CONTENTS

Introduction

"Let us calculate!"

"TO STRENGTHEN our social fabric and bring the world closer together." This, maintains Mark Zuckerberg, CEO of Facebook Corporation, is his enterprise's reason for being. Yet it would not take a particularly critical mind to notice that strengthening the social fabric and bringing the world closer together are not, in fact, what Facebook is doing. No, Facebook and the other big tech companies are, plainly, tearing the social fabric to threads, and pulling people apart.

Just fire up your computer and marvel at the news of the day, at all the angry people behind the avatars, fighting with one another and with bots about the news, and about the meaning of the news. Witness how global and local politics have been corrupted into a form of unrelenting disinformation warfare. See organized trolling campaigns fomenting violence against minority groups throughout the world. Observe the mobbing of political dissidents by mass campaigns from below, and the repression of the same dissidents by state surveillance from above. Revisit 2016 and watch the technologies of the new big tech companies mobilize to propel a disreputable internet troll into the highest office of the most powerful country in the world. Tremble before the online rage addicts who daily band

together in search of new targets: someone caught on video in a moment of indiscretion, who is then summarily doxxed (that is, has their personal information revealed on the internet), shamed, fired, or ostracized; some young adult on the cusp of success who is brought low when shown to have used hateful language as a teenager in a chat forum; some clueless normie (slang for a normal, mainstream person, oblivious to the rhythms and insiderisms of online culture) ruthlessly ridiculed for not yet having adopted the terminology for a given identity group that was ratified by social-media vanguards only a short time before. There is no sign that anyone has a clear plan, or the necessary power, to abate the chaos these technologies have unleashed. We are living in a crisis moment of history, in the true sense of "crisis": things might get better eventually, but they will never be the same.

As recently as ten or fifteen years ago, one could still sincerely hope that the internet might help "to bring people together and to strengthen the social fabric." When the revolutions of the so-called Arab Spring began to break out, many of us, myself included, declared that this was the power of social media being unleashed, hailing a new era of democracy and egalitarianism throughout the world.

The arc of such utopian hopes is long, and it has decisively bent in the direction of defeat only in the last decade. The dream of a rationally governed society, freed of passionate human conflicts through the outsourcing of decision-making procedures to machines, is one that the German philosopher Gottfried Wilhelm Leibniz already articulated as early as the 1670s. In a text in which he develops an artificial and formal language for the exact expression of all natural-language terms, the philosopher envisions a near future in which, "if controversies were to arise, there would be no more need of disputation

between two philosophers than between two calculators. For it would suffice for them to take their pencils in their hands and to sit down at the abacus, and say to each other: *Let us calculate!*"[1] The "abacus" in question is not a real abacus, but any tool that might aid in processing the formal language, though in principle Leibniz also thinks, as he conveys in this passage, that the language can be deployed using only a pen and paper (just as one might do long division either by hand or by using some sort of calculator).

This hortatory third-person-plural use of the Latin verb "to calculate"—*Calculemus!*—might well serve as the motto of Leibnizian optimism, of the belief that all problems can be resolved simply by clarifying our terms and rationally following the logical consequences of our commitments. This optimism extends not just to disputes between philosophers arguing over abstractions about the nature of substance or the immortality of the soul, but also to diplomats representing empires on the brink of war. For Leibniz, the development of a universal formal language is a key part of the imminent attainment of world peace, a part that would continue to capture imaginations in a more demotic form well into the twentieth century, where artificial languages such as Esperanto, Volapük, and Ido often appealed to peace activists of various strains, some of whom, notably Bertrand Russell (an advocate of Ido), also owed a deep philosophical debt to Leibniz.[2]

The history of artificial languages and the history of computing go hand in hand, and while the reckoning engine that Leibniz developed (which we will discuss on several occasions below) was only intended for arithmetical calculations, he well understood that in principle such a machine could also be used to process any information at all. In part this understanding was deepened by his important contributions to the development

of the binary calculus, which makes it possible to encode any proposition in a sequence of zeroes and ones, and thus to process language using the same tools with which one might also process numbers. In part, Leibniz's awareness of the possibility of concept-crunching machines, and not just of number-crunching machines, came from the fact that he was working in an already centuries-long tradition of thinking about such devices, some of which were merely fantastical, some of which may have actually existed.

Thus, in the early fourteenth century, the Majorcan polymath Ramon Llull designed a machine made of paper, consisting of several concentric discs marked with symbols on the edges denoting various attributes of substances. By rotating these discs one could, Llull hoped, exhaustively survey all of the combinatoric possibilities for the kinds of being in (and beyond) the world. Leibniz took Llull as an important predecessor in the history of formal-language processing, and Llull had his own influential predecessors too, notably Aristotle, as well as other sources in the Jewish and Islamic mystical traditions of Al-Andalus. While we might be tempted to see Leibniz, perhaps along with his contemporary Blaise Pascal, as the "father" of computer science, in truth computers have no father, or mother, and for any starting point you might attempt to choose in history, you can always find other predecessors with whom the thinker standing at that starting point was already in conversation, to whom that person was responding, who served as their starting points.

What happens with Leibniz is not the proper beginning of anything, but rather—a metaphor to which we will be returning frequently—it is the *weaving together* of several ideas into a filament thick enough to serve further on as a bright guiding thread through the rest of modern history up to the present

day: the idea that natural language can be formalized; the idea that formal language can be processed by machines; the idea that human reason can be outsourced to these machines to make decisions for us; the idea that all things are interconnected, and that therefore a change in one thing in the world is able to bring about an instantaneous change in all others, no matter what the physical distance; the expectation that we might work collectively toward creating a publicly shared compendium of all knowledge for the betterment of the lot of all humanity; the belief that knowledge is pursued and increased by individuals working within a much vaster network of other like-minded people; the conviction that collective, machine-aided labor toward the realization of reason as the governing principle of society will bring about a new era of enlightenment and lasting peace.

Although this is not a book about Leibniz specifically, he does make repeated visits, and even where he is not the subject of discussion, there is an implicit conviction that he, more than any other modern thinker, represents the spirit of the internet, the ideals that guided the first period of its development, and perhaps the best hope for its ultimate future. But the early Leibnizian spirit of the internet, as it extends, let us say, from roughly 1678 to roughly 2011, has of late fallen into existential peril. The call to "calculate" has not brought world peace. Far, far from it. Leibniz, with all due respect, was much too optimistic.

"Pessimism" about the promise of new technologies to ameliorate our condition is of course not new. To this day, no matter how careful a person is to articulate solid reasons, they still risk being called a "Luddite" in response to concerns about mechanization, recalling Ned Ludd's (likely fictional) radical resistance against the rising robotic workforce that began to emerge already at the beginning of the industrial revolution (though of

course not yet called by that name). In the early 1960s, Norbert Wiener was sharply aware that the possible apocalyptic results of modern technology might result simply from our loss of control over machines to which we have outsourced decision-making processes, and thus to teach a machine to play chess may already give it more responsibility than it can handle over war, peace, and human destiny. "There is nothing more dangerous to contemplate than World War III," Wiener writes in a supplementary chapter of the second edition of his *Cybernetics*, to which we will be returning throughout this book.[3] And, he adds: "It is worth considering whether part of the danger may not be intrinsic in the unguarded use of learning machines."[4]

A general wariness of modern technology pervades much mid-twentieth-century existential and phenomenological philosophy, frequently, as in Martin Heidegger, with discomforting undertones, and sometimes outright explicit claims, of the conflict between technological enhancement of our social lives, on the one hand, and "authentic" living on the other. This pessimism continues to echo in late twentieth-century psychological, psychoanalytic, and social-scientific engagement with the problem of modern "alienation" and the ways in which technological enhancements remove us from the human bonds and natural attachments that make life meaningful. In the 1970s, sociologists such as Manfred Stanley warned against the rise of "technicism" in interpreting human actions and motivations, and in so doing were criticized by others for their "pessimism." Yet like Stanley, and unlike Heidegger or some caricature of the "Luddite," I am interested here in "eschewing apocalyptic frenzies of doom or salvation in favor of calmer analysis."[5] While strongly opposed to the "technicist mystification of personal consciousness under conditions of modern industrial civilization"[6] and concerned to salvage "human dignity" under these

conditions,[7] I am likewise concerned to show that the greatest problem is not one of unstoppable technological determinism, or of a determinism that can only be countered by "flipping the off switch," but rather in clarifying the nature of the force with which we are contending, and understanding the limits of thinking that proceeds by analogy between human beings and machines. Stanley's approach is largely through the analysis of language, while mine is through history, but in both cases the aim is to engage in lucid criticism while avoiding the pitfalls of pessimism or authenticity-mongering.

I have been using the term "internet" in an overtly non-technical way. The internet, after all, is the entire network of networks that are connected by the Internet protocol suite. The "World Wide Web" that we commonly access through our familiar browsers is only one small part of this network. And the sites that will be of principal interest for us in the pages to follow are only one small part of what may be accessed on the World Wide Web. I am not centrally concerned, here, with the social implications of our new ability to access, say, digitized medieval manuscripts held by the Bibliothèque Nationale in Paris (though such new possibilities do become the center of attention in chapter 5), but with the more familiar sites of daily use by billions of people: Facebook, Google, and so on. Thus, "internet" serves as a sort of reverse synecdoche, the larger containing term standing for the smaller contained term. The reason for adopting this terminology is that it seems to agree with actual usage among current English speakers; on Twitter, for example, you will often see users declaring exasperatedly that their antagonists need to "get off the internet" and "touch

grass." Here, they don't really mean the whole internet; they mean Twitter.

To put this another way, I am concerned with the "phenomenological internet": the one we know directly through its appearances to us, and the one we commonly describe by that name; I date my own first use of "the internet" to a certain day in 1997, which was the first time I saw an html-based homepage, though I had been sending e-mails for five years before that and had connected to the networked computer service known as "The Source," on my father's old Kaypro, plugging our landline phone into an auxiliary suction-cup modem, as early as 1980. It seems reasonable terminologically to follow actual usage, and it seems conceptually justified to focus on the small corner of the internet that is phenomenologically most salient to human life, just as when we speak of "life on earth" we often have humans and animals foremost in mind, even though all the plant life on earth weighs over two hundred times more than all the animals combined, in terms of total biomass. Animals are a tiny sliver of life on earth, yet they are preeminently what we mean when we talk about life on earth; social media are a tiny sliver of the internet, yet they are what we mean when we speak of the internet, as they are where the life is on the internet.

Let us imagine, if we are able, a not-so-distant future in which the internet, or some suitable representative of this diffuse entity, finds itself in the dock, under prosecution for all the harms it has unleashed upon our fragile world. Let us not focus on its minor transgressions, the particular industries it has killed off or is threatening to kill off: journalism, music, film, higher education, publishing. In such cases we are only seeing what the

tech enthusiasts like to call "disruption," of the sort we see after the introduction of any new technology. And just as photography disrupted various practices, including book illustration, portraiture, and so on, without ultimately stunting or limiting our ability to represent the world around us, so too, for nearly every human practice threatened by the internet, we are already witnessing exciting and promising new practices that expand rather than shrink our potential. Newspapers for example were good in their day, but there is nothing about the electronic dissemination of news that is in principle incompatible with the social good once provided by a trusty old broadsheet.

The principal charges against the internet, deserving of our attention here, instead have to do with the ways in which it has limited our potential and our capacity for thriving, the ways in which it has distorted our nature and fettered us. Let us enumerate them.

First, the internet is addictive and is thus incompatible with our freedom, conceived as the power to cultivate meaningful lives and future-oriented projects in which our long-term, higher-order desires guide our actions, rather than our short-term, first-order desires. Second, the internet runs on algorithms, and shapes human lives algorithmically, and human lives under the pressure of algorithms are not enhanced, but rather warped and impoverished. To the extent that we are made to conform to them, we experience a curtailment of our freedom. Third, there is little or no democratic oversight regarding how social media work, even though their function in society has developed into something far more like a public utility, such as running water, than like a typical private service, such as dry cleaning. Private companies have thus moved in to take care of basic functions necessary for civil society, but without assuming any real responsibility to society. This, too, is a

diminution of the political freedom of citizens of democracy, understood as the power to contribute to decisions concerning our social life and collective well-being. What Michael Walzer said of socialism might be said of democracy too: that "what touches all should be decided by all."[8] And on this reckoning, the internet is aggressively undemocratic. Fourth, the internet is now a universal surveillance device, and for this reason as well it is incompatible with the preservation of our political freedom.

I shall have more to say about some of these indictments than others; in particular I am most interested in the first of them, the addictive power of the internet, which is one dimension of what we may call "the crisis of attention." But they all overlap in complex ways: increasingly, for example, social-media behavior in the form of likes for certain songs or artists, which might only have come to one's attention as a result of algorithmic processes over which one has no say, can also in turn place a person on the radar of law enforcement agencies or state security apparatuses as a potential terrorist, gang member, or other species of socially disadvantaged undesirable.

All of the major charges are related to one another, moreover, in contrast to the minor charges we are passing over concerning the destruction of this or that industry or art form, in that they involve, again, a threat to human freedom. Freedom is a difficult concept, in part because there are many different species of it. A Uighur in a Chinese detention camp, or a migrant in Texas with an ICE ankle monitor, is unfree, and so, in a different but somewhat related sense, is a hiker whose leg is caught under a fallen tree. A heroin addict is unfree in yet another distinct but related sense, and so are a wage laborer, a lay-about so entranced by soap operas as to never realize innate human potentials, and anyone else at all who, because of either inner weakness of the

will or objective outer forces, fails in some way to become what they could have been, fails to achieve full human thriving. We are all unfree in some of these respects. The charge here is that the internet contributes to the limitation of freedom in all of these respects. As such, the internet is anti-human. If we could put it on trial, its crime would be a crime against humanity.

Things were not always expected to turn out this way. Figuring out what went wrong will be the principal concern of this book. But in order to do this, we will need to think deeply not just about the past few years of what the internet has wrought in politics, culture, and economics. This ground has been well covered by many lucid scholars and critics. We will rather need to focus on what the internet is, ontologically speaking, on the nature of this new thing we already so easily take for granted; and we will need to focus on what the internet is genealogically speaking, too, on its place in the vast sweep of human and even natural history. Only in so doing can we begin to see what the internet might yet become.

A few words are in order concerning "methodology." This book will strike some readers as peculiar, in that it purports to be a "philosophy of the internet," yet spends most of its time dwelling on thinkers, texts, and problems from centuries ago. This is intentional; this *is* the methodology. I am, by training, a historian of philosophy and science, with a particular long-standing interest in the intersection of philosophy and the life sciences in Europe in the seventeenth and eighteenth centuries, and with an abiding interest as well in philosophical aesthetics and the many points of contact between philosophy, science, and art throughout history. I also have a strong sympathy for

some dimensions of the work of Michel Foucault, who well understood that some problems are best studied genealogically, that is, that we come to understand the essence of a thing by understanding how it develops over the course of history. This is thus in some respects a contribution to the genre of scholarship that Ian Hacking has called "historical ontology,"[9] that, namely, regards history as of central importance in any effort to understand what there is in general, or what the nature of a given thing that is, is. Thus, for example, if you want to offer up a "philosophy of cinema" (such an antiquated undertaking!), if you want to give an account of what cinema in its essence is, you must spend a good deal of time considering such things as nineteenth-century shadow plays and the narrative techniques of novelists such as Balzac or Flaubert.

Unlike Foucault, however, I am less inclined here to assent to the idea that different historical epochs are characterized by their own, radically distinct "epistēmēs." Indeed, my argument about the history of technology points much sooner in the opposite, perennialist direction: notwithstanding the enormous changes in the size, speed, and organization of the devices we use from one decade or century to the next, what these devices are, and how they shape our world, has been substantially the same throughout the course of human history (and, as we will see, even longer than that). So the book amounts to a kind of reverse Foucauldianism, or, if you will, a perennialist genealogy: bringing history to bear on a thing important enough to warrant philosophical attention, and determining through this historical-philosophical inquiry that the thing is more or less stable across the ages, and not a discursive product forever trapped within the confines of a single epoch's epistēmē, even if the current epoch does present us with some truly novel challenges.

In this short book we will range widely in topic and time, permitting ourselves to linger far from some of the questions that internet users and tech analysts today consider most pressing: the outsized power of the tech monopolies; the racism built into AI applications in security, social media, and credit-rating algorithms; the variations on the trolley problem to which self-driving vehicles give rise; the epidemic of disinformation and the corollary crisis of epistemic authority in our culture; internet mobs and the culture wars; and so on, ad nauseam. For the most part, this aloofness is intentional. This book does describe itself as a "philosophy" of the internet and, while there will be much disagreement about what that might mean, most of us can at least agree that a philosophy of something, whatever else it may be, has the right to zoom out from that thing and to consider it in relation to its precedents, or in relation to other things alongside which it exists in a totality.

But let us not suppose that zooming out can hold no practical lessons for the present day. Such an assumption is in part how we got into this whole mess in the first place. By treating the internet as a short-term problem-solver, we created for ourselves some new, very big problems; by allowing the internet to compel us to attend to a constant stream of different, trivial things, we have become unable to focus on the monolithically important thing that it is.

INDEX