly collected from each of a billion human beings.

This isn't even to mention possible future applications like automated contracts that execute themselves (Diedrich, 2016). Or self-operating farms without human farmers (*CNN Business*, 2 April 2019). Or an automated Earth Bank of Codes that could store and sell to businesses digitized information on the genomes and capacities of every plant, animal and single-cell organism on earth (Castilla-Rubio, 2017; *The Economist*, 23 January 2018).

How these technologies actually work is often pretty hard to understand. For example, how is it that hundreds of millions of pieces of data crunched by super-speedy processors enable computers to predict many aspects of your own behavior better than you can yourself? How is new "deep learning" software teaching itself how to foresee the future even better? What are the new developments in cryptography that make Bitcoin possible? Such technicalities are mystifying to a lot of ordinary people.

What's easier to understand are the frightening news stories that occasionally float up to the surface of the new wave of mechanization. A company called Cambridge Analytica is caught using personal information secretly mined from Facebook to find the right audiences to bombard with tailored propaganda for the Trump and Brexit campaigns (Cadwalladr and Graham-Harrison, 2018). China's government is reported to be planning a centralized computer system that assigns a financial and reputation score to each of its citizens (Kobie, 2019). Google's pioneer automated image recognition

system stirs outrage by identifying black people as gorillas, while a similar system being sold to police departments by Amazon is criticized for commonly misidentifying black women as men (Singer, 2019). The World Economic Forum, meanwhile, reports that the data-processing warehouses run by Bitcoin – to say nothing of those belonging to Facebook, Google, Amazon, Microsoft, Apple, IBM, Alibaba, Baidu and Tencent – could suck up the equivalent of the current world's energy production by 2025 (Jezard, 2017; see also Mora, Rollins et al., 2018). We might not understand exactly what's going on when we read such stories, but we sense that something might be wrong. Even that some sort of injustices might be afoot.

But what and where exactly are the injustices?

Is it just that the technologies don't yet quite work the way they're supposed to and need to be improved? On this view the answer is simple: just call in the technicians to fix the machines. Maximize the good things that the new mechanization provides while minimizing the bad.

Or is it maybe just that the technologies are being controlled by the wrong people? If so, the answer again seems simple: put them under democratic control, with proper safeguards in place.

Or maybe the problem is that the benefits of the new wave of automation aren't being spread around equitably. Again, the proposed response is almost a reflex: let's just make sure that the state regulates it right.

But what if the issue of technology injustice with respect to the new computation goes deeper? What if it goes beyond fixing a few circuits or making a few new rules for distributing the goodies? What if it goes so deep that some forms of conventional regulation could make underlying injustices worse? After all, that's already happened with attempts to regulate new financial technologies with new rule books (Dowd, 2014), or efforts to regulate fossil-fuelled technologies with carbon trading and carbon taxes (Lohmann, 2018).

This article will argue that this deeper approach is needed in order to evolve realistic strategies for dealing with injustices connected with the new mechanization. It will contend that the deep issues of injustice with respect to cutting-edge 21st-century computer-led automation are identical to the deep issues of injustice with respect to 19th-century industrial mechanization. It will propose that if we can understand one, we can understand the other – as well as the fundamentals of the social struggles that surround, and will continue to surround, both. And it will suggest that if regulation never had any chance of “fixing” the deeper justice issues connected with 19th-century mechanization, it's not going to have any better luck with the new 21st-century forms.

The article does not deny that there are innumerable injustices in who gets to use and benefit from artificial intelligence or so-called “cryptocurrencies” like Bitcoin – just as there were innumerable injustices in the distribution of the benefits of 19th-century factory or transport machinery. It does not deny that many injustices result from the ways advanced computing can be, and is being, used by assorted bad guys – just as many injustices have resulted from the malevolent use of 19th-century technologies. Nor does it deny that at least a few of the biases, rights violations, skewed effects and accidents associated with the latest wave of automation might be avoided by technical improvements. But it does contend that a sober understanding of the central social questions involved has to start at a more basic level: with an account of the logics of the new mechanization, its origins, and the relations involving the activities of human and nonhuman beings that define it. That account has to begin with the “old mechanization” of the 19th century.

Work and mechanization in the 19th century and after

What has been popularly referred to, since around 1950, by the word “technology” – the constellation of overwhelmingly fossil fuel-powered patterns and processes of mechanization so evident in manufacturing, transport, construction, domestic activities, commerce, extraction, exploration, war and computation – got its start in a 19th-century social movement led by business.

In the long-gone days before fossil fuels and thermodynamic energy became prominent, business hadn't yet really figured out the concept of worker “productivity”: that is, how to increase the output of workers per unit of time. Instead, the usual strategy of firms under pressure to lower costs and sell more goods than the competition was just to make workers work longer hours and chase up new ones who could do the same. But that idea hit a wall with worker resistance and foot-dragging and labour legislation restricting working hours. Replacing conventional factory devices with what later came to be called “technology” was what was needed to change the terms of engagement (Malm, 2016). The rush was on to develop the kind of machine network that could give business more power to dial up worker productivity whenever and wherever worker resistance or lack of cheap wage labour were making things sticky. Here there was an advantage in developing and getting control over a form of energy that was mobile, centralizable, and endlessly interconvertible. Step forward the conversion engine, which could take a stock of concentrated, transportable fuel under business's control and transform it into heat, which could then in turn be transformed into mechanical energy, electricity, magnetism, kinetic energy or heat again, and so on. Naturally the technical details took up a lot of intellectual work: hence thermodynamics, electrodynamics and the whole modern concept of abstract “energy resources” (Lohmann and Hildyard, 2014). But in the end, fossil-powered mechanization gave business unprecedented abilities to get a lot more out of labour in a lot less time. Indeed, in theory, business could now increase indefinitely the profit it made from each worker's hour, the number of workers it could concentrate and organize in one place, and the volume and speed of the sales it could make around the world.

Of course, the mechanization that began in the 19th century has never come free, any more than business' earlier – and continuing – strategy of separating workers from their land has been without costs. But one of the great things for capital about the arrangement has always been how much of these new costs can be covered by the unpaid work of what Maria Mies calls “women, nature and colonies” (Mies, 2014), and how little of the burden need fall on business. That is, not only did 19th-century fossil-based mechanization create millions more paid workers and put them under tighter, cheaper discipline, indeed spreading the wage labour relation throughout the world (Huber, 2009). It also was able to enlist new armies of workers, human and nonhuman, and mostly unpaid, tasked with “tending the machines” in a far broader sense. Industrial capital couldn't survive a single day without people on and off the factory floor, inside and outside offices, constantly maintaining, repairing, cleaning up, updating, adapting, checking, teaching, interpreting, de-digitizing, excusing, and managing and absorbing the wastes from machines – in general, making usable the actions that machines perform even when that means adjusting human actions to conform more closely to a machine model. Indeed, the more machines you have, the more living human work you will need – probably outside your own organization – to ensure that they interface productively with society and don't go “feral”, becoming golems or sorcerer's apprentices. Nor could mechanized business survive a single day without calling on and depleting huge volumes of the work the organisms of the past have provided for free in forming coal, oil and gas, or that the domesticated plants and animals of the present provide in the form of the exercise of the adaptability and flexibility they have acquired via evolutionary interactions that again

predate the age of capital by thousands or millions of years. In expanding the extent and scope of unequal labour relations, in short, 19th-century mechanization also involved the expansion of business control over pools of work stretching across new continental spaces and reaches of geologic time stretching back to the Carboniferous Era and beyond.

Technology, in other words, has never been about “replacing” labour or “saving” it. It’s about increasing it. The overall character of the labour involved may be modified, via “deskilling,” “reskilling,” “simplification” and “diversification” of all kinds, but as long as business seeks “machine fixes” for its never-ending struggles with workers, the sheer volume of the additional human labour it turns out to require to make those “fixes” work will never go down. Still less will the sheer volume of the unpaid nonhuman work it needs to carry on its struggle to make profits.

That’s not to say that specific jobs and occupations are not constantly lost to technology. Nor is it to say that individuals have never been replaced by machines. They often were, and will continue to be, by the millions. The closer the job that you do already imitates routine machine actions (Collins and Kusch, 1999), the more likely it is that machines will be doing it next year. But the overall thrust must always be to enroll and transform more of humanity and nature into business’ workforce. It’s not unlikely that much of the new labour will be what is commonly referred to as shitwork, or low-paid or non-paid precarious work of a devalued or invisible kind: feeding and tending the growing mass of machines (Hornborg, 2011), quietly absorbing wastes, fixing the inabilities of machines, enduring disasters, caring for human casualties. But whatever its character, it is set to grow in volume, and not despite but because of mechanization.

By the same token, the idea that the story of technology is all about “mechanizing tasks” – machines coming to do the “same” things that humans do, only better or faster – is somewhat misleading. The fossil-powered loom is doing something, and that something results in an output of fabric, but the machine is not doing the same thing that the handloom artisan is doing, no matter how similar the cloth at the end of the process might look. Because the fossil loom freezes selected bits of living knowledge about fabrics built up over generations simply in order to execute a few of the rules embedded in that knowledge at precariously high speeds, it needs a different kind of human tending, especially when something unexpectedly goes wrong. It requires quantities of thermodynamic energy that the artisan wouldn’t know what to do with, and that also needs human organization. It demands the professional management of humans and nature in extraction zones, and so on.

Similarly, Google Translate is doing something, and that something results in an output of translated sentences, but it is not doing the “same thing” that human interpreters do, even if one outcome is the identical translated sentence. Instead, Google Translate gloms onto billions of digitized, machine-ready sentences on the internet – representations of oceans of the living work of past and present generations of humans and nonhumans. It then subjects this “big data” to unprecedented computer power in order to predict – probabilistically rather than linguistically – which sentence-to-sentence equivalences would be most acceptable to human translators, especially those working in international business. And it constantly corrects its own procedures on the basis of new data provided free of charge by users of electronic devices around the world. With its giant, publicly-subsidized, energy-hungry data centres, Google Translate, again, needs vast quantities of thermodynamic energy that human interpreters wouldn’t know what to do with, while its wastes call for further armies of uncompensated human and nonhuman cleanup workers. Again, there is no “thing” that stays constant through the process of “being mechanized” any more than there are discrete objects called “technology” that, when sprinkled onto humanity, helps it attain its desires but otherwise leaves it just as it is.

In all of these senses, injustice is not something incidental to technology. It's not something that you can somehow refine out of it. Instead, it's sort of the whole point. Fossil-fuelled mechanization and the thermodynamics that accompanies it exist because 19th-century business had to find a new kind of hold over workers that was capable of sustaining and augmenting the process of getting something for nothing that is essential to the cycle of profit-making, reinvestment and more profit-making. It was part of their design that they were not to be used justly. They were not supposed to be detachable objects that could be put in everybody's hands, or even under any kind of democratic control. Just the opposite: they constituted a better way of compelling ordinary people to contribute to business a lot of things that they didn't necessarily want to contribute, and to reinforce the class divisions needed for capital's cycle. And, of course, they brought with them further cascades of injustice in the form of greater calls on unpaid work and common lands, up to and including the work of the earth in maintaining its climate. As will be explored further below, computer-intensive 21st-century mechanization is animated by the same politics, only applied to new domains. None of these injustices are accidents merely waiting to be regulated away by future generations. To treat them as such is often to undermine communities at the front lines of struggles against them.

This is not to be "anti-technology". Indeed, it's hard to make out what it would amount to for anybody to try to be either "anti-technology" or "pro-technology", or why they would bother. Battles over technology injustice are joined not when people make abstract, indiscriminating pronouncements about how terrible – or how wonderful – electricity or computers or nuclear power are, but rather when concrete efforts are made to engage with particular struggles over mechanization that are actually in progress. Among millions of other examples are movements to "keep the oil in the soil" in Latin America or Africa or to "keep coal in the hole" in Europe – which in Britain date back to 1605 or before (Nef, 1966: 316) – and campaigns to curb the efforts of software-driven "platform capitalists" such as Uber and Airbnb to open up new frontiers of labour exploitation or rent-seeking. Mechanization, in short, is a sprawling, complex arena, constantly being remodeled, where everyone has to find their feet from birth. To note that profitable injustice is its foundational *raison d'être* is merely to state its address.

New fields to mechanize

As the previous section indicates, the computer era of the 21st century has not transcended the conversion engine – the steam engine, dynamo, turbine, electric motor, or internal combustion device that changes one kind of energy into another and whose development went hand in hand with the fossil fuel boom that began in the 19th century. As Andrew Ross observes acerbically, although it has become a pop cliché that data and computer-enabled social networking are the oil of the 21st century, "so far, at least, it looks as if oil is still the oil of the 21st century" (Ross 2013: 29). New technologies like Facebook, Bitcoin, Amazon, Google and the rest have not broken the link between fossil fuels and mechanization, but strengthened it. Instead of getting rid of the heat engine (the archetype of 19th-century mechanization), the new age of automation has merely married it to the Turing Machine, the archetype of 20th-century mechanical computation (Caffentzis, 2013). The clacking, thundering repetitions of 19th-century factory machines, in other words, have not gone away. They have merely been supplemented by the whirring, droning repetitions of ranked thousands of computer processors housed in giant warehouses together with their cooling devices and energy links.

But what are all these new energy-hungry machines doing, exactly? If this is mechanization, what is being mechanized? This might be the hardest thing to grasp about the new technologies. Again, it can be illuminated by analogies from the 19th century.

Imagine, for example, that what was mechanized in the 19th century was the sort of activity seen in this photograph (right) of a Tibetan cotton spinner.

It would be pretty difficult to try to give a “full” description of what’s happening in photos like this. Certainly an exercise of skills, but what skills? Following certain learned routines of eye, finger, thread, wood and arm; making commodities for sale to neighbours or tourist or other markets; clothing and contributing to or sustaining a community; having fun or being sociable; being creative; keeping certain customs and senses of time alive – many things could be part of the picture.



In any case, whatever it is that mechanical spinners do (below), it is quite different. Not only is the machine faster and more “productive” of specific commodities. Not only is it more energy-intensive, with both its inputs and outputs linked to markets and to time in different ways. It also has different qualitative relations to human labour. The people directly taking care of the machine and correcting its errors are exercising different skills and have less opportunity to relax and converse. Equally important, there is a largely-invisible infrastructure of extra labour, paid and unpaid, tending the machine from a distance. Industrial Revolution factories, for instance, tended to need armies of pieceworkers, “outworkers” or contract labourers taking on jobs in their homes that the machines within the factory and their human assistants couldn’t do, or do cheaply enough (Gray and Suri, 2019: 38-63). No less important were the colonial plantations and mines extracting the masses of raw materials that the factories so voraciously consumed at such a high cost in justice to humans and the land.



Despite these differences, the usual convention is to say that fossil-powered spinners and other devices “mechanize” the spinning activity, doing nothing more than codifying and “energizing” some selection of the human being’s spinning skills so they can be repeated faster and more accurately. That oversimplification makes it easier to look at the photo of the human spinner and instead of seeing “livelihood”, “survival”, “flourishing”, or “improvisation”, to see instead “folklore”, “inefficiency”, or “tradition”.

By analogy, what is being increasingly mechanized in the 21st century is the sort of activity communicated by this photo of an ordinary conversation:



Again, it would be pretty hard to give a full description of what’s going on in photos like this one. But we can identify at least some of the familiar skills that are in play. Being able to recognize an individual; being able to interpret or translate what they are saying; being able to understand them (which includes knowing what kind of questions to ask if you think you’ve misheard or have missed part of the context); being able to search for information; being able to argue; being able to negotiate; being able to trust; being able to recognize where you are – all are almost certainly active.

What’s happened over the last decade or so is that it’s become possible for the first time for business to look at these activities and skills in something like the same way it looks at the activities and skills of the spinner – as a lucrative field for mechanization. Business is learning, that is, to see the two people in the photo as performing a kind of labour that can be mechanized to increase profits and rents and speed up circulation while getting more control over the whole work process. Part of the motive, as always, has been crisis. The financial crash of 2008, reduced profits from traditional industrial production, the limits of austerity measures, and the precarious nature of current speculative investments in infrastructure and financial products all give capital good reason to seek new horizons

of profit. On this front, what could be more innovative than finding ways to use mechanization to speed up and exploit realms of human work that might not have seemed before even to *be* work?

One good blanket term for these kinds of newly-mechanized labour is *interpretation*. For example, the mutual work of *recognition* that is a part of the story of the two people conversing in the photo above can now be performed – sort of – by machines that can distinguish our faces and voices from those of others. The potential for profits on the security market alone is enormous.

The labour of *understanding* between the two friends can also be mimicked by machines, as when Amazon gives you shopping suggestions based on data it has collected about you. Indeed, business visionaries are now contemplating speeding up turnover still further by sending customers products before they have even ordered them, since individualized automated forecasts of what they might accept are becoming so cheap and reliable that the economic risk of irate involuntary buyers returning the goods is seen to be rapidly diminishing. If the human work of shopping, already having been speeded up by micro-packaging and self-checkout, could be transferred to machines, business could go straight from producing stuff to shipping it without having to pass through retailing at all.

Not least, schemes are underway to mechanize much of the labour of *trust*, which up to now has been largely a matter of precisely the kind of time-intensive, cumbersome, interpretive, “inefficient” personal relationship-work visible in the photo of the two friends above. The idea is to be able to manufacture trust (or trust-substitutes) and make them available on demand as cheaply and quickly as possible via what *The Economist* magazine (31 October 2015) calls a global “trust machine.” Instead of making it possible to produce huge volumes of cars, paper or cement quickly and cheaply, with minimal trouble from labour, the “trust machine” is designed to make it possible to monitor, execute, record and enforce huge numbers of global *transactions* quickly and cheaply, no matter how tiny, with minimal human oversight.

That seems to make the “trust machine” an ideal lubricant for trade, banking and consuming, directly addressing business’s old objective of reducing transaction costs. Put together cryptocurrencies with mechanized contracts housed on the internet that automatically execute payments when sensors embedded in shipments signal that inspections have been passed, tax paid, or borders crossed, and you have a recipe for speeding the flow of goods along lengthy infrastructure corridors such as China’s One Belt One Road project. True to mechanization’s traditional aims, the technology also makes possible new forms of labour control (along the way making inevitable new forms of labour rebellion). One supply-chain executive, for instance, looks forward to “utilizing artificial intelligence and predictive analytics technology” to shift the task of optimizing deliveries “to algorithms rather than tribal employee knowledge” (*Supply Chain* 24/7, 17 December 2018).

Other business visionaries celebrate a coming age when machines can “encode the collective knowledge of management science” as lucratively as mechanical spinners encoded some of the knowledge of human spinners (Tapscott and Tapscott, 2016: 140). Once “day-to-day decision-making” is “programmed into clever code” (ibid.) running on computers and the internet, whole layers of management could be stripped away from corporations, boosting efficiency and kicking up the neoliberal project of “maximizing shareholder value” one more notch. With workers employed, paid and micro-supervised by algorithms, labour could increasingly be seen on the model of Uber or Glovo, as temporary subcontracting without benefits, with the additional boon that workers couldn’t sue companies for any abuses that had not been written into the computer code beforehand.

But the new mechanization of interpretation offers a lot more than just fresh possibilities for lubricating and controlling commerce. It also promises to expand its range immensely. The technology most fundamental to the “trust machine” – known as blockchain – has already enabled the rise of cryptocurrencies like Bitcoin, which, in facilitating intermediary-free payments, are ultimately supposed to replace the work of many bankers and bureaucrats. But the next step would be an environment in which you could safely own and trade private property without having to involve other human beings at all. Without the need for a burdensome, expensive human infrastructure of trust involving lawyers, judges, marriage alliances, word of mouth, family networks, state agencies, credit card companies, bankers, credit scorers, and so forth, those mechanical contracts housed on the internet – set up very cheaply in their bespoke millions – could become capable of directly, and relatively smoothly, channeling the wealth of some of the (monetarily) poorest people on earth to some of the world’s largest transnational corporations.

For example, a unique mechanized contract located on the internet could securely transfer rights to the genome of a particular forest frog species in an indigenous territory to a faraway drug company while guaranteeing in return that a micro-payment was automatically transferred to the account of the local leader whenever a sale was made of a product based on the frog data. Indeed, any fraction of any entity anywhere – a third of a customary irrigation system in Pakistan, a quarter of the solar electricity produced on a rooftop in Uganda, a fifth of a squatter’s patch of land in Honduras – could become instantly and cheaply exchangeable worldwide for any fraction of anything anywhere else, releasing a continuously-flowing tide of liquid assets the likes of which has never seen before.

Mechanized contracts could also be devised that revised themselves automatically in response to new data from the outside world, mechanically sensed and interpreted without need for human monitors and renegotiators. For example, with a fast, reliable global database constantly feeding fine-grained risk data into millions of mechanically-secure, automatically self-adjusting insurance contracts, insurance firms could reduce the chances that less vulnerable customers – or the firms themselves – end up subsidizing customers that the companies learn too late are more risky than they thought. The large, crudely-rationalized “risk pools” that the insurance and welfare sectors have had to work with in the past in order to spread costs and exposure and stabilize premiums – and to make possible coverage for high-risk as well as low-risk customers – could become more and more outdated.

Some of these plans for automation might seem far-fetched. But let no one imagine that the 21st-century mechanization of interpretation has not already channeled benefits to business as tangible and sizable as those that accrued to the tycoons of earlier versions of mechanization: the textile magnates of 19th-century Manchester, for example, or Henry Ford with his assembly line, John D. Rockefeller with his hold on the requisite energy, and all the other business titans of the golden age of heavy industry. Six out of the ten richest individuals in the world today – Jeff Bezos, Bill Gates, Mark Zuckerberg, Larry Page, Sergey Brin and Larry Ellison – have made their fortunes largely out of the automation of interpretation, and there are surely many more billionaires soon to come out of the same mould.

New means for mechanizing

If the *motives* for today’s new wave of mechanization are becoming clearer bit by bit these days (not least to business itself, which has required and will require a lot of time to grasp them), what is providing the *opportunity*? Who is providing the kind of technical groundwork for the new projects of mechanizing interpretation that figures like James Watt, Lord Kelvin or Thomas Edison provided for the old projects of mechanizing spinning, say, or car manufacture?

Part of the opportunity has been provided by a remarkable “perfect storm” of coinciding technical developments in computation that has arrived during the last decade or so. These include a recent renaissance of the deep-learning approaches mentioned earlier: software that can continuously teach itself what algorithms or recipes are best at predicting whatever it wants to predict.

Also important has been the recent construction of huge libraries or stocks of digitized bits of information out of undigitized flows of human culture – for example, labeled and encoded JPEG images, sentence and sentence-pairs expressed as series of ones and zeros, and so forth. Naturally, immense amounts of human labour have gone into this. At the beginning of its education process, and at certain points afterwards, a machine can’t know that the picture on your smartphone is a picture of Elvis and not an Elvis impersonator unless a human tells it so. Some of this kind of work is paid, if precarious – as when a woman in Hyderabad sitting at her kitchen table with a computer is asked via the internet to second-guess an Uber computer about whether two separate photos used by one of its drivers in Chicago to identify himself in fact match (Gray and Suri, 2019: xv). The great bulk of this labour, however, is unpaid – as when Google Translate or police image-recognition software parasitize the voluntary, almost unconscious linguistic work of hundreds of millions of smartphone owners exchanging gossip or snapshots on Facebook or other platforms. All of this human labour is ongoing, because interpretation itself is an ongoing process, and interpretation machines – irremediably oriented to a past they have no choice but to entomb continuously – won’t be able to make a profit for their proprietors if they can’t keep up with the flow of culture.

A final ingredient of the “perfect storm” of technical computing advances of the last decade are the big increases in computer processing speed and capacity that make it possible to convert all those growing mountains of “big data” into cheap, accurate predictions in breathtakingly short times, as well as carry out the advanced cryptographic operations necessary for non-state currencies like Bitcoin. Indeed, it’s these processors – arranged in racks of thousands upon thousands in giant, custom-built, megawatt-hungry, artificially-cooled data centres hidden in the back landscapes of several continents – that probably provide the best visual analogue to the spinning machine in one of the photos above, with its rows of bobbins performing the same actions over and over again.



As a representation of what it is that is supposedly “mechanizing” the labour of the two people sitting in conversation in the previous illustration, this photograph is probably even more misleading than the picture of the spinning machine. The picture of the servers with their obligatory human attendant simply leaves out too many things. It fails to show the hundreds of millions of teaching-assistant “outworkers” labeling and feeding digitized images and sentences to the machines for free so that they can “learn” to “recognize” things as well as the two people on the sofa. It omits the millions of “ghost workers” (Gray and Suri, 2019) receiving pittance for correcting the machines’ continual, inevitable, multiplying errors and misplaced predictions. The server photo also fails to show the work of countless “iSlaves” (Qiu, 2016) who manufacture the computer hardware required in giant Chinese factories, using minerals extracted around the world in indigenous territories and elsewhere. It fails to reveal the handouts from taxpayers required to set up and maintain the warehouse-like installations that contain and organize the server racks like those in the photo. And it leaves out the vast work of the Carboniferous-Era organisms being used up to run the trillions of repetitive floating-point operations per second that the machines require to deliver a facsimile of human interpretative labour.

But the deceptiveness of the photograph – which, if treated as a symbol of what mechanization is, has a way of taking us all in – does serve a purpose. If we come to believe that machines are able to perform labour – and make people rich – all by themselves, we are deprived of one of the main resources for achieving justice: knowledge of how the world works and how technology performs within it. Examined closely enough, today’s new wave of mechanization may provide fresh materials for dismantling such old illusions.

Conclusion

A recurring romantic story about technology, as the introduction to this special issue of *Development*, suggests, generates “spaces of comfort where we underplay the transformational shifts we need to address.” Successive waves of mechanization, the story goes, may well have brought a lot of pain at first, but in the end, once they had been regulated correctly, were revealed as an essentially benign force that lifted the burdens of all. So, too, in the future, we are told, technology, rather than being integrally related to continuing struggles involving class, gender, race and nature, will continue to develop “from outside” as an independent superhero-like object that might be able to sweep in someday to rescue us from such conflicts and their consequences.

This article has helped list a few of the holes in this romantic story by looking at the innovations in mechanization of the last decade or so in the light of the 19th century. In both periods, it argues, mechanization has been structured in ways that increase labour rather than “saving” it. It provides evidence that the common claim that machines “replace” what humans do, whether viewed as threat or promise, is misplaced. Technology injustice, it adds, isn’t simply a matter of lack of regulation. Rather, it’s bound up in the way both the industrial mechanization of the 19th century and the “information” technologies of the 21st have functioned in exhausting more and more human and nonhuman work in pursuit of competitive advantage for business.

Accordingly, the article suggests, rather than repeatedly telling ourselves that technology injustice is a question of how mechanization is “used”, or what its “implications” are, it might be better to spend more time trying to figure out what mechanization actually *is*. Looking at technologies such as blockchain, image recognition software, machine translation and artificial intelligence in the light of steam engines, early assembly lines, steel mills and industrial plantations may help in doing that.

References

- Cadwalladr, Carole and Emma Graham-Harrison (2018) 'Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach', *The Guardian*, 18 March, <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>.
- Caffentzis, George (2013) *In Letters of Blood and Fire: Work, Machines and the Crisis of Capitalism*, Oakland: PM Press.
- Castilla-Rubio, Juan Carlos (2017) 'Nature-Inspired Design: How the Amazon can help us solve humanity's greatest challenges', World Economic Forum, 25 June, <https://www.weforum.org/agenda/2017/06/bio-inspired-designamazon-technology/>.
- CNN Business* (2019) 'How 5G will Change the Future of Farming', 2 April, <https://edition.cnn.com/2019/04/01/business/5g-farming/index.html>.
- Collins, Harry and Martin Kusch (1999) *The Shape of Actions: What Humans and Machines Can Do*, Cambridge, MA: MIT Press.
- Diedrich, Henning (2016) *Ethereum: Blockchains, Digital Assets, Smart Contracts, Decentralized Autonomous Organizations*, Wroclaw: Wildfire Publishing.
- Dowd, Kevin (2014) 'Math Gone Mad: Regulatory risk modeling by the Federal Reserve', *Policy Analysis* 754.
- Gray, Mary L. and Siddharth Suri (2019) *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*, Boston: Houghton Mifflin Harcourt.
- Hornborg, Alf (2011) *Global Ecology and Unequal Exchange: Fetishism in a Zero-Sum World*, New York: Routledge.
- Huber, Matthew (2009) 'Energizing Historical Materialism: Fossil fuels, space and the capitalist mode of production', *Geoforum* 40 (1): 105-115.
- Jezard, Adam (2017) 'In 2020 Bitcoin Will Consume More Power than the World does Today', World Economic Forum, 15 December, <https://www.weforum.org/agenda/2017/12/bitcoin-consume-more-power-than-world-2020/>.
- Kobie, Nicole (2019) 'The Complicated Truth about China's Social Credit System', *Wired*, 7 June, <https://www.wired.co.uk/article/china-social-credit-system-explained>.
- Lohmann, Larry (2018) 'Should We Put a Price on Carbon? That depends: who are "we"?', presentation for the Panel "Should We Put a Price on Carbon?" Sheffield University Festival of Debate, 11 May, http://www.thecornerhouse.org.uk/sites/thecornerhouse.org.uk/files/SHOULD%20WE%20PUT%20A%20PRICE%20ON%20CARBON%20--%20THAT%20DEPENDS%20--%20WHO%20ARE%20WE_0.pdf.
- Lohmann, Larry and Nicholas Hildyard (2014) *Energy, Work and Finance*, Sturminster Newton: The Corner House, <http://www.thecornerhouse.org.uk/sites/thecornerhouse.org.uk/files/EnergyWorkFinance%20%282.57MB%29.pdf>.
- Malm, Andreas (2016) *Fossil Capital: The Rise of Steam Power and the Roots of Global Warming*, London: Verso.

Mies, Maria (2014) *Patriarchy and Accumulation on a World Scale: Women in the International Division of Labour*, London: Zed Books.

Mora, Camilo, Randi Rollins et al. (2018) 'Bitcoin Emissions Alone Could Push Global Warming above 2°C', *Nature Climate Change* 8: 931–933.

Nef, J. U. (1966) *The Rise of the British Coal Industry*, London: Routledge.

Qiu, Jack Linchuan (2016) *Goodbye iSlave: A Manifesto for Digital Abolition*, Urbana: University of Illinois Press.

Ross, Andrew (2013) 'In Search of the Lost Paycheck', in Trebor Scholz (ed.) *Digital Labour: The Internet as Playground and Factory*, pp 24-51, Routledge: New York

Singer, Natasha (2019) 'Amazon is Pushing Facial Technology that a Study Says Could be Biased', *New York Times*, 25 January, <https://www.nytimes.com/2019/01/24/technology/amazon-facial-technology-study.html>.

Supply Chain 24/7 (2018) 'New Digital Supply Chains are Powered by Artificial Intelligence and Predictive Analytics', 17 December, <https://247customsbroker.com/index.php/2018/12/18/new-digital-supply-chains-are-powered-byartificial-intelligence-and-predictive-analytics/>.

Tapscott, Don and Alex Tapscott (2016) *Blockchain Revolution: How the Technology behind Bitcoin is Changing Money, Business and the World*, New York: Penguin.

The Economist (2018) 'Sequencing the World', 23 January.