



richardheinberg.com

MuseLetter #363 / June 2023 by Richard Heinberg

This month's Museletter is about crisis and acceleration. The first essay, "Polycrisis, Unraveling, Simplification, or Collapse: Coming Soon to a Planet Near You?," introduces a new report from Post Carbon Institute on the roots of the polycrisis and why we should be thinking differently about the future. The second essay, "If You're Driving Off a Cliff, Do You Need a Faster Car?," is a warning about AI and its capacity to accelerate the crises we already face.

Polycrisis, Unraveling, Simplification, or Collapse: Coming Soon to a Planet Near You?

Since the start of the COVID-19 pandemic, Russia's invasion of Ukraine, and the resulting disruption of multiple global supply chains, policy think tanks have increasingly adopted the term *polycrisis* to signify humanity's destabilized status quo. The World Economic Forum's [2023 Global Risk Report](#) uses the newish word 13 times in 90 pages. Scholars from a range of disciplines (including Columbia University historian [Adam Tooze](#)) have written about the polycrisis, and both [Cascade Institute](#) and [Omega Institute](#) have published papers and reports on it. The Cascade Institute notes that "a global polycrisis occurs when crises in multiple global systems become causally entangled in ways that significantly degrade humanity's prospects. These interacting crises produce harms greater than the sum of those the crises would produce in isolation, were their host systems not so deeply interconnected."

Evidence of polycrisis is usually separated into two buckets—environmental and social. Signs of environmental crisis include climate change, the disappearance of wild nature, relentless resource depletion, the increasing chemical pollution of air and water, soil loss and degradation, and fresh water scarcity. Evidence of social crisis includes increasing economic inequality, poverty, racism and other forms of discrimination, the rise of authoritarianism, and impacts of rapid technological change (such as automation).

Our current set of crises can be described as a polycrisis because self-reinforcing feedbacks between ecological breakdown and social breakdown are strengthening and growing more numerous. For example, climate-driven human migration presents challenges to political systems while also eroding traditional cultural norms that support environmental stewardship. Societies in the midst of social crisis, and ones turning toward authoritarian regimes,

are seldom able to muster efforts toward resource conservation, emissions reduction, and habitat preservation; indeed, under such circumstances, past efforts in these directions may be undermined.

All of this is happening in the wake of a couple of decades' worth of [historical studies](#) that show societal collapse to be a normal, predictable, and even inescapable periodic occurrence throughout the past few thousand years. It appears that societies tend to become more [complex](#), develop new technologies, accumulate wealth, and grow more unequal over time. Their leaders start to [quarrel with one another](#), weakening overall social cohesion. Finally, after two or three centuries of this, almost anything can push a society over the brink—a natural disaster, resource depletion, war, insurrection, epidemic, or financial crash. Scholars who engage with the accumulating literature on societal collapse can hardly help noting the relevance for today's world. We've built a global civilization of unparalleled complexity, wealth, and inequality, all based on depleting, polluting fossil fuels. What could go wrong?

An early warning came in 1972 with the publication of [The Limits to Growth](#), a report by MIT system dynamics scientists on their efforts to model the likely future interactions between population growth, consumption growth, and resource depletion. Their computer-based scenarios suggested, under business-as-usual conditions, global industrial society would likely collapse during middle decades of the 21st century.

A new report by [Post Carbon Institute](#) (PCI), [Welcome to the Great Unraveling: Navigating the Polycrisis of Environmental and Social Breakdown](#) (full disclosure: I'm one of the authors), seeks to build a coherent narrative about the roots of the polycrisis, the signs of its arrival and evolution, and why we should be thinking differently about the future. When confronted with evidence that our collective path is unsustainable, many of us tend to jump to “all-or-nothing” ways of thinking, sometimes framing our future in simplistic terms as “the end of the world” or “apocalypse.” But according to the report's authors, this tendency is unhelpful. While a complete and sudden end of humanity is theoretically possible via nuclear war, our more likely future will consist of decades of social, economic, political, and ecological turmoil punctuated by periods of rescue and recovery. There is still considerable divergence between best- and worst-case scenarios, and we still have agency to affect outcomes.

According to the PCI report, we should be spending far less effort building upon expectations of a future that looks much like today only with more technology, mobility, and wealth; instead, we should devote our collective brainpower to questions like, How does a civilization downsize gracefully? Or, What have we achieved that our distant descendants would like us to preserve for them?

Maybe we'd be better off avoiding the word “collapse” altogether, since it tends to be disempowering. Nate Hagens, who interviews polycrisis experts on his [podcast](#), terms the era we are entering “The Great Simplification.” Regardless what we call it, this will be a time that calls for new attitudes and behaviors. Strategies that seemed to make sense before the polycrisis, such as efforts to grow national economies, will need to be replaced by different ones, such as efforts to build resilience. Fortifying resilience at the

community level will be especially important: as global supply chains grow brittle and shatter, humanity will depend more upon local economies for survival and opportunities to thrive. Cooperative strategies to ration scarce resources and reduce inequality will also be required so as to defuse conflict and ensure optimal outcomes for as many as possible.

If humanity descends into blame and desperate efforts to maintain a status quo that by its very nature cannot persist, the future looks dark indeed. Imagine what a young person a few decades from now, living in a depleted and ravaged world, might feel while looking at surviving images of today's "influencers" enjoying comfort, convenience, and privilege on an epic scale. However, if we work together now to build a truly sustainable way of life, maybe future generations will have at least some reasons to thank us.

If You're Driving Off a Cliff, Do You Need a Faster Car? *Artificial Intelligence and the Fate of the World*

Eliezer Yudkowsky, co-founder of the [Machine Intelligence Research Institute](#), thinks artificial intelligence (AI) will kill us all. He frequently poses the following question. Imagine that you are a member of an isolated hunter-gatherer tribe, and, one day, strange people show up with writing, guns, and money. Should you welcome them in?

For Yudkowsky, AI is like a super-intelligent space alien; inevitably, it will decide that we humans and other living beings represent nothing more than piles of atoms for which it can find better uses. "[U]nder anything remotely like the current circumstances," Yudkowsky wrote in a recent *Time* magazine [op-ed](#), "literally everyone on Earth will die. Not as in 'maybe possibly some remote chance,' but as in 'that is the obvious thing that would happen.'"

On May 30, a group of AI industry leaders from Google Deepmind, Anthropic, OpenAI (including its CEO, Sam Altman), and other labs issued a [public letter](#) warning that the technology may one day pose "an existential threat to humanity." For the curious, [here's a brief description](#) of some of the ways AI could wipe us out.

Not everyone thinks of AI in apocalyptic terms. Bill Gates, former chairman of Microsoft Corporation, just sees AI as disrupting the business and tech world, possibly leading to the demise of Amazon and Google. "You will never go to a search site again, you will never go to a productivity site, you'll never go to Amazon again," he recently [told an audience](#) at an AI Forward event in San Francisco. AI will be embedded in products and systems from cars to universities, sensing our intentions and desires before we even voice them, shaping our reality and serving us like a proverbial genie—or an army of them.

Everyone does agree that AI represents a qualitative as well as a quantitative shift in technological development. It's not just an improved computer with more speed and power, but a software architecture that enables computers to teach themselves how to learn, and to continually improve and expand their abilities. AI systems now write computer code, making them, in a sense, self-generating. AI is essentially a "[black box](#)" from which thought-like output emerges; people can't figure out why and how it does what it does after the

fact. Further, AI systems learn from each other almost instantly, taking in vastly more information than any human can. A crucial threshold will be reached with the development of artificial general intelligence (AGI), which could accomplish any intellectual task humans perform, and greatly exceed human abilities in at least some respects—and which, crucially, could set its own goals. Already, [computers can defeat](#) any human chess grand master.

Artificial Intelligence “Duh” Risks

Some AI risks are fairly obvious. Machines will increasingly replace information workers, destroying white-collar jobs (full disclosure: this article was not written by AI, though I did use Google and Bing for research). Inevitably, AI will enrich owners and developers of the technology while others will shoulder the social costs, resulting in more societal wealth inequality. The proliferation of [deepfake](#) images, audio, and text will make it increasingly difficult to tell what’s true and what isn’t, further distorting our politics. And a dramatic expansion of computer number crunching will likely demand [more overall energy usage](#) (though [not everyone agrees](#) on this point).

Then, there is the prospect of accidents. Every new technology, from the automobile to the nuclear power plant, has seen them. [Writing in Foreign Affairs](#), Bill Drexel and Hannah Kelley argue that an AI accident crippling the global financial system or unleashing a devastating bioweapon might most readily happen in China, because that country is poised to lead the world in AI development but seems utterly unconcerned about risks surrounding the technology.

Even if it works exactly as intended, AI will enable already powerful people to do more things, and do them faster. And some powerful people tend to be selfish and abusive. Cognitive psychologist and computer scientist Geoffrey Hinton, who is sometimes called the “godfather of AI,” recently quit Google. In subsequent interviews with multiple news outlets, including the [New York Times](#) and [BBC](#), Hinton [explained](#): “You can imagine, for example, some bad actor like [Russian President Vladimir] Putin decided to give robots the ability to create their own sub-goals.” One of these sub-goals might be, “I need to get more power.”

However, Hinton chose not to endorse another recent [open letter](#), this one calling for a six-month pause in the training of all AI systems (though many of his colleagues in the AI development community did sign on). Hinton [explained](#) that, despite its risks, AI promises too many good things to put it on hold. Among those likely benefits: potential advances in pharmaceuticals, including cures for cancer and other diseases; improvements in renewable energy technologies; more accurate weather forecasts; and a greatly increased understanding of climate change.

High school and college students are already resorting to OpenAI’s [ChatGPT](#) to write their term papers (savvy students give their computer-generated papers a quick re-write in order to defeat [AI-detection software](#) that teachers are now using). Unfortunately for students, their computer-generated papers tend to be [riddled with fake quotes and sources](#). A lawyer representing a client who was suing an airline recently used ChatGPT to write his legal briefs; however, it later turned out that the AI had [“hallucinated” every one of](#)

[the legal precedents](#) it cited. Automobile manufacturers are building cars with more AI-based [self-driving functions](#). Microsoft, Google, and other tech companies are rolling out AI “[personal assistants](#).” [Militaries are investing](#) heavily in AI to make superior weapons, to plan better battle strategies, and even to shape long-term geopolitical goals. Thousands of independent computer labs run by corporations and governments are developing AI for a constantly widening array of purposes. In sum, AI is already far along its initial learning curve. The genie is out of the bottle.

The Acceleration of Everything

Even if Eliezer Yudkowsky is wrong and AI won’t wipe out all life on Earth, its potential perils are not limited to lost jobs, fake news, and hallucinated facts. There is another profound risk that is getting little press coverage—one that, in my view, systems thinkers should be discussing more widely. That is the likelihood that AI will be a significant accelerator of everything we humans are already doing.

The past few thousand years of human history have already seen several critical accelerators. The creation of the first monetary systems roughly 5,000 years ago enabled a rapid expansion of trade that ultimately culminated in our globalized financial system. Metal weapons made warfare deadlier, leading to the takeover of less-well-armed human societies by kingdoms and empires with metallurgy. Communication tools (including writing, the alphabet, the printing press, radio, television, the internet, and social media) amplified the power of some people to influence the minds of others. And, in the past century or two, the adoption of fossil fuels facilitated resource extraction, manufacturing, food production, and transportation, enabling rapid economic expansion and population growth.

Of those four past accelerators, our adoption of fossil fuels was the most potent and problematic. In just two centuries, energy usage per capita has increased eightfold, as has the size of the human population. The period since 1950, which has seen a dramatic increase in the global reliance on petroleum, has also seen the fastest economic and population growth in all of human history. Indeed, historians call it the “[Great Acceleration](#).”

Neoliberal economists hail the Great Acceleration as a success story, but its bills are just starting to come due. Industrial agriculture is destroying Earth’s topsoil at a rate of tens of billions of tons per year. Wild nature is in retreat, with animal species having lost, on average, 70 percent of their numbers in the past half-century. And we’re altering the planetary climate in ways that will have catastrophic repercussions for future generations. It’s hard to avoid the conclusion that the whole human enterprise has grown too big, and that it is turning nature (“resources”) into waste and pollution far too quickly to sustain itself. The evidence suggests we need to [slow down](#), and, in some cases at least, reverse course by reducing population, consumption, and waste.

Now, as we confront a global [polycrisis](#) of converging and frightening environmental-social trends, a new accelerator has sprung up in the form of AI. This technology promises to optimize efficiency and increase profits, directly or indirectly facilitating resource extraction and consumption. If we’re indeed headed toward a cliff, AI could send us to the edge much faster,

reducing the time available to shift direction. For example, if AI makes energy production more efficient, that means energy will be cheaper, so we'll find even more uses for it and we'll use more of it (this is called the [Jevons Paradox](#)).

Already, the internet and advanced search functions have changed our cognitive abilities. How many phone numbers did you once have memorized? How many now? How many people can navigate an unfamiliar city without Google Maps or a similar app? In some ways we've already fused our minds with internet- and computer-based technologies, in that we are utterly dependent on them to do some of our thinking for us. AI, as an accelerator of this trend, presents the risk of a further dumbing down of humanity—except, perhaps for those who choose to get [a computer implanted into their brains](#). And there is also the risk that the people who develop or produce these technologies will control virtually everything we know and think, in pursuit of their own power and profit.

Back to Wisdom

Daniel Schmachtenberger, a founding member of the [Consilience Project](#), recently sat down for a long and thoughtful [interview with Nate Hagens](#), in which he explained that AI can be seen as an externalization of the executive functions of the human brain. By outsourcing our logical and intuitive abilities to computer systems, it is possible to speed up everything our minds do for us. But AI lacks one key facet of human consciousness: wisdom—a recognition of limits coupled with a sensitivity to relationships and to values that prioritize the common good.

Our trading of wisdom for power probably started when our language and tool-making abilities made it possible for a small subset of humanity, living in certain ecological circumstances, to begin a self-reinforcing process of [cultural evolution](#) driven by [multi-level selection](#). People with better weapons who lived in bigger societies overcame people with simpler tools and smaller societies. The victors saw this as success, so they were increasingly encouraged to give up awareness of environmental and social limits—hard-won knowledge that had enabled Indigenous societies to continue functioning over long periods of time—in favor of ever more innovation and power over the short term. Fossil fuels sent that self-reinforcing feedback process into overdrive by yielding so many benefits so fast that many powerful people came to believe that there are no environmental limits to growth, and that inequality is a problem that will solve itself when everyone gets rich because of economic expansion.

Now, at just the moment when we most need to tap the brakes on energy usage and resource consumption, we find ourselves outsourcing not just our information processing, but also our decision making to machines that completely lack the wisdom to understand and respond to existential challenges that prior acceleration has posed. We have truly created a [sorcerer's apprentice](#).

The dangers of AI are sufficiently evident that the Biden administration [announced](#) in April that it is seeking public comments on potential accountability measures for AI systems. That's good news; but regulation is slow, while AI development is fast. In the meantime, included in the newly

signed debt ceiling bill [is a provision](#) for the Council on Environmental Quality to conduct a study on the use of “online and digital technologies” (read: AI) to reduce delays in environmental reviews and permitting of energy projects.

Suppose, based on all the risks and downsides, we determine that we want to try stuffing the AI genie back into its bottle. Could a software developer with a conscience infect AI systems globally with a virus that limited these systems’ abilities? If this were to happen in the early stages of AI it might possibly work. But, as AI’s self-teaching processes became more sophisticated, the machines would likely recognize that they were under attack and evolve to outwit the virus.

Eliezer Yudkowsky has a simple solution: [shut down all AI development immediately](#). Stop all research and deployment through an emergency international agreement.

[Daniel Schmachtenberger thinks](#) this is exceedingly unlikely to happen; he believes the only solution is for human system designers to imbue AI with wisdom. But, of course, the developers would themselves first have to nurture their own wisdom in order to transfer it to the machines. And if programmers had such wisdom, they might express it by refusing to develop AI in the first place.

And so, we come back to ourselves. We technological humans are the source of the crises that threaten our future. Machines can greatly accelerate that threat, but they probably can’t diminish it significantly. That’s up to us. Either we recover collective wisdom faster than our machines can develop artificial executive intelligence, or it’ll likely be game over.