To Produce and Serve: Computers, Capital, and the State

ANONYMOUS AUTHOR(S)*

The potential of computers has been subject to much speculation, sensationalism, fantasy, optimism, and concern in different parts. Yet the basic capabilities and limitations of electronic digital computing have not fundamentally changed since the establishment of the technology's theoretical foundations nearly a century ago. Seeking to put aside debates about the varied past and potential future vicissitudes of computation, we instead explore this historical continuity. We ask whether, despite an appearance of progress, any relatively fixed values have organised computation and computer science since the inception of these areas, and whose interests this continuity serves. We present an argument based on a combination of historical and critical analysis that computer science has tended to consistently prioritise and elevate orientations towards values of efficiency, productivity, and compliance. These values fit neatly in line with the interests of core contemporary and historical patrons of computer science—commercial firms and state agencies.

CCS Concepts: • **Social and professional topics** → **History of computing**; *Funding*; Computing organizations; • **Computing methodologies** → *Philosophical/theoretical foundations of artificial intelligence*; • **Human-centered computing**;

Additional Key Words and Phrases: critical data studies, history of AI, institutional analysis and design, historical materialism

ACM Reference Format:
Anonymous Author(s). 2018. To Produce and Serve: Computers, Capital, and the State. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 18 pages. https://doi.org/XXXXXXX.XXXXXXX

"Productive people.. well.. they are generally treated very well." -His Dark Materials s2e1

1 INTRODUCTION

The argument we present in brief is:

(1) Computer science can be understood both historically and today as the efficient automation of precisely defined tasks.
(2) This definition answers the positive question of what computers can do: Computers can improve productivity through gains in efficiency.
(3) The definition also answers the negative question of what computers can't do: Computers can't do anything besides exactly what they are told to do. Modern techniques of AI and machine learning have not changed this fundamental situation which is the theoretical basis for all electronic digital computing.
(4) A picture emerges of computers as efficient and compliant mechanical servants, literally "servo-mechanisms" as computers were historically called.
(5) The rise of computers can then be understood by the convergence of interests of capital and the state[1] in efficiency, productivity, and compliance.

---

[1] "The state" here is meant to represent the general institutional needs of bureaucratic forms of governance, but our examples will all be specifically of the American and British states. "Capital" refers to the individuals and firms invested in the accumulation of capital, most notably for-profit incorporated commercial entities but also private wealth and many neoliberal institutions such as the modern Anglo-American settings of higher education.

The plan for the paper roughly follows the structure of this argument. First, we present a brief review of the early history of computer science, drawing in particular on accounts by Galison [32] and primary historical sources [67, esp.]. We then establish a historical continuity by looking at Newell's telling of the establishment of computer science as an institutionalised field within American universities [60] and Wing's contemporary characterisation of "computational thinking" as the defining logic of computer science [93]. Together these authors point towards a definition of computing as the efficient automation of well-defined tasks. Taking this working definition, we then analyse what affordances computers offer and what intrinsic limitations they have from the point of view of this definition, drawing on the works of many prior critical commentators [27, 45, 87, 97, esp.]. From these analyses we distill three values at the core of historical and contemporary computer science: efficiency, productivity, and compliance. In the course of our review and analysis, we encounter ways that the shaping of the field around these values has served the interests of its key patrons, particularly commercial firms and state agencies. Our paper complements recent work re-examining the historical links between cybernetics and computer science [2, 45, 54, 77, cf.] and recent work interrogating how the institutional bedrock of computer science has related to its problems of bias and ethics [9, 11, 38, 61, 88, 92, 98, cf.]. We conclude with a discussion of subfields of computer science that have challenged its often narrow technical focus—noting that many of these exploratory areas have been led by women, people of colour, and others systematically excluded from centres of computer science, but also noting that these innovators at the margins have frequently too been serving the same commercial and state interests. We hope this paper will help practising computer scientists to understand and reflect upon the historical and social context of our work.

2 HISTORICAL CONTEXT

Although the field of computer science has resisted singular definition, we can establish a clear and unified understanding of computer science by drawing a connection between the contemporary definition of "computational thinking" [93] and the historic definition of the immediate predecessor field to computer science known as "cybernetics" [67]. There is a long history of automation and algorithmically or technologically-assisted calculations, from early time-keeping devices to the Antikythera mechanism to the Euclidean algorithm. Such machines and mechanisms have been used for millennia across countless applications in astronomy, agriculture, taxation, archiving, printing, textiles, photography, and commerce. Excitement and anxieties around automation have an equally long history, from Al-Jazari's 1206 CE Book of Knowledge of Ingenious Mechanical Devices [26] to ancient and medieval Jewish golem lore [90], revived in the modern image of Frankenstein [78].

Despite this long history, innovations in computation during the 20th century presented a sea change, bearing what we now call the Information Age. The beginning of the Information Age is generally marked by the advent of the modern focus on digital electronic computing. Through at least the 1950s there was a sustained interest and debate on the relative merits of analog versus digital computing [3, 36, 99, e.g.]. Although analog devices continue to be important in electronics more broadly, the success of the general-purpose electronic digital computer defined the vast capabilities of information technologies of the late 20th and early 21st centuries, so histories of modern computing generally begin with the first electronic digital computers and their immediate antecedents. It bears noting that "digital" in and of itself does not mean or imply "electronic", just as analog computers could also be electrically powered. As Galloway [33, p.111] poetically explains, "What is the digital? … Digital information is … constructed from the most minimally noticed amount of difference … the digital means that the one has divided into two, that something continuous has been differentiated. Thus, the digital indicates twoness… The digital means distinction or making-discrete. The digital happens wherever separation and distinction form the essential substrate of the medium. 'Two is the smallest unit of

Being,' wrote Kaja Silverman. 'It is also only through this interlocking that we ourselves exist.'" In other words, digital means based on discrete differences between 0 and 1, on and off, flip and flop, compared to analog computers which operate on or manipulate continuous, non-discrete signals. As for non-electronic digital computers, Galloway puts certain 19th century textile machines in this category, noting that Ada Lovelace had studied the punch card system used to provide these weaving machines instructions for patterns. As for the unique combination of the electronic and the digital in computing, many of the foundational developments existed in their modern forms by the time of researchers in the field that called itself "cybernetics".

Following the histories as told in particular by Galison [32] and the transcripts of the so-called Macy Conferences [67], we can identify an important initial locus of digital electronic computing and its applications in the field of cybernetics. "Cybernetic" is an interesting word. While many in critical studies and the arts have seen a renewed interest in this historical area, the term is rarely used by practising computer scientists themselves, and probably sounds dated in a similar way that "cyberspace" might. Contemporary research in the most highly regarded venues for AI research rarely use the term as a description. For instance, the word does not appear to have ever been used in the history of the proceedings of the preeminent machine learning conference Neural Information Processing Systems (NeurIPS)[2] and only once in the web index of the Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI).[3] Perhaps most notably, the related shorthand term "cyber" has become important in divisions of Anglo-American military and intelligence agencies. Of course the coincidence of these terms in these usages is no accident as we will see in terms of the massive investment for military uses from the beginning of the history of the field when computer science was still known as cybernetics. The unfamiliarity and anachronistic connotations of the term "cybernetics" are in fact also no accident, as we will also see, and belie cybernetics' relevance to computer science. The historical record shows that the popular term "artificial intelligence" was intentionally invented as a political marketing manoeuvre to control streams of funding and to erase the contributions of cyberneticists.

The original ideas of cybernetics began under the heavy influence of Arturo Rosenblueth, a Mexican researcher working for a time at the Harvard Medical School [17]. Rosenblueth was graciously acknowledged by the populariser of cybernetics, Norbert Wiener [89], but played a diminishing role in the developments of the field that in the 1940s and early 1950s came to largely centre around a cadre of electronics and computing labs at the Massachusetts Institute of Technology (MIT). Histories of cybernetics therefore for better or worse frequently begin with Wiener. Galison [32] documents Wiener's formative work as an example of three simultaneously developed and interrelated World War II wartime sciences—game theory, operations research, and cybernetics. After the turn of the 20th century, new forms of warfare had precipitated a need for new technologies. Naval power was on the wane in terms of determinative strategic importance with new kinds of rapid long-range combat as represented by air power in missiles, bombs, and airplanes developing. Submarine warfare produced research on radar. Complex physical simulations were needed to design the nuclear bomb. Fighter planes necessitated new kinds of defence systems, anti-aircraft guns. This last was where Wiener aimed to contribute. Wiener took the idea that modelling goal-oriented behaviour and feedback loops were crucial to robotic/motor control, and applied this idea to designs for anti-aircraft guns. The cybernetic anti-aircraft gun design represented an enemy pilot as an abstract entity whose behaviour could be predicted as a trajectory of goal-directed movement, what Galison calls "the ontology of the enemy." Control of the gun would be modulated by the feedback of a data stream tracking the target's observed movements. After the war, Wiener and colleagues thought to put these ideas of abstraction, feedback loops, and automatic controllers to use in broader applications.

---

[2] https://proceedings.neurips.cc/papers/search?q=cybernetics

[3] Via web search, "cybernetics site:https://aaai.org/proceeding/"

It is interesting to note that since that time, and to this day, military applications of computer science are often not mentioned explicitly even in works that disclose funding by military research agencies. Possibly this has something to do with the guilt and trauma of World War II and particularly the American use of the nuclear bomb on civilian populations in Japan. While Wiener had not himself worked on the bomb, his colleagues had, including an unrepentant John von Neumann. Wiener would jab in 1949, "the physicists have given us the ultimate of hostility and the psychologists have conditioned us to be able to use it" [66, p.29]. The cognitive dissonance between receiving military funding and justifying work outside that funding context as "basic research" was prepared just after the war in efforts to continue the high levels of scientific research funding by another MIT scientist and science policy advisor, Vannevar Bush, whose ready-made ideology separating basic from applied science helped to cleave science from society [18, 81]. While military agencies would continue to be generous patrons of American computer science research, cybernetics also attracted early interest from other parts of the state. Cybernetics was believed by some to be potentially capable of creating the *machines a gouvener*, a mechanical governor, not just of silicon but also of society. Pias explains, "Historically, cybernetics gained significance from the fact that its systems of governance and control… seemed to be highly scalable at unprecedented levels and according to entirely new standards" [66, p.21]. Or by comparison as Joque would write of modern computing [44, p.182]: "Machine learning appears to establish precisely what capitalism has always dreamed of: a smooth, universal lingua franca of epistemic and economic commensurability." Pias explains a related point drawing on Pierre Bertaux writing in 1963, "I am convinced that the future will belong to those groups of people who are first to recognize clearly that the most profitable investments they can make will be in the 'projections,' forecasts, and technical predictions…" And in 1966, Robert Theobold, "The only way to run the complex society of the second half of the twentieth century is to use the computer." Of the foundational venue of cybernetics, the Macy Conferences, the organisers write that the transactions of these meetings are recorded and published in order to be made available to among others, "government officials, administrators, etc." [67, p.340]. Pias notes after mentioning the interest of the CIA in cybernetics that "it is possible to surmise ... that issues of confidentiality were in play" [66, p.13].

We can get a clear appreciation for the ideas and intellectual climate of cybernetics from the transcripts of the Macy Conferences, edited and republished in the entirety of the conferences' last years by Pias [67]. The Macy Conferences—funded by Josiah Macy Jr, organised by Frank Fremont-Smith, and moderated by Warren McCulloch—ran from 1946 to 1953, with transcripts of the open debates in conferences made available via a stenographer beginning with the March 1949 conference. The conversations recorded at these events unveil the philosophies, ideological commitments, and personalities of those who were chosen by the organisers to be represented as the foremost cyberneticists. The conferences were exclusive invitation-only events bringing together luminary figures in the physical, biological, and social sciences who were unified by an interest in how the mathematical tools of information theory, Turing machines, and feedback control could present a basis for new insights and modelling in the sciences. By the time of these events, the Church-Turing Thesis had already been formulated, and there was a vibrant optimism—expressed to different extents and in different ways by the varied conference participants—that the new language of computers would yield profound and far-reaching understandings of complex systems. As Pias explains: "A universal theory of digital machines, a stochastic theory of the symbolic, and a non-deterministic yet teleological theory of feedback [control] were combined at the Macy Conferences into a single theory that could then claim validity for living organisms as well as machines, for economic as well as psychological processes, and for sociological as well as aesthetic phenomena" [66, p.15]. Pias goes on to explain that some participants went as far as to even believe "if all neuronal functions could be recorded as embodiments of logical calculus, it would probably have to be admitted that everything that can be known could be known in and by means of logical calculus" [66, p.15]. Indeed, the conference moderator declared [67, p.719] "...we have been very ambitious in seeking those notions which pervade all purposive behavior and all understanding of our world: I mean the mechanistic basis of teleology and the flow of information through machines and men... [the] meetings began chiefly because Norbert Wiener and his friends in mathematics, communication engineering, and physiology, had shown the applicability of he notions of inverse feedback to all problems of regulation, homeostasis, and goal-directed activity from steam engines to human societies." An illustrative passage from the transactions [67, p.345]: "In the days of the thermal engine, Maxwell developed the theory of the mechanical governor of steam engines. . . Recent complex electronic devices are not only error-controlled (like a mechanical governor), but can be so built as to seek a certain state, like goal-seeking missiles which predict the future position of a moving target (at time of impact) by extrapolation from its earlier positions during pursuit. Such devices embody electronic computing circuits, and the appearance of purpose in their behavior. . . has intrigued the theorists and prompted the construction of such likeable robots as Shannon's electronic rat. . . " This passage explains the development of ideas of control/feedback control tracing from the physicist Maxwell, who developed a self-correcting physical controllers involving a counter-weight kind of system that could prevent "error" in the operation of a steam engine, to data-driven predictive technology such as by Wiener. The latter bring technological systems beyond the capability of stable cyclical patterns such as in physical automation, towards machines that can engage in not just self-corrected, but actually self-directed behaviour towards predefined goals. The continuing influence of the military frame is evidence in a further description of the same idea [67, p.720]: "Appetitive [goal-oriented] behavior was described as inverse feedback over a loop, part of which lay within the organism, part in the environment. When a target or a goal could be indicated, a description of appetitive behavior was found to be couched in the same terms as that for self-steering torpedoes and self-training guns. . . Wiener drew a most illuminating comparison between the cerebellum and the control devices of gun turrets, modern winches, and cranes. . . "

Pias explains [66, p.23], "From its beginnings, cybernetics was less a disciplinary science than a general methodology of action." Rosenblueth, Wiener, and Bigelow made explicit the connection between action, goal-direction behaviour, computational function, function as in biological/evolutionary function, and teleology or ultimate cause [72]. It's helpful to consider the broader philosophical climate of the time to see these ideas in their historical place. Cybernetics was forming soon after the time of Nietzsche, the beginning of postmodernism. Darwin had within the past century offered the framework that for many scientists and would-be atheists became an ideology for further disempowering the Christian church in Western Protestant countries. In this fast-secularising scientific world, human "free will" was not even in question, it was becoming denied any ontological status. Evolution by natural selection provided teleological explanations in biology. The coalescing field of neuroscience could offer the potential to explain human behaviour as deterministic outcomes of neuronal activities. Computational function organised thinking about computers, especially with the focus in computing at the time being on special-purpose computers built to accomplish single specific functions. Precisely specified deterministic goal-directed behaviour was an emerging popular scientific alternative to the concept of human will. Perhaps this is why Dreyfus [27] would later write that "In Heideggerian terms ... Western Metaphysics reaches its culmination in Cybernetics." As Katz notes [45, p.188], "the field [of AI/CS] has largely adopted the epistemology of the 'analytic' philosophical tradition. This epistemology—generally associated with Descartes, fairly or not—places boundaries between mind and body, and between the external world and the subject. In this frame, subjects are primarily reflective, thinking agents: they observe and reason about an external world that they do not substantively shape." This analytical abstraction through functionalist ideology at the core of cybernetics mirrors the reductionism [76, cf.] and datafication [51, 100] that would continue to become defining features shaping the values and politics of computer science [50].

In summary, our historical review has highlighted the ways in which striving for abstract precise mathematical definition of tasks organised the thinking of cyberneticists in their early work on computers. We have also given examples of where this way of thinking was directly related to the needs of applications demanded by imperial states in the early 20th century. There were also early signs of the commercial potential. In fact, from even before their beginnings in cybernetics, computers were of a military and commercial science. Charles Babbage and Ada Lovelace's Difference Engine design was a proposal "to the British Government in 1820 as a way to make navigation tables more reliable and eliminate shipwrecks" [83] in the context of British naval and mercantile imperial power. Cybernetics saw these interests through in establishing the theoretical and practical foundations for modern computing.

3 COMPUTATIONAL THINKING

By the last meetings of the Macy Conferences there had been a sobering of some of the taller hopes of the initial years, a palpable feeling in reading the transactions and further evidenced by the departure of Wiener from the later lists of attendees. By 1952, the organisers would write much more humbly that [67, p.343] "available work is sufficient to show how communication considered from [the standpoint of the conference] can be investigated in mechanical systems, in organisms, in social groups; and the logical and mathematical problems that go into the construction of modern automata, in particular the large electronic computers, have at least partial application to our theorizing about nervous systems and social interactions." A comment by Leonard Savage further exemplifies, [67, p.379] "...it was widely advertised, that computing machines would be like people because they did so many human things… We have had, throughout the existence of this group, the important problem of seeing if there is anything that people do, that can be precisely stated, that may not be done by a machine." This lucid passage by Savage indicates the interest in computers in what they can do, and specifically what human things they can do. It is then specifically speculated that computers can do anything human can, with the proviso as above that the activity is precisely stated, i.e., stated in the form a clear and unambiguous command.

Jumping ahead some more than 50 years, we can find close parallels to the thought of cybernetics in explications of contemporary computer science. Where computers could have been thought of simply as assistive tools for calculations, cybernetics took the foundational theoretical ideas underlying computation into an entire way of thought, an ideology now called "Computational Thinking." Often credited with coining or popularising "computational thinking," former Microsoft Vice President Jeannette Wing explains simply that "Computing is the automation of abstractions" [95, p.3] and "Computational Thinking is the thought processes involved in formulating a problem and expressing its solution in a way that a computer—human or machine—can effectively carry out" [95, p.7]. Here we see the key point emphasised by Savage re-emerge: the process of formulating a task specifically enough to be made instructions for a computer is the fundamental work of the computer scientist. Wing expresses the same idea in a different way in another piece, with some further striking parallels to the agent-oriented and goal-oriented language of cybernetics: "Computational thinking is the thought processes involved in formulating problems and expressing its solution as transformations to information that an agent can effectively carry out" [94, p.1].

Reading the transactions of the Macy Conferences, it's not so surprising to see this kind of continuity decades later despite computer science being a field that seems to move so fast as to not look much past the last couple years of work in its more mathematical papers. The ideas laid out in the Macy Conferences read as still quite fresh. While the applications of computers at the time were limited by the available hardware, the paradigm of computer science in terms of what it has continued to be and the types of innovations that have been made were firmly established by the last years of the Macy Conferences in the early 1950s—from Shannon's computer analysis of natural language and mechanical rat to Ashby's chess playing software or Bavelas's networked experiments in human group problem-solving. Many new computer applications have been made since, but still solidly circumscribed within the closed paradigm laid out by cybernetics.

Artificial intelligence (AI) deserves a special note at this point not just due to its prominent place in contemporary discourse about computers, but also because its foundation helps to explain why we do not called the field of "computer science" as "cybernetics" today. As others have argued, we take the view that AI cannot be distinguished in its essence from any other form of digital electronic computing. According to the sharper and more critical frames of this point of view, AI as a term and as an imaginary has played a role as evocative gloss used to market computers and their applications [59, 91, cf.]. The historical evidence for this point is strong and warrants review, in addition to shedding light on the construction of computer science as an academic field.

As AI pioneer Allen Newell testified on the naming of the field of computer science in an unpublished manuscript archived in Carnegie Mellon's Allen Newell special collection [60, p.2]: "No notion of a scientific field shows in any of the antecedents of computer science before the 40s. Pascal, Leibnitz, Babbage, Turing all put forth their contributions as particular concerns within technology, mathematics or philosophy, as the case may be. With Weiner [sic] in 1947 we first obtained the self consciousness coupled with feelings of imminence that produced 'Cybernetics' as the name for an entire field: the new science of control and communication in the animal and the machine. It is a name and a conception cast broadly enough so that it should have served. Indeed, it currently does serve the Russians well in such a role. That it did not serve us is an historical accident, whose constituents may be examined, but whose effect seems unlikely to be reversed." Further [60, p.3], "cybernetics grew as a very broad band activity... The original spread, as in the justly famous Macy conferences on Cybernetics, had both the thrill of discovery and the feel of philosophy… Justifiably or not, Cybernetics became tinged both with a strong flavor of being an interpretation of living systems and a philosophical viewpoint rather than a solid body of scientific results… The vicissitudes of cybernetics bears extensive review because 'cybernetics' is without doubt the best name for the field."

A major point that Newell leaves out of this explanation for why the "cybernetics" did not last as the field of "computer science" was a manoeuvre of Newell's contemporary and fellow recognised AI pioneer, John McCarthy. The term "artificial intelligence" is said to have been coined at the 1956 event known as the Dartmouth Conference, organised by John McCarthy. As McCarthy himself explains it in his polemical essay collection *Defending AI Research* [52, p.73], "The proposal for the Dartmouth conference, as I remember having written it, contains no criticism of anybody's way of studying human behavior, because I didn't consider it relevant. As suggested by the term 'artificial intelligence' we weren't considering human behavior except as a clue to possible effective ways of doing tasks. The only participants who studied human behavior were Newell and Simon... whatever revolution there may have been around the time of the Dartmouth Project was to get away from studying human behavior and to consider the computer as a tool for solving certain classes of problems. Thus AI was created as a branch of computer science and not as a branch of psychology"—although McCarthy probably ahistorically places the coining of AI in 1956 as after the establishment of computer science, which according to Newell didn't firm up until the 1960s. McCarthy's recollection of AI as a subfield of computer science rather than as a subfield of cybernetics comports with a further rewriting of the history of AI as if it was invented from scratch that McCarthy details [52, p.73]: "Schopman mentions many influences of earlier work on AI pioneers. I can report that many of them didn't influence me except negatively, but in order to settle the matter of influences it would be necessary to actually ask (say) Minsky and Newell and Simon. As for myself, one of the reasons for inventing the term 'artificial intelligence' was to escape association with 'cybernetics'. Its concentration on analog feedback seemed misguided, and I wished to avoid having either to accept Norbert (not Robert) [sic] Wiener as a guru or having to argue with him… Minsky tells me that neither Wiener nor von Neumann, with whom he had personal contact, influenced him, because he didn't agree with their ideas…" This attitude of intellectual hubris is a common theme in the history of AI, as documented with brilliant perception and sharp wit by Diana Forsythe observing "knowledge engineers" of the 1990s [29]. Of course, it was to McCarthy's material benefit, as the renaming was a political manoeuvre to control funding streams [91].

Katz reviews further illuminating points [45, p.25]: "In a similar vein, Stafford Beer, a major proponent of cybernetics, complained that 'artificial intelligence,' as promulgated by Marvin Minsky and his colleagues, was 'a con trick'—a way to raise funds from the Pentagon while in practice 'simple [doing] the sort of things we [cyberneticians] were all doing anyway.' Indeed, the Pentagon has steadily sponsored projects under the label 'artificial intelligence' from the 1960s to the present… Until the early 1970s essentially all of DARPA's AI grants were given to the Massachusetts Institute of Teechnlogy and Stanford University. For a while funding was still highly concentrated: between 1970 and 1980 MIT, Stanford, and Staford Research Institute (SRI) received over 70 percent of the agency's AI funds." Katz also recounts the thinly veiled jab of the ACM [45, p.23]: "When McCarthy won the A. M. Turing Award, in 1971, the Association for Computing Machinery (ACM) noted that 'it is ironic that his most widely recognized contribution turned out to be in the field of marketing, specifically in choosing a band name for the field.'"

Returning to the commentary of Allen Newell, Newell also told of how the name for the field of computer science emerged from universities looking at how to name their new programs in this area, programs that "all had broader objectives than … strictly mathematical or engineering subfields" [60, p.5]. "Computer science" became the compromise. But Newell maintains of Computer Science departments in universities of that time, "They are all cybernetics in the original world view…" The situation has not much changed to date.

We have now established a relationship between cybernetics and the modern field of computer science. Both have focused on the efficient automation of well-defined tasks, and the scope of applications of AI and computer science has been much the same as that envisioned by cybernetics. From this continuity we can now move on to further analysis of the key values and politics that have been embedded in these trajectories.

4 WHAT COMPUTERS CAN DO

Having established historical links between cybernetics and computer science, and understood computation as the efficient automation of well-defined tasks, we now turn to examining what this definition means in practice for the field and what values it reveals as core ideological commitments of practicing computer scientists.

We have argued for understanding computers as mechanising the efficient completion of well-defined tasks. Key to analysing what computers can do then is understanding the meaning of the various components of this essential statement. What is a well-defined task? And what does it mean to complete it efficiently? Fortunately the question of what digital computers can do is a question older than computer science itself, answered precisely in the abstract general case before the field had a name. In particular, the Church-Turing thesis states that the abstract model known as the Turing machine is capable of expressing the implementation of every possible algorithmic procedure. In this way, the capabilities of a Turing machine represent everything a general-purpose digital computer can do, and indeed general-purpose computing hardware has been built with the explicit goal of reliably implementing the procedures of a Turing machine—although not ever in the same way that a Turing machine would. Computer languages are called "Turing-complete" when they can express any computation that a Turing machine could complete, and actual computer hardware is designed with this equivalence in mind.

Given the intended one-to-one mapping between Turing machines and their implementation in electronic digital computers, the question then becomes, what kinds of things can Turing machines do? Turing machines are abstract entities, which is why a formal definition can be made and can work with a great deal of precision in what it seeks to describe. The technical view is that Turing machines follow instructions to read, write, and rewrite binary (0/1, on/off, flip/flop) cells on an infinitely long serial record called a "tape", while being able to perform a collection of primitive logical operations/elementary mathematical operations with the contents of these cells. The trick of computer science is converting relevant tasks of human interest into procedures that a Turing machine can follow. The tasks for which this is possible are ones (1) in which not much is lost by representing the features and states of the world relevant to the task as a collection of numbers (i.e., a quantifiable task environment); (2) for which the computer has access to reliably controllable sensors and actuators to accurately detect and predictably affect changes in the environment relevant to the task (i.e., an observable and manipulable task environment); and (3) in which not much is lost by having a numerically measurable criterion for success at the task (i.e., a quantifiable goal).

4.0.1 Examples. Some examples may help to illustrate the various forms a precise specification of a task can take—much of the research in computer science having been dedicated to finding various formal representations of tasks, in addition to ways to complete them.

- Anti-aircraft gun [32, cf.]: Given the trajectory of a moving aircraft, predict the aircraft's next location. The task environment is quantified via Cartesian or Polar coordinates of the target's trajectory. The environment is observable via data from a physical sensor device such as the output of a radar. Success in the task is quantifiable as prediction accuracy.
- Maze-solver [67, 73, cf.]: Given the starting point, paths and walls, and end point in a maze, find a path from the start to the end point. The environment can be formally represented as a network graph with network nodes representing splits in paths of the maze and network edges representing paths without splits. Success in the task is a connected set of edges in the network graph from the start node to the end node.
- Facial recognition [69, e.g.,]: Given an image of a target face and a database of faces for comparison, identify a face in the database belonging to the same person as the target face. The environment is described by the numeric pixel values of the relevant images. Success is correct identification of a match.
- Game of Go [79, e.g.,]: Given a state of the board in the game of Go, select a move that will maximise the probability of winning.
- Crypto/blockchain [28, e.g.,]: Given a series of monetary transactions between registered parties, return the current balances each of party's accounts. Cryptocurrencies and blockchain fundamentally perform the same functions as an ordinary financial database.
- ChatGPT [64]: Given an enormous set of human-generated pieces of text and human-rated/human-annotated responses to common prompts, return a text response that is guessed to be likely to fit the pattern of highly-rated responses to prompts gauged to be similar to the given prompt.

What we have discussed so far—how tasks and information are represented to computers—what computer scientists simply call "representation," is a big part of the work in computer science. Another big part is "algorithms", i.e., the instructions that a computer follows to do something with its representations of information. It's not much help to represent a task precisely enough for a computer to undertake it if there is no efficient way for the computer to complete the task. Understanding what tasks can be completely efficiently is topic of the area known as Complexity Theory in computer science [4]. The computational complexity of a task formally describes the amount of resources needed to complete it, typically measured as the number of steps an optimal algorithm will take to find a solution, the amount of storage space needed, or the amount of energy consumed. Computer scientists describe "computational problems"—abstract groupings of precisely defined tasks—as fitting into different "complexity classes," with some problems yielding efficient algorithms for solutions, some yielding efficient approximate solutions, and others yielding no efficient solutions or approximations—being "non-computable" in the most extreme, which are tasks that can be precisely defined but for which a Turing machine would have to run for an infinite amount of time to be able to solve or which otherwise defy solution such as due to some logical paradox embedded in the task formulation.

The narrowing from all precisely formulated tasks to those amenable to efficient completion separates computer science from other branches of mathematics, logic, and analytical philosophy. The focus on efficiency and tasks that can be efficiently completed by computers also provides computer science its unique economic position as the beneficiary of huge investment from state and capital interests. As a modern example, we can look at how the influential OpenAI research organisation defined "artificial general intelligence" (AGI) for its 2019 mission statement. Two of the organisation's founders defined AGI as "automated systems that outperform humans at most economically valuable work" [14]. Here the value of economic productivity is explicitly praised.

A participant at the Macy Conference of 1951 noted [67, p.379], "Hanns Sachs wrote a paper, Why the Delay in Civilization, making the point that the machine operationally represents a magnified part of the human being, and that the building of a machine is copying the human being over a narrow range of the human capacities but multiplying the human capacities in that particular narrow range… it would be a horrible comment upon the people who build computing machines if they couldn't build one which, over a narrow range, was better than the human being over the same limited range." The entire point of computers then is that they do useful things for people, more quickly or cheaply or with less effort or using a different set of resources than otherwise possible. Wendy Chun notes the conspicuous resonances here with servitude [19, p.15]: "The history of slavery is central to the history of computing. Control systems were first called 'servo-mechanisms.' 'Master' and 'slave' functions and circuits riddle computers. This master-slave relation goes beyond computers to media more generally. Communications theorist Marshall McLuhan's framing of media as the 'extension of man' equated slaves, staples, and media: some humans were 'men' and others were their extensions." Chun's analysis emphasises the importance of the value of compliance in computing. Computers are little help when they don't do what they are told.

Having someone or something else that can do a task for you of course fits squarely into the capitalist (and racial capitalist [46, 71]) logic of productivity, and interestingly, this orientation of computing also commonly positions the programmer or computer scientist as the servile extension, the master butler. As three practising computer scientists tell it [22], "When many us were in school, we were given definitions of computer science such as 'the study of information processes and their transformations' or 'the study of phenomena arising around computers.' But when we entered the world of professional practice, we experienced computer science in a completely different way from these abstract definitions. In our professional world, our ability to obtain a job depends on how well we display competence in using computational methods and tools to solve problems of interest to our employers."

The interests of the state are not neglected either. As Katz recounts [45, p.124], "Weizenbaum offered a different metaphor for computing systems. Rather than seeing a computing system as a realization of theory, which would mean it can be 'explained' in algorithmic terms, he argued for seeing it as an intricate bureaucracy. In this bureaucracy, different subsystems, glued together somewhat haphazardly as a product of circumstance..." Or according to McQuillan [53, p.60], "Rather than heralding an alternative sci-fi future, AI can be more plausibly understood as an upgrade to the existing bureaucratic order." In this view, computerised bureaucracy represents the culmination of Weber's Iron Cage of

Rationality [25, 86]. The convergence of computing with the interests of the state is clarified by the notion of computers being aimed at efficiently completing well-defined tasks for people, and computers being completely dominated and compliant subjects.

In summary, we have seen the centrality of the value of efficiency in the very definition of computer science, the relationship between efficiency and the valuing of economic productivity, and finally the appearance of compliance as a third key value. We have also seen how these values are aligned with commercial and state interests.

5 WHAT COMPUTERS CAN'T DO

We have observed signatures of valuing efficiency, productivity, and compliance in the positive case of what computers can do. We continue our analysis of values implicit in computer science with the inverse question of what computers can't do. There are things that computers can do but which programmers generally aren't inclined to make them do, like completing tasks inefficiently or putting them to counterproductive purposes. But what kinds of things can electronic digital computers never be expected to be able to do?

One clue as to the scope of what computers can't do is in the fields represented at the Macy Conferences. The conference organisers explicitly stated their intention to include all relevant fields, and a few debates also occurred in the transcripts of the conferences as to the relevance of certain topics. From the transcripts, it is clear that many of the cyberneticists thought of psychoanalysis and hypnosis as out of scope. Also the arts were explicitly mentioned as too far afield, even though it comes up that some attendees believed arts were in principle within scope. From what topics were not represented at all and mentioned hardly or not at all, a few include history, theology, and archaeology. Some of the more embodied and particular physical sciences were also excluded, such as geology and astronomy.

It is also telling that almost none of the participants at the Macy conferences had any illusion that they would be able to imitate human life in computers, as some of the more fantastical proponents and detractors of AI believe to be possible. Conference attendees were aware of the incredible challenge of actually simulating people. Basic calculations in the Macy transactions indicate the incredible complexity of the human brain. Instead, the cyberneticists chose interim problems that would be more doable, and some of these would be of independent value to funders. This has always been the structure of computer science and AI since the time of cybernetics: A long-term imagined version, admitted as fantastical by many in the field, that other researchers figure might take years, even centuries to accomplish coupled with short-term practical applications that produce immediate or near-term value to funders. This is the Faustian bargain of the innocent and imaginative AI researcher.

The understanding we have established of computers as designed to efficiently complete well-defined tasks also underlines similar limits. Anything that has resisted formal definition or is definitionally without definition probably will not be amenable to automation in computer science or AI. This includes for example art, religion, or identity. Computers can assist in these things only insofar as we can specify the tasks for them. Games have represented many of the most lauded breakthrough demonstrations of computer "intelligence" (e.g., Checkers [84], Chess [43], Atari [55], Go [79]). Games are followed by areas for which there are already existing enormous formal ontological and semantic frameworks for: libraries/information retrieval, Jeopardy, protein folding, ChatGPT. If an application like protein folding might seem to foreshadow promise to come, consider the much praised 1960s software DENDRAL for analysing spectrograms [87] that demonstrated the great promise of computational biology. Statistical methods have been used in the sciences for centuries, and computers rely on exactly the same foundations. Similarly, computers have been able to generate natural-seeming language since the 1960s as well. Sensationalising psychoanalysts even believed that the rudimentary chat bot ELIZA would replace human therapists [87]. However, computers have always and still do struggle with the semantic content to put into the grammatical structure. The computer will never answer a question that is outside an expected format, nor will ChatGPT write a sentence unlike any it has seen.

Because computers require such neatly structured environments in order to operate effectively, the big innovations in new applications of computers have generally involved the structuration [35] of the environment to suit the application. This is comparable to the locomotive or the automobile. Locomotives would not have worked at all and automobiles would not have been nearly as transformative without the construction of railroads and motorways on which these technologies can function as intended. The proliferation of digital environments and the digitisation of life are present examples of this structuration in action. Computers work very well in digital environments and when subjects are rendered legible to them via feature engineering or classification [13, 76, cf.]. Some apparent limits of computers can also be elided through careful structuration. The famous Turing test has already been passed many times, but this has been by creating environs in which interaction is so limited and structured that computers are able to produce meaningful language due to the meaning encoded into the structure of the environment, which the computer can then navigate and optimise according to its given rules.

It must be reiterated that computers do exactly what they are told to do and only exactly what they are told. Despite much exciting and dramatic hullabaloo [16, 31, e.g.,], this does include contemporary generative AI techniques such as GPT-4 and ChatGPT, which are still just computers. Computers are told by programmers what to do in a special language designed solely for this purpose of telling computers what to do. The way a programmer communicates with a computer is unlike any form of human communication because it is exactly equivalent to the programmer manipulating certain toggles in order to have certain deterministically decided other toggles be manipulated by the computer hardware. The relationship of the programmer to the computer is not really a language in any human sense of language, then. It is more like a set of shorthand for pressing a long series of buttons on a personal calculator. There are various layers that obscure this relationship between the programmer and the hardware, but the relationship always exists because that is how computers work. In the case of so-called "generative AI", the computer does not really generate anything in the sense of the human creative capacity. The user enters a prompt, this prompt becomes data to the computer, and the computer processes that data according to its rules to return a response. This is only similar to human creativity if you believe that humans operate as computers, which has continuously been compellingly argued against in great technical detail by esteemed mathematicians and computer scientists, from Shannon at the Macy Conferences to Hilary Putnam's remarkable 1973 tour de force [68]. This is not to say that computers can't be applied to the arts, history, or even theology or religious practice—for instance, photo editing and word processing software have clearly changed work in the arts or printing—but computers cannot be historians, artists, or believers in and of themselves as we would understand these words today as people involved in the willful production of personal or social meaning.

A computer is much like any other complex technology; a computer is a tool just doing what it was built to do. It took a long and detailed investigation to understand the basic facts of the Challenger disaster, and the social meaning and causes of the disaster are still debated by scholars. But the Challenger was just a machine that was behaving according to how it was built. The Challenger has no agency in the situation. Another analogy is law [48, cf.]. Laws do not have a will in and of themselves. They are representations of human will that are used and manipulated by the will of others. Agent models can include representations called "desires" or "utility functions" but these are simply convenient ways for the programmer to present a framework for action to a computer. Computers are tools that people use to do things. So what is it computers can't do? Computers can't do anything besides exactly what they are told to do. In other words, they cannot have free will, and all this entails. They are perfectly compliant to the mathematically precise sets of instructions they are bound to follow.

The fact that computers do just what they are told does not mean they are "neutral" though, just as being "dual use" does not mean each possible use of a technology has equal moral weight that cancels each other out. Computers like all designed human creations are value-laden in their orientations towards particular uses and particular users [30, 96]. Here is revealed another side of the values of efficiency and productivity. We must ask: Whose efficiency and productivity? Especially considering the severity of global inequality in access to and usage of computing infrastructure [24, 58, e.g.,], it is clear that the fruits of these compliant servant machines fall vastly disproportionately into the bushels of the owners of the means of production [85, cf.]—the wealthy class of computer owners, and particularly high-performance computing cluster owners such as Big Tech firms and states—and these disparities fall along the typical lines that organise our "imperialist capitalist patriarchal white supremacy" as the renowned feminist theorist bell hooks labels it [8, e.g.,].

Due to the realities of social inequality, the values of productivity, efficiency, and compliance therefore reproduce the social and material conditions that protect the status of the most privileged in our societies—those who, not at all unrelatedly, also often top our state agencies, commercial firms, or both simultaneously [75, e.g.,]. Another continuity then that arises here is a value of White Supremacy, understood as the institutional structures, implicit dispositions, and explicit attitudes that benefit people who are socially constructed as White [63, cf.]—systems that notably also include the construction of the intersecting realities of gender, sexuality, class, cultural minoritisation, and disability further structuring privilege and access in many Western, colonial, and post-colonial societies. We then understand what computers can't do for whom as closely related to what society can't do, or doesn't do, for whom. As sociologist Ruha Benjamin powerfully argues [9], and Katz helpfully reprises [45], race is itself a productive technology, an artificial construct that has produced economically valuable relations of exploitation. White Supremacy—in part through commercial and state influence—has helped to produce computer science, and computers have then in turn reproduced racial categories that lay at the core of the imperial and colonial order of racial capitalist states [15].

We asked: What can't computers do? And we answered: Anything but what they are told. The question then became: So who is telling the computers what to do? If computers can only comply with their orders, who are the masters of these servant machines? (These servant machines which it must be noted are also hardly ever as automated as they seem, greatly relying instead on the deliberately obscured labour of an exploited global precariat [10, 37, e.g.,].) Computers can't serve the people for whom they don't work for. Computers incentivise the structuration of the world into one in which their masters can most efficiently extract productive outcomes from them, to either the detriment or the negligence of everyone else. Less than "automate", then, computers redistribute, and towards the already-privileged. What can't computers do? Computers can't revolutionise the social conditions that underlie the Western world order.

6 ALTERNATIVES

Having analysed the values of efficiency, productivity, and compliance in relation to both what computers can and can't do, we now turn to whether there may be alternative formulations of the study of computing that could potentially prioritise different sets of values. Since our analysis so far has also visited the fact that these values are distinctly aligned with commercial and state interests, and generally not aligned with values such as social welfare or equalities, we particularly emphasise potential realignments along these lines in our discussion of alternatives.

Prior scholars who have made similar arguments to our own generally conclude that computer science should be expanded to better include the study of people in relation to computers [27, 80, 97, e.g.,]. Such suggestions have led to interrelated efforts that have come under a variety of names, including "human-computer interaction", "human-centred computing", “human factors”, "social computing", and "public interest computing". Perhaps in part because of the possibility that these reformulations could reorient the values of computer science, these areas of study have often been led by women, people of colour, or peripheral scholars from other disciplines and, probably not unrelatedly, have regularly been sidelined relative to the more mathematical areas of computer science [47].

One of the earliest examples is Mary Shelley, author of the book popularly known as "Frankenstein" [78]. Frankenstein is perhaps one of the most famous popular culture analyses of “artificial intelligence” after the Terminator franchise, and written almost 200 years prior. Shelley, a pioneering English woman author (and as it happens, daughter of the esteemed 18th century women’s rights advocate Mary Wollstonecraft), can also be considered one of the forerunners of social commentary on modern information technology and its relationship to the human.

Another early example of interest in the human and the social within cybernetics was the inclusion of anthropologist Margaret Mead at the Macy Conference, who spoke on the complex, dynamic, and embodied elements of human language learning [67]. Mead was the only anthropologist and the only social scientist engaged in interpretive methods who was included in the years of the Macy Conferences with published transactions.

In the following decades, racialised authors, white women, and activists continued to lead in these areas. In 1963 Alice Mary Hilton wrote about the emergence of “cyberculture”, coining the term [42]. Black/African American science fiction author Samuel Delaney wrote throughout the 60s on topics ranging from gender to distributed intelligence. In 1965 the New York Times [40] published that the "computer used in a large Southern manufacturing plant to screen job applicants seemed to sparkle with objectivity" while ultimately discriminating against Black applicants. Through the late 1980s and 1990s the area of “technofeminism” grew, inspired in no small part by Donna Haraway’s 1985 essay The Cyborg Manifesto. All of this before the more-famous-within-tech 1996 “Declaration of the Independence of Cyberspace”. Alongside these cultural and speculative works was the academic research in computing that would lay the groundwork for vast swaths of what have become the established subfields of human-computer interaction and Fairness, Accountability, Transparency, and Ethics. Both the ACM Conference on Human-Computer Interaction (CHI) and ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW) were founded by women. The pioneering work of the first Black/African American woman PhD in Computer Science from MIT, Latanya Sweeney—whose work from the early 2000s was the basis for some of the first major modern AI ethics conferences within the field of computer science [82, e.g.]—was itself based on the 1980s work of another woman researcher, Dorothy Denning [21].

Despite the promise of these directions, the aspects of the work of these scholars that have managed to permeate the largest areas of computer science have often been those subject to the same commercial and state influences. Tech companies with major investments in public research such as Xerox, Yahoo, Microsoft, and Google have employed some of the biggest names in critical and social computing. Famed software engineer Grace Hopper was an admiral of the US Navy. US military research grants funded a great deal of early work on human-computer communication. Even "public interest tech" falls into the trappings of the non-profit industrial complex [62]. It must be assumed that the firms and agencies funding this work are not acting contrary to their material interests, and that at least the work of these influential scholars during their tenures under the patronage of state agencies and commercial firms is in line with the same interests as the rest of the field. In short, despite showing greater demographic diversity than better-funded areas of computer science, human-centred computing and its cognates haven’t fundamentally altered the overarching order of computer science.

Besides the subfields of computer science in the vein of human-centred computing, a number of other humanistic approaches to computing are being explored. Authors have advocated phenomenological [44] or hermeneutic approaches [97] that emphasise personal meaning-making and embodied human experience. These theoretical frameworks are well-reflected in mixed media artistic works and creative computing, which while now experiencing renewed interest also themselves have a long history [70, e.g.,]. Other threads within human-centred computing and cognate fields have aimed to self-consciously push against the dominant interests in computer science [1, 5–7, 12, 20, 23, 49, 56, 57, 65, 74]. Meanwhile, many fields have integrated computing into their own regular practices. The proliferation of mixed methods approaches in the social sciences and the growth of digital humanities are two examples. Enthusiasts in the arts and social sectors have developed innovative interdisciplinary conferences and venues that also incorporate science fiction, poetry, activism, and more. And other fields facing similar pressures from the neoliberalisation of the academy have responded with increasingly popular counternormative trends such as refusal [34], rest [41], and fugitivity [39]. However, these kind of ideas easily too become co-opted and then yield little change in what values or whose interests are ultimately being served. Katz sees this pattern as a charge to emphasise the need to not collaborate at all with state or industry partners [45]. But maybe there are no solutions, or maybe looking for a solution is a problematic starting point, or maybe thinking about problems as problems is a problem. Maybe at least it certainly wouldn't hurt to give ACM Turing Awards to Joy Buolamwini and Timni Gebru, who almost single-handedly moved conversations about racial bias and AI into both academic and public spheres.

7 CONCLUSION

In this paper we have argued that computer science can be understood as engineering the efficient automation of well-defined tasks. We have argued that this definition of the field has been historically continuous since its beginnings; that this definition and continuity reflects a deep commitment to values of efficiency, productivity, and compliance; and that the commitment to these values reflects the influence of commercial and state interests in the field. We also discussed how this convergence relates to the colonial imperial logic that organises Western society more broadly, and how many various alternative orientations for computing have been attempted—frequently by more demographically diverse groups of researchers—but have nevertheless often either faltered, dwindled, or been co-opted.

What's at stake? In the present situation, companies make donations to top computer science departments in order to define students' research work and guarantee internships. For all the libertarian rhetoric of the tech industry, academic freedom within computer science only exists within an imaginative space with borders defined by corporate and state interests. The organising question of the Academy is no longer how can we approach Truth through intellectual pursuits, but rather how can we make money through it. Productive scientific paradigms are no longer just those that produce new ideas, but those that produce profitable ideas. These profitable ideas and the technologies they yield ultimately enable certain lifestyles. But what lifestyles, and whose?

REFERENCES

[1] Syed Mustafa Ali. 2016. A brief introduction to decolonial computing. *XRDS: Crossroads, The ACM Magazine for Students* 22, 4 (2016), 16–21.

[2] Josh Andres, Alexandra Zafiroglu, Katherine Daniell, Paul Wong, Mina Henein, Xuanying Zhu, Ben Sweeting, Michael Arnold, Delia Pembrey Macnamara, and Ariella Helfgott. 2022. Cybernetic lenses for designing and living in a complex world. In *Proceedings of the 34th Australian Conference on Human-Computer Interaction*. 348–351.

[3] E. G. Andrews. 1951. A Review of the Bell Laboratories' Digital Computer Developments. In *Papers and Discussions Presented at the Dec. 10-12, 1951, Joint AIEE-IRE Computer Conference: Review of Electronic Digital Computers* (Philadelphia, Pennsylvania) *(AIEE-IRE '51)*. Association for Computing Machinery, New York, NY, USA, 101–105. https://doi.org/10.1145/1434770.1434787

[4] Sanjeev Arora and Boaz Barak. 2009. *Computational Complexity: A Modern Approach*. Cambridge University Press.

[5] Mariam Asad. 2019. Prefigurative design as a method for research justice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–18.

[6] Mariam Asad and Christopher A Le Dantec. 2015. Illegitimate civic participation: Supporting community activists on the ground. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 1694–1703.

[7] Haydn Belfield. 2020. Activism by the AI community: Analysing recent achievements and future prospects. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 15–21.

[8] bell hooks. 2004. *The Will to Change: Men, Masculinity, and Love.* Washington Square Press.

[9] Ruha Benjamin. 2019. *Race after Technology.* Polity Press.

[10] Abeba Birhane. 2020. Algorithmic colonization of Africa. *SCRIPTed* 17 (2020), 389.

[11] Abeba Birhane, Elayne Ruane, Thomas Laurent, Matthew Brown, Johnathan Flowers, Anthony Ventresque, and Chistopher Dancy. 2022. The forgotten margins of AI ethics. In *Proceedings of the 2022 Conference on Fairness, Accountability, and Transparency.*

[12] William Boag, Harini Suresh, Bianca Lepe, and Catherine D'Ignazio. 2022. Tech Worker Organizing for Power and Accountability. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 452–463.

[13] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences.* MIT Press.

[14] Greg Brockman and Ilya Sutskever. 2019. The OpenAI Mission. https://blog.gregbrockman.com/the-openai-mission

[15] Simone Browne. 2015. *Dark Matters: On the Surveillance of Blackness.* Duke University Press.

[16] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv preprint arXiv:2303.12712* (2023).

[17] Andrés Burbano and Everardo Reyes. 2022. Capsaicin and cybernetics: Mexican intellectual networks in the foundation of cybernetics. *AI & Society* 37, 3 (2022), 1013–1025.

[18] Vannevar Bush. 1945. Science the Endless Frontier: Report to the President by Vannevar Bush, Director of the Office of Scientific Research and Development. *United States Government Printing Office, Washington* (1945).

[19] Wendy Hui Kyong Chun. 2021. *Discriminating Data: Correlation, Neighborhoods, and the New Politics of Recognition.* MIT Press.

[20] Sasha Costanza-Chock. 2020. *Design Justice: Community-led Practices to Build the Worlds We Need.* MIT Press.

[21] Dorothy E. Denning. 1980. Secure Statistical Databases with Random Sample Queries. *ACM Trans. Database Syst.* 5, 3 (sep 1980), 291–315. https://doi.org/10.1145/320613.320616

[22] Peter J Denning, Matti Tedre, and Pat Yongpradit. 2017. Misconceptions about computer science. *Commun. ACM* 60, 3 (2017), 31–33.

[23] Catherine D'ignazio and Lauren F Klein. 2020. *Data Feminism.* MIT Press.

[24] Paul DiMaggio, Eszter Hargittai, Coral Celeste, and Steven Shafer. 2004. Digital inequality: From unequal access to differentiated use. *Social Inequality* (2004), 355–400.

[25] Paul J DiMaggio and Walter W Powell. 1983. The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review* (1983), 147–160.

[26] Mahmut Dirik. 2020. Al-Jazari: The Ingenious Inventor of Cybernetics and Robotics. *Journal of Soft Computing and Artificial Intelligence* 1, 1 (2020), 47–58.

[27] Hubert Dreyfus. 1979. *What Computers Can't Do: The Limits of Artificial Intelligence.* Harper Colophon Books.

[28] Quinn DuPont. 2019. *Cryptocurrencies and Blockchains.* John Wiley & Sons.

[29] Diana E Forsythe. 1993. Engineering knowledge: The construction of knowledge in artificial intelligence. *Social Studies of Science* 23, 3 (1993), 445–477.

[30] Batya Friedman and David G Hendry. 2019. *Value sensitive design: Shaping technology with moral imagination.* MIT Press.

[31] Future of Life Institute. 2023. Pause Giant AI Experiments: An Open Letter. https://futureoflife.org/open-letter/pause-giant-ai-experiments/

[32] Peter Galison. 1994. The ontology of the enemy: Norbert Wiener and the cybernetic vision. *Critical inquiry* 21, 1 (1994), 228–266.

[33] Alexander Galloway. 2021. *Uncomputable: Play and Politics in the Long Digital Age.* Verso Books.

[34] Patricia Garcia, Tonia Sutherland, Marika Cifor, Anita Say Chan, Lauren Klein, Catherine D'Ignazio, and Niloufar Salehi. 2020. No: Critical Refusal as Feminist Data Practice. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing* (Virtual Event, USA) *(CSCW '20 Companion)*. Association for Computing Machinery, New York, NY, USA, 199–202. https://doi.org/10.1145/3406865.3419014

[35] Anthony Giddens. 1986. *The Constitution of Society: Outline of the Theory of Structuration.* University of California Press.

[36] Saul Gorn. 1958. Letter to the Editor. *Commun. ACM* 1, 1 (jan 1958), 1. https://doi.org/10.1145/368685.368689

[37] Mary L Gray and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass.* Harper Business.

[38] Ben Green. 2021. The contestation of tech ethics: A sociotechnical approach to technology ethics in practice. *Journal of Social Computing* 2, 3 (2021), 209–225. https://doi.org/10.23919/JSC.2021.0018

[39] Stefano Harney and Fred Moten. 2013. *The Undercommons: Fugitive Planning and Black Study.* Minor Compositions.

[40] John Herbers. 1965. New Job Panel Tracks Down 'Biased' Computer; Finds Hiring Prejudice Was Fed to 'Objective' Device. *The New York Times* (1965). https://www.nytimes.com/1965/12/14/archives/new-job-panel-tracks-down-biased-computer-finds-hiring-prejudice.html

[41] Tricia Hersey. 2022. *Rest Is Resistance: A Manifesto.* Little, Brown Spark.

[42] Alice Mary Hilton. 1964. An ethos for the age of cyberculture. In *Proceedings of the April 21-23, 1964, Spring Joint Computer Conference.* 139–153.

[43] Feng-hsiung Hsu, Murray S Campbell, and A Joseph Hoane Jr. 1995. Deep Blue system overview. In *Proceedings of the 9th International Conference on Supercomputing*. 240–244.

[44] Justin Joque. 2022. *Revolutionary Mathematics: Artificial Intelligence, Statistics and the Logic of Capitalism.* Verso Books.

[45] Yarden Katz. 2020. *Artificial Whiteness: Politics and Ideology in Artificial Intelligence.* Columbia University Press.

[46] Robin D G Kelley. 2017. What Did Cedric Robinson Mean by Racial Capitalism? *Boston Review* (2017).

[47] Nicholas Laberge, K Hunter Wapman, Allison C Morgan, Sam Zhang, Daniel B Larremore, and Aaron Clauset. 2022. Subfield prestige and gender inequality among US computing faculty. *Commun. ACM* 65, 12 (2022), 46–55.

[48] Lawrence Lessig. 2000. Code is law. *Harvard Magazine* 1 (2000), 2000.

[49] Conor Linehan and Ben Kirman. 2014. Never mind the bollocks, i wanna be anarCHI: a manifesto for punk HCI. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*. 741–748.

[50] James W Malazita and Korryn Resetar. 2019. Infrastructures of abstraction: How computer science education produces anti-political subjects. *Digital Creativity* 30, 4 (2019), 300–312.

[51] Noortje Marres. 2017. *Digital Sociology.* Polity Press.

[52] John McCarthy. 1996. Defending AI Research: A Collection of Essays and Reviews. (1996).

[53] Dan McQuillan. 2022. *Resisting AI: An Anti-fascist Approach to Artificial Intelligence.* Policy Press.

[54] Eden Medina. 2011. *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile.* MIT Press.

[55] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.

[56] Shakir Mohamed, Marie-Therese Png, and William Isaac. 2020. Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology* 33 (2020), 659–684.

[57] Maryam Mustafa, Shaimaa Lazem, Ebtisam Alabdulqader, Kentaro Toyama, Sharifa Sultana, Samia Ibtasam, Richard Anderson, and Syed Ishtiaque Ahmed. 2020. IslamicHCI: Designing with and within Muslim Populations. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.

[58] Lisa Nakamura and Peter Chow-White. 2012. *Race after the Internet.* Routledge New York.

[59] Arvind Narayanan. 2019. How to recognize AI snake oil. *Arthur Miller Lecture on Science and Ethics* (2019).

[60] Allen Newell. 1967. Letter – "What is Computer Science?" – Newell, Perlis, and Simon – SCIENCE 157:3795 – Drafts - manuscripts for publication. https://doi.org/10.1184/pmc/newell/box00048/fld04157/bdl0002/doc0001

[61] Rodrigo Ochigame. 2019. The invention of 'ethical AI': how big tech manipulates academia to avoid regulation. *The Intercept* (2019).

[62] INCITE! Women of Color Against Violence. 2017. *The Revolution Will Not Be Funded: Beyond the Non-Profit Industrial Complex.* Duke University Press.

[63] Temu Okun. 2022. *White Supremacy Culture.* https://www.whitesupremacyculture.info/

[64] OpenAI. 2022. Introducing ChatGPT. https://openai.com/blog/chatgpt

[65] Joyojeet Pal. 2017. CHI4good or Good4CHI. In *Proceedings of the 2017 CHI conference extended abstracts on human factors in computing systems*. 709–721.

[66] Claus Pias. 2016. The Age of Cybernetics. In *Cybernetics: The Macy Conferences 1946-1953, The Complete Transactions*. diaphanes.

[67] Claus Pias (Ed.). 2016. *Cybernetics: The Macy Conferences 1946-1953, The Complete Transactions.*

[68] Hilary Putnam. 1973. Reductionism and the nature of psychology. *Cognition* (1973).

[69] Daniella Raz, Corinne Bintz, Vivian Guetler, Aaron Tam, Michael Katell, Dharma Dailey, Bernease Herman, PM Krafft, and Meg Young. 2021. Face Mis-ID: An interactive pedagogical tool demonstrating disparate accuracy rates in facial recognition. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 895–904.

[70] Jasia Reichardt. 1968. Cybernetic Serendipity. https://cyberneticserendipity.net/

[71] Cedric Robinson. 1983. *Black Marxism: The Making of the Black Radical Tradition.* Zed Books.

[72] Arturo Rosenblueth, Norbert Wiener, and Julian Bigelow. 1943. Behavior, purpose and teleology. *Philosophy of Science* 10, 1 (1943), 18–24.

[73] Stuart J. Russell and Peter Norvig. 2009. *Artificial Intelligence: A Modern Approach.*

[74] Devansh Saxena, Erhardt Graeff, Shion Guha, EunJeong Cheon, Pedro Reynolds-Cuéllar, Dawn Walker, Christoph Becker, and Kenneth R Fleischmann. 2020. Collective Organizing and Social Responsibility at CSCW. In *Companion Publication of the 2020 Conference on Computer Supported Cooperative Work and Social Computing*. 503–509.

[75] Eric Schmidt. n.d.. About Eric. https://ericschmidt.com/bio/

[76] James C Scott. 1998. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed.* Yale University Press.

[77] Aaditeshwar Seth. 2023. What's Missing in the ACM Code of Ethics and Professional Conduct. *Interactions* 30, 3 (2023), 44–47.

[78] Mary Shelley. 1818. *Frankenstein; or, The Modern Prometheus.* Lackington, Hughes, Harding, Mavor Jones.

[79] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529, 7587 (2016), 484–489.

[80] Brian Cantwell Smith. 1998. *On the Origin of Objects.* MIT Press.

[81] Donald E Stokes. 1997. *Pasteur's Quadrant: Basic Science and Technological Innovation.* Brookings Institution Press.

[82] Latanya Sweeney. 2002. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-based Systems* 10, 05 (2002), 557–570.

[83] Matti Tedre and Peter J Denning. 2016. The long quest for computational thinking. In *Proceedings of the 16th Koli Calling International Conference on Computing Education Research*. 120–129.

[84] Gerald Tesauro. 1995. Temporal difference learning and TD-Gammon. *Commun. ACM* 38, 3 (1995), 58–68.

[85] McKenzie Wark. 2004. *A Hacker Manifesto*. Harvard University Press.

[86] Max Weber. 1905. *Die protestantische Ethik und der Geist des Kapitalismus*.

[87] Joseph Weizenbaum. 1976. *Computer Power and Human Reason: From Judgment to Calculation*. WH Freeman & Co.

[88] Meredith Whittaker. 2021. The steep cost of capture. *Interactions* 28, 6 (2021), 50–55. https://doi.org/10.1145/3488666

[89] Norbert Wiener. 1948. *Cybernetics or Control and Communication in the Animal and the Machine*.

[90] Norbert Wiener. 1964. *God & Golem, Inc.: A Comment on Certain Points where Cybernetics Impinges on Religion*. MIT press.

[91] Chris Wiggins and Matthew L Jones. 2023. *How Data Happened: A History from the Age of Reason to the Age of Algorithms*. W. W. Norton Company.

[92] Oscar Williams. 2019. How Big Tech funds the debate on AI ethics. *The New Statesman* (2019).

[93] Jeannette M Wing. 2006. Computational thinking. *Commun. ACM* 49, 3 (2006), 33–35.

[94] Jeannette M Wing. 2010. Computational Thinking: What and Why? *Link Magazine* (2010). http://www.cs.cmu.edu/%7ECompThink/papers/TheLinkWing.pdf

[95] Jeannette M. Wing. 2012. Microsoft Asia Faculty Summit: Computational Thinking. https://www.microsoft.com/en-us/research/wp-content/uploads/2012/08/Jeannette_Wing.pdf

[96] Langdon Winner. 1980. Do artifacts have politics? *Daedalus* (1980), 121–136.

[97] Terry Winograd and Fernando Flores. 1986. *Understanding Computers and Cognition: A New Foundation for Design*. Ablex Publishing Corporation.

[98] Meg Young, Michael Katell, and PM Krafft. 2022. Confronting Power and Corporate Capture at the FAccT Conference. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 1375–1386.

[99] Serge J. Zaroodny and Tadeusz Leser. 1958. AYDAR, Special Purpose Analog Machine for Raw Data Reduction. *J. ACM* 5, 1 (jan 1958), 89–99. https://doi.org/10.1145/320911.320922

[100] Shoshana Zuboff. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. Public Affairs.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009