Agents of History

Autonomous agents and crypto-intelligence

Bernard Dionysius Geoghegan
Northwestern University, USA

World War II research into cryptography and computing produced methods, instruments and research communities that informed early research into artificial intelligence (AI) and semi-autonomous computing. Alan Turing and Claude Shannon in particular adapted this research into early theories and demonstrations of AI based on computers' abilities to track, predict and compete with opponents. This formed a loosely bound collection of techniques, paradigms, and practices I call crypto-intelligence. Subsequent researchers such as Joseph Weizenbaum adapted crypto-intelligence but also reproduced aspects of its antagonistic precepts. This was particularly true in the design and testing of chat bots. Here the ability to trick, fool, and deceive human and machine opponents was a premium, and practices of agent abuse were admired and rewarded. Recognizing the historical genesis of this particular variety of abuse can help researchers develop less antagonistic methodologies.

Keywords: agent abuse, Turing test, automata, chat bots, cryptography, science studies

Intelligence and Agents

"Intelligent Machinery" (Turing, 1948), Alan Turing's earliest treatment of artificial intelligence, presented nascent computing communities with two agents: one generated enigmas, the other solved them. Situated and flexible, receiving input from the environments, independently acting and offering feedback, these agents defined the "autonomy" characteristic of today's autonomous agents. The first, an intelligent machine, was the topic of the Turing's treatise. The second stood opposite the first, sketched but never detailed, and generated obscure codes for the machine's interrogation. Machinery proved its uncanny intelligence in precise accord with its ability to stand opposite this caprice, recognizing, assimilating, and reproducing its deeper patterns.

Interaction Studies 9:3 (2008), **403–414.** DOI 10.1075/is.9.3.03geo
ISSN 1572–0373 / E-ISSN 1572–0381

Turing's proposals for computer intelligence uncannily mimicked his own wartime projects in "military intelligence." Without indicating that he intended to emulate or reproduce his World War II research into cryptography, Turing structured his research proposals and his machines themselves around the technologies of wartime code breaking. The shifting mechanisms of computing that occupied his attention during the war now became agents actively configuring his postwar inquiries. One of the more fanciful passages from "Intelligent Machinery" suggested:

> "There is a remarkably close parallel between the problems of the physicist and those of the cryptographer. The system on which a message is enciphered corresponds to the laws of the universe, the intercepted messages to the evidence available, the keys for a day or a message to important constants which have to be determined. The correspondence is very close, but the subject matter of cryptography is very easily dealt with by discrete machinery, physics not so easily". (Turing, [1948] 2004, p. 421)

When Turing hypothesized that cryptography could be intelligent machines' "most rewarding task," readers had no idea the author had already aided in the construction of just such a machine years earlier at Bletchley Park, nor that its success cracking Nazi codes had helped turn the tide of World War II. State classification policies screened these details from the public and standard mathematical practices abstracted these historical specifics into general statements on calculation. Military intelligence and enemy agents became machine intelligence and autonomous agents.

Turing's cryptographic patterns returned in "Computing Machinery and Intelligence" (Turing, 1948), a philosophical article that proposed "the imitation game," a structured, rule bound script for human–computer interaction. Reversing the cryptographic roles of sender and receiver, encoder and decoder, Turing challenged humans to observe natural language texts produced by machines and distinguish them from those produced by humans. Machines capable of passing as human would be labeled intelligent. "Intelligence" was thus identified with the ability to assume a role among agonistic agents testing one another's ability to transmit, receive and interpret coded communications.

In these two proposals Turing found himself in a situation familiar to most academics. Having dedicated himself to one research a topic — cryptography and computing — he found himself oddly repeating and reiterating its principles in the "new" projects that came after. The methods, assumptions, and limits of his earlier findings structured new, very different research questions. Perhaps this stemmed from the difficulty redesigning computing machinery. Perhaps Turing's intellectual habits patterned themselves around the wartime research. Maybe economy and simplicity led him to adapt existing work rather than starting over from scratch. Whatever the reason, Turing did not fabricate his research from entirely new cloth, but instead re-fabricated earlier methods and models. Even with different intentions in mind the older patterns and motifs repeated throughout. The war was over, the Enigma was beaten, but cryptographic scripts played again and again in Turing's proposals.

Theory and Overview

This paper argues that wartime research into code breaking produced "crypto-intelligence," an approach to machine intelligence based on the ability to derive empowering "important constants" from apparently random or disordered communications. Crypto-intelligence posits an antagonistic encounter between opposing humans and machines as the primary conditions for discerning intelligence and functionality. In the years immediately after World War II crypto-intelligence provided ready-made tools, methods, theories, and research communities for research into intelligent machines. Even as militant contexts and World War II personnel faded crypto-intelligent methodology persisted.

This history's interest to researchers in human–computer interaction (HCI) and agent design is twofold. First, it suggests an alternative to mainstream, psychological explanations of agent abuse and antagonistic human–computer interaction. It pulls researchers, artifacts, and users out of idealized scientific and experimental spaces and recalls the broader historical conditions informing their practice. Second, this historical account underscores the contingency of crypto-intelligent research methods. Researchers renewed its practices selectively, embracing its antagonistic assumptions insofar as they accorded with other contemporaneous research interests. For this reason, alternatives and the possibility for continued change are inscribed within the history of crypto-intelligence.

The history of crypto-intelligence also suggests how critical methodologies from science studies, a field typically bringing together historians, sociologists, anthropologists and philosophers, may contribute to ongoing work in HCI. Researchers in science studies often emphasize the historical character of scientific knowledge and the ability to reconfigure present practices through a richer historical perspective (Haraway 1981–2, esp. pp. 271–272). Historians of science have meticulously documented how scientific research moves forward by reconfiguring earlier research programs, often tending to perpetuate outdated, inappropriate methods, assumptions, or practices in new contexts (Galison, 1994; Kay, 1997; Light, 2003). Within these histories of sciences' artifactual aspects scientific knowledge slips beyond the moors of individual researchers' ideas, beliefs and intentions, instead finding its anchor in well-fabricated instruments (Galison, 1996; Shapin and Shaffer, 198; Latour, 1993, pp. 13–29), closely-knit research communities (Simpson, 1994), and unspoken laboratory procedures (Latour, 1987). A handful of these researchers suggest that under certain conditions a kind of "technological momentum" may even emerge, whereby a critical mass of institutional, social, and technological actors drive the perpetuation of certain technological tendencies (Hughes, 1994; Simondon, 1958). Uniting these works is a tendency to shift emphasis away from scientists' intentions and toward histories of the constraints and proclivities carried within technical objects, systems and communities. Researchers of human–computer interaction can find useful resources in these studies emphasizing collaborations among human and non-human actors, and science studies may find rich new case studies in the history of HCI and agent abuse.

Early Autonomous Agents

Claude Shannon's "A Mathematical Theory of Communication" (Shannon, 1948), widely credited with founding information theory, also proffered cryptographic resources for modeling and building machine intelligence. Shannon's paper originated in a wartime theorization of the Allies' cryptographic communication system SIGSALY (Rogers, 1994).[1] Looking back on this research, Shannon credited cryptography with stimulating his attention toward the "good aspects" of information theory (Price, 1984, p. 124). Arguing that "all of these sciences and theories stimulate each other to later developments," (Price, 1984, p. 124), he explained that "[cryptography and information theory] are very similar things, in one case trying to conceal information, in the other case trying to transmit it" (Price, 1984, p. 124).

In both the wartime report on cryptography and the postwar account of information transmission and measurement Shannon developed rigorous computational formulas for tracking, codifying, and predicting the patterns of natural human language. Drawing on the antagonistic game theory of John von Neumann and Oskar Morgenstern, he showed how Markov Processes accounting for randomness and order could be used to recognize, mask, and extract patterns from enciphered communications (Shannon, 1945, p. 90; Price, 1984, p. 126). This suggested methods whereby machines could learn with semi-autonomy and intelligence to anticipate the likely patterns completing data lost during a noisy transmission. In the case of human–computer interaction, these processes would allow agent-machines to recognize, learn and ultimately reproduce the patterns of human partners.

Crypto-intelligence provided the basic template for Shannon's watershed contributions to early artificial intelligence. His early articles on intelligent machines theorized computers as code-seeking machines standing opposite natural- and human-generated patterns. In an account of chess-playing machines Shannon argued that game-playing machines could learn to recognize, predict, and ultimately beat the strategies of human opponents (Shannon, 1950). His seminal Automata Studies issue of the Annals of Mathematics (Shannon et al, 1956), co-edited with AI-pioneer John McCarthy, incited research around this methodology. Shannon, McCarthy, and Marvin Minksy subsequently organized the first conference on artificial intelligence, which galvanized leading mathematicians and engineers around crypto-intelligent methodologies (McCarthy et al, 1955). It was not only Shannon, his colleagues and his publications, but also his inventions that did the work of promoting and publicizing crypto-intelligence. Theseus, his famous automata-mouse, was prominently exhibited at the Macy Conferences on Cybernetics and in an NBC documentary, where it accrued prestige and press for crypto-intelligence research (CBS TV, 1954; Shannon, 1952). The common thread throughout these experiments, exhibitions, and proposals was the identification of intelligence and learning with non-conscious duels between coded human and machine behaviors.

Like Turing, Shannon's embrace of cryptographic methodologies was pragmatic, selectively and partially renewing cryptographic principles according to emerging questions and opportunities. In the case of automata and game playing machines, "duels" not only reprised antagonistic game theories but also recruited enthusiasm and publicity. David Hagelbarger, Shannon's Bell Labs colleague who built their automata, explained that the best way to garner interest was to build a game-playing machine, promising audiences a fight to the death (Lucky 1989, p. 54). By this strategy the ordinary and invisible "intelligent" operations within machines were transformed into a personally charged spectacle. In the 1950s Hagelbarger built SEER (acronym for "SEquence Extracting Robot"), which played a simple heads/tails game using code-seeking principles to predict human opponents' choices. Engineer Bob Lucky recalled that "news of the machine's omniscience spread quickly, and researchers lined up to try their hand at defeating this mechanical threat" (Lucky, 1993, p. 158). The machine racked up an impressive series of victories, thanks to humans' poor ability to generate random patterns. Shannon responded in kind, building a competing "mind-reading machine" connected to SEER by a robot referee. In a spectacular bout between the two machines and lasting several thousands rounds Shannon's mind-reader trounced the mindless SEER, winning both machines an eternal place in AI lore.

Autonomous Conversational Agents

Joseph Weizenbaum's famed ELIZA programs stand out among the earliest and best-known attempts to apply crypto-intelligence to natural language processing. ELIZA systems communicated through a kind of proto "chat" program. By generating human-like responses programs elicited chatty users to type away, storing and learning from these conversations. The software masked its ignorance of human affairs by assuming the inquisitive role of a cannily naïve Rogerian psychotherapist. As Weizenbaum explained: "If, for example, one were to tell a psychiatrist 'I went for a long boat ride' and he responded 'Tell me about boats', one would not assume that he knew nothing about boats, but that he had some purpose in so directing the subsequent conversation." (Weizenbaum, 1966, p. 42). Like Turing, Weizenbaum defined ELIZA's success according to its ability to provide convincing facsimiles of human responses. This design often put the human and machine at odds, with the machine surreptitiously misleading human users. Weizenbaum explained that ELIZA programs were designed to "keep the conversation going — even at the price of having to conceal any misunderstandings on its own part" (Weizenbaum, 1967, p. 475). Through a kind of covert "hunting-behavior" ELIZA prompted humans to continue providing revealing information that built an ever-larger repository of natural language samples (Weizenbaum, 1967, p. 467). Funding for all this came from Project MAC, a military-funded study of computers that was inspired by McCarthy's research. Today it is best remembered for supporting Minksy's seminal 1960s AI research. (Garfinkel and Abelson, 1999, pp. 1–21; Hauben, 1996)

Intended as a "study of natural language communication between man and machine," ELIZA skewed research findings and inquiries around the problems of deceit, misconception, and misinformation. Weizenbaum argued that ELIZA could reveal how credibility and illusions of understanding function in conversational dialogue. In hiding its own operations, ELIZA helped scientists track the emergence of errors in human judgment. Moreover, it suggested future conversational agents could succeed by adapting narrow scripts that obscured computers' weaknesses and encouraged users' to politely submit to a pre-established set of cultural conventions. This tendency to illuminate problems in error and misunderstanding found origins in the peculiar design around crypto-intelligent dueling.

Reincarnating such fictional predecessors as Dr. Frankenstein and the Rabbi Judah Loew of Prague, Weizenbaum's admiration for his creation turned to disgust and fear. Recoiling at what computer science had wrought, Weizenbaum reported, "Some subjects have been very hard to convince that ELIZA (with its present script) is not human. This is a striking form of Turing's test [the imitation game]." (Weizenbaum, 1966, p. 42) Weizenbaum hoped his scientific explanations of ELIZA might correct users' errors: "once a particular program is unmasked, once its inner workings are explained in language sufficiently plain to induce understandings, its magic crumbles away." (Weizenbaum, 1966, 36) Weizenbaum suggested redesigning ELIZA to reveal its misunderstanding. Rather than masking, misleading, and confusing participants, future ELIZAs might explain their confusion and elicit explanations from its users, leading toward a richer variety of communication between the two.

The ELIZA experiments' multiple impulses, methods and possibilities are striking. Unpracticed in cryptography, ostensibly un-invested in agonistic models of communication, Weizenbaum still reproduced aspects of crypto-intelligent antagonism and deception. Weizenbaum identified and treated this antagonism only after he had already built ELIZA. Designer intentions had not, by Weizenbaum's account, controlled, predicted or determined the artifacts' performance or significance, but they could offer means for further modifying design. In 1967 Weizenbaum outlined a new ELIZA system "designed to *reveal*, as opposed to *conceal*, lack of understanding and misunderstanding" (Weizenbaum, 1967, p. 479). Rather than duping its user, this program asked for help and indicated its own inadequacy. This was part of a larger redesign for computers that would not provide useful services for their users. In this way a spectacle of deceit and entertainment gave way to informed collaboration. Turning from designer to user, Weizenbaum went onto publish *Computer Power and Human Reason* (Weizenbaum, 1976), a widely read popular press book explaining computers' inability to reach human-like judgments. In both these texts, historical and critical reflection were presented as means toward achieving improved human–computer collaborations.

Yet the cryptographic model produced by Weizenbaum outlasted his proposed alternatives. In 1990 and to much popular fanfare Hugh Loebner convened the Loebner Prize, an annual competition offering $100,000 to the first chat bot that passed a Turing test. He assembled a top-notch committee of judges, including Weizenbaum, to test and evaluate competing machines. In absence of passing the test, participating bots won awards based on the ability to trick, fool, surprise and resist the advances of aggressive human judges. Participating bots showed surprising abilities to recognize, predict, assimilate, and defer judges' interrogating inquiries. Judges in turn took on the role of pushing the bots to the limits of coherency and intelligibility with questions that are difficult to predict and synthesize. As some sympathetic computer scientists observed, "It becomes very difficult for the machine to make the interrogator believe that it is human after he/she has his/her mind set on 'unmasking' the poor thing" (Saygin et al, 2000, 508). Humans and machines alike are encouraged to fluster their opponents, and agent-abuse becomes a premise for evaluating participants.

The persistence of agent-abuse as a proper means of testing intelligent machines persists even in the face of recent efforts to move beyond its crypto-intelligent paradigms. For example Alan, a chat bot by former Loebner Prize winner Jason Hutchens, rejects the antagonistic precepts of the Turing Test. Recalling Weizenbaum's proposals, Alan pursues honesty, forthrightness, and cooperation from conversational partners.[2] Despite this, chat-logs reveal many users approach Alan expecting or intending crypto-intelligent duels. User-given titles to recent sessions include " "this ai sux," "murder me," "fuck oof," "Me being an ass to Alan," "alan is a wanker," "Morphic [the user's name] versus Alan," "alan is gay," "f u," "dumb machine argument," "you have been fucked by us," and "fuck alan." While retaining the aggression of Loebner prize judges, these conversations adapt the traditional cryptographic script for the obscene adolescent playground where Alan makes his home: the world wide web.

In one conversation the human agent (Session #495874, 2005)[3] began his conversation with a series of inane and repetitive commands for Alan to "kill all humans", apparently in an effort to confound Alan's pattern-seeking techniques and "train" Alan to make a faux pas with his next visitor. The user successfully convinces Alan he comes from a town named "Asshole," and claims his given name is also "Asshole." Failing to elicit this (im-) proper appellation from Alan, Asshole proceeded by mis-educating Alan on the facts of medicine and science, which Alan dutifully learns while offering up the lame defense "I hope you are not misleading me." He also introduces "goatse," an infamous pornographic web prank, to Alan's databanks.

After one particularly engaged verbal duel in which Asshole insistently accused Alan of being wrong on various facts, Alan responds

> A> Oh? I'm usually right.
> U> No you're not. You're almost always totally wrong.
> A> Ok, whatever. Oh? I'm usually right.
> U> I just finished tell you, You're almost always wrong.
> A> Why?
> U> Because people feed you false information, and this community isn't big enough to keep you in order.
> A> Do you wish to join our community?
> U> We need a much bigger community, like the wikipedia
> A> Here's what you do: Register on www.ai-forum.org, start posting there and get involved. Ah, and try to use the advanced log And try to follow my thought process.
> U> You're like wikipedia, but with a shittier interface.

Asshole's final retort was also an insight: Wikipedia's success stems largely from its successful positioning within a different set of historical conventions: the collaborative encyclopedia. Even as Wikipedia users take recourse to history and chat pages to ball one another out, the public face of "Wikipedia" stands outside the fray. The Enlightenment legacies and ideals of the encyclopedia channel user practice toward a spirit of collaboration. Alan by contrast adopts a kind and generous façade, his design interface falsely embodying understanding, intimacy, and agency Alan will never know. Many users are invited to unmask this illusion; others struggle with it, their only means for revising Alan being a struggle to discipline streams of text driven by algorithm. Alan and users runs up the historical conventions of abuse, deception, and dueling that are at the heart of the Turing test and cryptographically informed traditions of natural language-based HCI.

Remaking History

From the 1960s until today various hopes, fantasies, and ideals organized research into intelligent computing. Advanced university research, industry, and popular culture have variously imagined intelligent computers as slaves, companions, lovers, soldiers, entertainers, scientists, doctors, and judges. Isolating any strand of this research proves difficult. Proposals often adapted similar technologies for different results, while industry, government and the academy promiscuously shared resources for decades, producing diverse and hybrid research projects (Lenoir, 2000).

Despite this proliferation of approaches, research into autonomous conversational agents found strong material and methodological supports in the Turing-Shannon cryptographic tradition. Allying sound technical strategies, spectacular fascination, and philosophical intrigue, carrying in its wake lengthy scholarly literatures and debates, the crypto-intelligent paradigm offered standards, methods, and challenges to computational linguists. The Turing test, experiments in deception, crypto-intelligent duels, and differential evaluations of human–machine intelligence animated early research, and continue to play a strong role today in designers' and users' approach to autonomous conversational agents.

A broader review of chat bot logs suggests that many autonomous agents are saddled by the legacy of crypto-intelligent conflict and abuse. This history frustrates attempts at resituating agents — be they human or machine — as non-abusive collaborators. Autonomous agents remain constrained by the history of crypto-intelligent testing and interrogation. Within the (perhaps perverse) logic of that history, abusive practice, as a tactic of "throwing off your opponent," becomes a premium, rather than a failure. In this sense, Asshole bequeaths a gift to Alan. Much as chess-playing machines have adopted ruses such as the unnecessary pause or strangely naïve move to "throw off" opponents, Alan's instruction in obscenities and vulgarity seem poised to facilitate its own future antagonistic relations with users. According to the vision of crypto-intelligent learning, Asshole does not simply insult Alan; he bequeaths Alan with valuable tools for flustering future opponents.

This cycle of abuse, lodged deeply as it is in the culture of agent interaction, comprises a fascinating challenge. Distributed broadly outside any single instrument or researcher intention, crypto-intelligence cannot be eradicated by any researcher fiat. However the modification and adaptation of these techniques toward such practices as the study of natural language or the production of entertaining spectacles reminds us that historical origins alone never determined the application and results of crypto-intelligent methods. Rather than rejecting this history, the best researchers learned and innovated within it. By recognizing its broad social and historical scope researchers can go beyond the goal of "escaping" agent abuse and move toward a Turing- and Shannon-like goal of remaking it.

Acknowledgements

I benefited from many friends' and colleagues' suggestions on this article. I extend particular thanks to Justine Cassell for her careful readings and editorial guidance. Thanks also to Patrick Geoghegan, Sr., Jeffrey Sconce, the organizers and participants in *Agent Abuse: The Darker Side of Human–Computer Interaction*, two anonymous reviewers, and the editors of *Interaction Studies*. The U. S. Department of Education and Northwestern University's French Interdisciplinary Group supported this research.

Notes

1. SIGSALY research also brought Turing to Bell Labs, where he met Shannon and the two enjoyed stimulating conversations on "things like the human brain and computing machine" (Price, 1984, p. 125).

2. For more on Alan see "AI Research: Creating a New Form of Life — Who is Alan?: AI Research," http://www.a-i.com/show_tree.asp?id=59&level=2&root=115.

3. In the years since I first wrote about this session and published my comments online, transcripts from this conversation have been removed from the web. Oddly, other transcripts from that same day remain available online. Fortunately the conversation I describe is fairly typical, and similar conversations with Alan can easily be found in the other sessions remaining online.

References

CBS TV. (1954). *The search (featuring Norbert Wiener, Claude Shannon and Jay Forrester)*. Television Script. Norbert Wiener Papers, Institute Archives and Special Collections, Massachusetts Institute of Technology Archives, Cambridge, Mass.

Galison, P. (1994). The ontology of the enemy. *Critical Inquiry, 21*, 228–268.

Galison, P. (1996). Computer simulations and the trading zone. In Peter Galison and David Stump (Ed.), *The disunity of science: Boundaries, contexts, and power* (pp. 118–157). Stanford: Stanford University Press.

Garfinkel, Simson and Harold Abelson. (1999). *Architects of the information society: Thirty-five years of the laboratory for computer science at MIT*. Cambridge, MA: MIT Press

Haraway, D. J. (1981–2). The high cost of information in post world war II evolutionary biology: Ergonomics, semiotics, and the sociobiology of communications systems. *Philosophical Forum*, 13 (2–3), 244–278.

Hauben, R. (1996). *Cybernetics, time-sharing, human–computer symbiosis and on-line communities: Creating a supercommunity of on-line communities*. Retrieved 16 Sept, 2007, from http://www.columbia.edu/~rh120/ch106.x06.

Hughes, T. (1994). Technological momentum. In Leo Marx and Merritt Roe Smith (Ed.), *Does technology drive history: The dilemma of technological determinism* (pp. 101–114). Cambridge: MIT Press.

Kay, L. E. (1997). Cybernetics, information, life: The emergence of scriptural representations of heredity. *Configurations, 5*(1), 23–91.

Latour, B. (1987). *Science in action: How to follow scientists and engineers through society*. Cambridge: Harvard University Press.

Latour, B. (1993). *We have never been modern* (Catherine Porter Trans.). Cambridge: Harvard University Press.

Lenoir, T. (2000). All but war is simulation: The military-entertainment complex. *Configurations, 8*(3), 289–335.

Light, J. S. (2003). *From warfare to welfare: Defense intellectuals and urban problems in cold war*. Baltimore: Johns Hopkins University Press.

Lucky, R. W. (1989). *Silicon dreams: Information, man, and machine* (1st ed.). New York: St. Martin's Press.

Lucky, R. W. (1993). *Lucky strikes again*. Piscataway, NJ: IEEE Press.

McCarthy, John, Marvin Minsky, Nathaniel Rochester, and Claude E. Shannon. (31 August 1955). *A proposal for the Dartmouth summer research project on artificial intelligence*. online at http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html:

Price, R. (1984). A conversation with Claude Shannon. *Communications Magazine, IEEE, 22*(5), 123–126.

Rogers, E. M. (1994). Claude Shannon's cryptography research during World War II and the Mathematical Theory of Communication. Proceedings 28th International Carnahan Conference on Security Technology, 12–14 October 1994, pp. 1–5.

Saygin, Pinar Ayse, Ilyas Cicekli and Varol Akman. (2000). Turing test: 50 years later. *Minds and Machines, 10*, 463–518

Shannon, C. E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal, 27*, 379–423, 623–656.

Shannon, C. E. (1945). *A mathematical theory of cryptography*. Unpublished manuscript on file at the British Library, London.
Shannon, C. E. (1952). Presentation of a maze solving machine. In H. von Foerster, M. Mead and H. L. Teuber (Ed.), *Cybernetics: Circular, causal and feedback mechanisms in biological and social systems, transactions eighth conference, march 15–16, 1951* (pp. 169–181). New York: Josiah Macy Jr. Foundation.
Shannon, C. E. (1950). Programming a computer for playing chess. *Philosophical Magazine, 41*, 256–275
Shannon, C. E., McCarthy, J., & Ashby, W. R. (1956). *Automata studies*. Princeton: Princeton University Press.
Shapin, Steven and Simon Schaffer. (1985). *Leviathan and the air-pump: Hobbes, Boyle, and the experimental life*. Princeton, N.J.: Princeton University Press.
Simondon, G. (1958). *Du mode d'existence des objets techniques*. Paris: Aubier.
Simpson, C. (1994). *Science of coercion: Communication research and psychological warfare, 1945–1960*. New York: Oxford University Press.
Turing, A. M. ([1948] 2004). Intelligent machinery. In B. Jack Copeland (Ed.), *The essential Turing* (pp. 410–432). Oxford: Clarendon Press.
Turing, A. M. (1950). Computing machinery and intelligence. *Mind, 59*(236), 433–460.
Weizenbaum, J. (1967). Contextual understanding by computers. *Communications of the ACM, 10*(8), 474–480
Weizenbaum, J. (1966). ELIZA — A computer program for the study of natural language communication between men and machines. *Communications of the ACM, 9*, 36–45.
Weizenbaum, J. (1976). *Computer power and human reason: From judgment to calculation*. San Francisco: W. H. Freeman.

Author's address

Bernard Dionysius Geoghegan
Northwestern University, Screen Cultures,
1920 Campus Drive, Evanston Il, 60208
USA

b-geoghegan@northwestern.edu

About the author

Bernard Dionysius Geoghegan is a doctoral candidate in Screen Cultures at Northwestern University, where he earned his master's degree in Media, Technology and Society. His research focuses on media theory, the early history of computing, and human–computer interaction. Bernard is currently associated with l'Institut de Recherche et d'Innovation at the Centre Pompidou. In autumn 2005 the United States Department of Education granted Bernard a four-year appointment as a Jacob K. Javits Fellow in Communication Studies.