TALK WITH MACHINES, REDUX

By Lucy Suchman

"The designer's project in this sense is to imbue the machine with grounds for behaving in ways that are accountably rational; that is, reasonable or intelligible to others including, in the case of interaction, ways that are responsive to the others' actions."

Suggested citation:
Lucy Suchman, Talk with Machines, Redux. *Interface Critique* 3 (2021), pp. 69–80.
DOI: 10.11588/ic.2021.3.81328.

Preface

With Claude Shannon's translation of numerical analysis into "The Mathematical Theory of Communication" in 1949[1], the foundations were secured for an imaginary of the computational sensorium. Framed through the trope of communication, moreover, these foundations were posited as the model not only for signal processing but also for human cognition and social relations. In the ensuing decades the intelligent, interactive machine has become an established figure in discussions of information technology, from scientific and professional discourse to popular media representations. At the same time, scholarship in feminist science and technology studies, post/decolonial theory, media studies, and related fields has taught us to question the figure, to trace its genealogies and identify its legacies. Figuration, we now understand, is consequential not only for what is posited to be known but also for what comes to matter, as well as what is ignored and what differences that matter are erased.[2]

Since the 1980s I have been following the figure of the intelligent, interactive machine with the aim of slowing down its facile claims, not only for kinship with (more than) human being but for pride of place among sensing media. In the beginning my impetus to resist was based in an immersion in ethnomethodology and conversation analysis, through the work of a community of researchers profoundly attentive to the exquisite, mundane improvisations through which human interlocutors reflexively co-construct their mutual intelligibility.[3] In light of the loving care with which these scholars traced the emerging, contingent arc of socially organised, materially embodied human action and communication, demonstrating the astonishing range and depth of competencies involved in even the most ordinary encounter, the trope of signal processing seemed to offer a weak analogy. But it was only gradually that I came to appreciate the strength of its rhetorical grasp.

The work revived in this contribution began its life over three decades ago, in the context of a PhD dissertation submitted to the Department of Anthropology at the University of California at Berkeley in 1984. It was an unusual dissertation for the discipline of anthropology at the time, addressing developments in computing and the emerging fields of cognitive science, artificial intelligence (AI), and human-computer interaction (HCI). Based at Xerox's Palo Alto Research Center, my 'fieldwork' involved a critical

1 Claude Shannon, *The Mathematical Theory of Communication* (Champagne-Urbana 1949).

2 See Donna Haraway, *Modest _Witness @Second_Millenium.FemaleMan_Meets_OncoMouse™: Feminism and Technoscience* (New York 1997); Sara Ahmed, *Differences that Matter: Feminist Theory and Postmodernism* (Cambridge 1998); Anne Balsamo, Reading Cyborgs Writing Feminism, in: *The Gendered Cyborg*, ed. Gill Kirkup (New York 2000), pp. 148 – 158; Londa Schiebinger, Taxonomy for Human Beings, in: *The Gendered Cyborg*, ed. Gil Kirkup (New York 2000), pp. 11 – 37; Claudia Castañeda, *Figurations: Child, Bodies, Worlds* (Durham 2002)

3 For a recent and exemplary text see Charles Goodwin, *Co-Operative Action* (Cambridge, MA 2017).

but also constructive engagement with some emerging projects in the creation of humanlike, interactive machines. What follows is an unpublished paper titled 'Talk with Machines,' presented at the conference *Talk and Social Structure*, a gathering of ethnomethodologists and conversation analysts (a community hereafter referred to as ethno/CA) at the University of California at Santa Barbara in March 1986.[4] This was one of the first conference papers I had ever presented, and definitely my first to an audience of ethno/CA scholars. To say that I was nervous would be an understatement; a condition not lessened when Gail Jefferson, one of the founding and most famous members of the field, commented to me before convening the panel on which I spoke that she had absolutely no idea who I was and so found the job of introducing me a (clearly somewhat annoying) challenge. To my great relief things improved from there, as the paper was well received by an audience delighted by the (novel at the time) idea of bringing conversation analysis to the interface of people and machines.

I rediscovered this paper as a hardcopy document in my basement file cabinet, which I then ran through a scanner and OCR software. This left me with a text that was amusing in the ways that it manifested the argument that it made, including errors of 'recognition' that no human reader would make. The text below has been edited for correctness but is otherwise as it was presented in 1986. I follow it with a brief afterword, reflecting on the book(s) in which versions of the talk were subsequently published, and what I think this all has to say to us now, amidst the massive transformations in the human/machine sensorium in the decades since.

4 For the publication based in this conference, not including my own paper, see Deirdre Boden and Don H. Zimmerman, *Talk and Social Structure* (Berkeley 1991).

Talk with Machines (as presented in 1986)

From questions that I've received over the last several days of the conference, it seems worthwhile to begin by saying something about where in the world it is that I do my work. Xerox began its Palo Alto Research Center in 1970, with the mandate to do research on computing machinery. PARC now comprises roughly 200 or so computer scientists, with a smaller complement of cognitive psychologists, linguists, and two anthropologists, grouped into several laboratories. The lab of which I'm a member includes researchers engaged in the project of building intelligent, conversationally competent machines. It's a preposterous enterprise on the face of it. And yet in the ways that it proceeds, its achievements and its troubles, it offers a unique setting in which to look in detail at what the practical activities of intelligence and conversation actually involve.

When I say that the project of the lab is to build machines that display intelligence and are capable of engaging in conversation, I use the terms intelligence and conversation advisedly. That is, I want neither to appropriate them in a facile way to the description of machines, nor simply to put scare quotes around them and argue that they are *a priori* unique to persons. I want rather to point out that they have now come to be applied to machines, and to suggest that the grounds for that application, how it gets done, represents a new embodiment of practical sociology and a new laboratory for social studies.

It was quite a few years ago that I first read Harvey Sacks' paper "On Sociological Description"[5]. Sacks' primary business in that paper was to consider what are the essential requirements for a science of the social world. Somewhat incidentally, as illustration, he sets up the following scene. Imagine that you are at a trade show where there is a machine that, as it is operating, provides a running account of its own operations. At the show as well are several typified observers, each of whom has a particular stance toward this self-explicating machine. Sacks develops his paper around the differences among the observers. But their differences are of less importance, for present purposes, than their common problem: for each of them the relevant issue in making sense of the machine is the relation between what the machine says, and what it can be seen to be doing.

Metaphorically, the observers at the trade show and the machine observed are members of the society, their common problem that of constructing the coherence of talk and actions. Through the metaphor Sacks is pointing, among other things, to something that seems uniquely identifying of the social world. That is the fact that we human beings have the ability both to produce intelligible actions and to find sense in the actions of others. We are self-explicating, in short, and treat others as such, and that is a fundamental premise for the mutual intelligibility of our interaction.

As life imitates art, recent directions in the development of technology turn Sacks' metaphor into a more nearly literal description of an occasion that any of us may in fact encounter. Specifically, researchers in the field of artificial intelligence are hard at work using the powers of computation to try to build machines that we no longer simply use, but that will interact with us. Over the past seven years I have become deeply interested in this notion of interactive artifacts; its use within the community of computer researchers and designers, its propagation out into the popular press and, most crucially, its basis in what actually goes on when people use computational machines.

The project of designing interactive machines is motivated in part by a practical concern with their usability. The designer of any artifact that is a tool must communicate the artifact's intended use and, in some cases, the rationale for its behavior to the user. This concern is dou-

5 Harvey Sacks, On Sociological Description. *Berkeley Journal of Sociology* 8 (1963), pp. 1–16.

bly relevant insofar as increasingly complex technology is to be usable by people with decreasing amounts of training. The preferred solution is that machines should somehow be self-explanatory; really, that users should be able to discover the machine's intended use solely from information found in and on the machine itself. In physical design, the designer anticipates certain questions such that, in the event, an answer is there ready at hand. So, for example, the user's question 'Where do I grab?' is answered by a handle fitted to the act of grabbing. In the traditional instruction manual, some further classes of inquiry are anticipated, and answers provided. The stepwise instruction set addresses the question "What do I do next?" and the diagram "Where?". In every case, the questions anticipated and answered must be those that any user of the machine might ask, and the occasion for both questions and answers is found by the user.

For the novice engaged in doing some procedural task with a machine, the guiding inquiry is some form of the question "What next?"[6]. Lynch, Livingston and Garfinkel, in their paper "Temporal order in laboratory work"[7] characterise the general task in following instructions as bringing standard descriptions of objects and actions to bear on the specific objects and embodied actions that the instructions describe, right here and right now, this time through. Social studies of the production and use of instructions have identified what Garfinkel calls the "irremediable incompleteness" of instructions and the nature of the work required to carry them out.[8]

The idea of a self-explicating artifact accords well with the notion that using a machine could be like interaction. The interactive machine, in this sense, represents the latest solution to the longstanding problem of providing the user of a tool with instruction in its use. There is also, however, a sense of machine interactivity that is more recent, and is uniquely tied to the advent of computing. The new idea is that the intelligibility of artifacts could be not just a matter of the availability to the user of the designer's intentions for the artifact's use, but of the intentions of the artifact itself. The designer's project in this sense is to imbue the machine with grounds for behaving in ways that are accountably rational; that is, reasonable or intelligible to others including, in the case of interaction, ways that are responsive to the others' actions.

In 1950, the mathematician A.M. Turing proposed a now-famous test for ma-

6 I would now note that the interface as imagined here is the site for inter/action in a rather different sense than that of our interfacing activities at the screen today, which are more ways of inhabiting a place of engagement (with texts, images, others). At the same time, the issues of un/familiarity that I highlight remain at least as salient.

7 Michael Lynch, Eric Livingston and Harold Garfinkel, Temporal Order in Laboratory Work, in: *Science Observed: Perspectives on the Social Study of Science*, eds. Karin D. Knorr-Cetina and Michael Mulkay (London 1983), pp. 205–238.

8 Briefly the argument is that instructions are indexical; that is, they assume the work of finding their relevance for the immediate task at hand, in this particular setting and circumstances, and that work cannot be fully specified given the contingencies of any actual time/place. For an extensive and quintessentially Garfinkelian discussion see Harold Garfinkel, *Ethnomethodology's program: working out Durkheim's aphorism* (Lanham 2002), Chapter 6.

chine intelligence, based on a view of intelligence as rational accountability. Turing argued that if a machine could be made to respond to questions in such a way that a person asking the questions could not distinguish between the machine and another human being, the machine would have to be described as intelligent.[9] Turing expressly dismissed the possible objection that although the machine might succeed in the game, it could succeed by means that bore no plausible resemblance to human thought. The Turing test became the canonical form of the argument that if two information processors, subject to the same input, produce indistinguishable output, then regardless of the differences in their internal operations one is essentially equivalent to the other.

The lines of controversy raised by the Turing test were drawn over a family of programs developed by Joseph Weizenbaum at MIT in the 1960s under the name ELIZA, and designed to support, in Weizenbaum's words, "natural language conversation" with a computer. Anecdotal reports of occasions on which people approached the teletype to one of the ELIZA programs and, believing it to be connected to a colleague, engaged in some amount of interaction without detecting the true nature of the respondent led many to believe that Weizenbaum's program had passed a simple form of the Turing test. Weizenbaum himself, however, denied the intelligence of the program – not on the basis of its interactional success, but on the basis of the underlying mechanism – in a paper that discussed the program's reliance on what Weizenbaum called 'a mere collection of procedures'[10]. In explicating the ELIZA programs, Weizenbaum was concerned with the inclination of human users to find sense in the computer's output, and to ascribe to it an understanding unwarranted by the actual mechanism. This was a process that Harold Garfinkel[11] was at the time identifying as the documentary method of interpretation; a method he discovered, among other places, in his study of students producing the sense of advice provided to them by a counselor. The method is, he argues, the basis for our commonsense knowledge of what are, in the terms of sociology, social structures.

As I said earlier, in looking at the case of ELIZA Weizenbaum was concerned not only with the behavior of his programs, but with the underlying mechanisms that he took to generate that behavior. His concerns elucidate the prevailing view of social structures held by those engaged in engineering machine intelligence. Roughly, the view is that action, and by extension interaction, are epiphenomenal. The structure of observ-

9 While this is the standard characterization of the test, it misses the crucial dimension of gender. For reflections on this lacuna see the preface to Katherine N. Hayles, *How we became posthuman: virtual bodies in cybernetics, literature, and informatics* (Chicago 1999); see also Jennifer Rhee, *The Robotic Imaginary: The human and the price of dehumanized labor* (Minneapolis 2018), pp. 13–14; and Lucy Suchman, Demystifying the Intelligent Machine, in: *Cyborg Futures: Social and Cultural Studies of Robots and AI*, ed. Teresa Heffernan (Basingstoke 2019), pp. 35–61.

10 Joseph Weizenbaum, ELIZA – a computer program for the study of natural language communication between man and machine. *Communications of the ACM* 9 (1966), pp. 36–45.

11 Harold Garfinkel, *Studies in ethnomethology* (Englewood Cliffs 1967).

able behavior is the reflection of underlying cognitive mechanisms that control behavior and give it its sense. On this view, the skillful use of the documentary method by observers or users is seen as a foil that belies the true – read underlying – nature of the artifact.

Commensurate with this view, Weizenbaum's critique of the ELIZA programs was little concerned with the question of conversation. While unmasking the apparent intelligence of his program, he continued to describe it as 'a program which makes natural language conversation with a computer possible.' Nevertheless, as part of his disclaimer regarding its intelligence, Weizenbaum does point to a crucial limit on ELIZA's behavior with respect to talk:

> *ELIZA in its use so far has had as one of its principal objectives the* concealment *of its lack of understanding. But to encourage its conversational partner to offer inputs from which it can select remedial information, it must* reveal *its misunderstanding. A switch of objectives from the concealment to the revelation of misunderstanding is seen as a precondition to making an ELIZA-like program the basis for an effective natural language man-machine communication system.*[12]

Twenty plus years later, conversation continues to elude the capabilities of interactive machines. In the time left to me, I'll tum to some preliminary data, and try to suggest why this might be the case.

The transcripts that you have on your handouts[13] are drawn from a corpus of videotapes of first-time users of a machine designed to be intelligent and interactive. The system is something of a hybrid of old and new technologies; a large photocopier, controlled by a computer-based system intended to act as an artificially intelligent 'expert' in the machine's use. The machine presents to the user a series of video displays on a computer screen attached to the photocopier, composed of text and drawings. Each display either describes the machine's behavior or provides the user with some next instructions. In the latter case, the final instruction of each display prescribes an action whose effect is detectable by the system – buttons pushed, paper taken in and out of trays – thereby triggering a change to the next display.

The objective of the system is that rather than providing a compendium of instructions and leaving decisions of their relevance to the user, instruction should be occasioned by and fitted to the user's actions. To meet this design objective, the system must in some sense be able to find that action's significance. To handle the problem of action interpretation, the designer of this machine adopts the view, common to the cognitive sciences, that the source of the order of situated actions is a cognitive structure – in this case a plan – that stands behind the action and gives it its sense. Action interpretation, on this view, is effectively plan recognition. From the user's response to an initial set of questions

12 Weizenbaum, ELIZA, p. 43.

13 The transcripts can be found in Chapter 9 of Lucy Suchman, *Human-Machine Reconfigurations: Plans and Situated Actions (Learning in Doing: Social, Cognitive and Computational Perspectives)* (Cambridge 2007).

her purposes in using the machine are identified with a goal, the goal invokes an associated plan, and the enactment of the plan is prescribed as a stepwise procedure. The prescribed procedure then provides the system with a ready-made template against which certain of the user's actions can be located in the plan, and the location of the user's action in the plan determines what the system does in response.

The most general aim of my analysis of these recordings was to find the locus of mutual understanding between users and the machine. More particularly, I wanted to compare the user's and the machine's respective 'views' of what happened over the course of events. To document the user's views, I adopted the simple device of asking two people, neither of whom had ever used the machine before, to collaborate on the production of copies. An artifact of such a collaboration is a kind of naturally generated 'think aloud' protocol.

In working to organize the transcripts of the videotapes, I arrived at the simple framework into which your transcripts are arranged:

The User		The Machine	
Actions not available to the machine	Actions available to the machine	Effects available to the user	Design rationale

The framework revealed that the sense of the users' actions was largely unavailable to the machine, and something of why that was the case. Beginning with the observation that what the user was trying to do was, somehow, available to the analyst, one could ask how that was so. The richest source of information for me, as a full-fledged 'intelligent' observer, was the talk, recorded in Column I. In reading the instructions aloud, users located the problem that they were working on. Their questions about the instructions identified the problem more specifically, and further talk provided their understanding of the machine's behavior and clarified their actions in response.

A second, but equally crucial resource was visual access to the actions in which the talk was embedded. Of all of the actions, one could clearly see the very small subset, recorded in Column II, that were actually detected by the system (that is, those that actually changed its state). From the machine's 'point of view,' correspondingly, one could see how it was that those available traces of the user's actions – the user's actions seen, as it were, through a pin-hole – were mapped onto the plan, under the design assumption that, for example, button x pushed at this particular point in the procedure must mean that the user is doing y.

The framework proved invaluable for considering seriously the idea that user and machine were interacting. By treating the central columns as the human-machine 'interface,' one could compare and contrast them with the sense made by the users as displayed in column 1 and the design rationale on which the machine state changes were based in column 4. This comparison located precisely the points of confusion, as well as the points of intersection or 'mutual intelligibility,' at least for practical purposes. I problematize 'mutual intelligibility' here insofar as it presupposes capacities for the detection and repair of troubles in conversation. This takes us back to the observation made by Weizenbaum with respect to the ELIZA program, that its intelligence was limited in the first instance by its inability to recognize misunderstanding on the part of its interlocutor or to reveal its own.

The idea that language ability is the mark of intelligence is found in the notion of competent member of the society, as used by Garfinkel and Sacks who explain that they "do not use the term [member] to refer to a person. It refers instead to mastery of natural language"[14]. At the same time, it would be misleading to take language as such as the problem for artificial intelligence research. The fact that the so-called natural language problem has proven to be such a difficult one suggests that language ability is not an isolable skill but is part and parcel of the more general competence that Garfinkel and Sacks identify as competent membership. It's instructive to view machines, in this way, as resource-limited participants in an interaction. In his recent book on ethnomethodology, John Heritage summarizes the findings of conversation analysis with respect to institutional interaction as (1) the selective reduction of the full range of conversational practices and (2) the specialization of particular procedures taken from ordinary talk.[15] The difference between this and human-machine communication is that institutional asymmetries, say between doctors and patients, are based in ordinary conversation. Mutual understanding continues to be founded in the latter, with further work being done through the addition of special institutionally established constraints. Comparing the interactional resources of user and machine, in contrast, reveals a fundamental asymmetry in the available means by which each produces and construes the features of their shared situation. Because of this asymmetry, human-machine interaction is less a matter of simulating human communication, than of engineering alternatives to talk's situated properties.

Afterword

The original book *Plans and Situated Actions: The problem of human-machine communication*[16] tried to make several, interconnected arguments, which might be helpfully summarized here. At the core (as the title suggests) was a reconceptualization of the relation of plans to the situated activities of their creation and use. Most importantly, I argued that while the plan-based AI dominant at the

14 Harold Garfinkel and Harvey Sacks, On formal structures of practical actions, in: *Theoretical Sociology: Perspectives and Development*, eds. John C. McKinney and Edward A. Tiryakian (New York 1970), p. 342.

15 John Heritage, *Garfinkel and ethnomethodology* (Cambridge 1984).

16 Lucy Suchman, *Plans and Situated Actions: The problem of human-machine communication* (Cambridge 1987).

time treated a plan as an algorithmic specification that determines action, plans are better conceptualized as a genre of artefact created as a resource for action and communication. Plans are made in anticipation of doing something, often referenced as the activity unfolds, and sometimes cited after the fact to account for how things went (or went differently). Plans may be only conceptual or discursive, or they may be materialized in a variety of media, as flow charts, lists, directions and the like.

One of the common mistranslations of this argument is that while sometimes our actions go as planned, often they don't; it's the latter case, on this reading, that calls for situated actions. But the argument is a much more fundamental one. The argument is that even in cases where everything does go 'according to plan,' the implementation of a plan is always, and necessarily, a situated activity not fully specified in the plan itself. In other words, given the contingencies of any actual occasion of action, every plan presupposes capacities of cognition and (inter)action that are not, and cannot ever be, fully specified. This isn't a problem for human actors, who rely on a range of ordinary (or extraordinary) competencies to bring plans into relation with the circumstances of action. But it is a profound, and unsolved, problem for computational machines.

The other central thesis of the original book addresses the implications of the argument just summarized for human-computer interface design. This begins with recognition that like action, human communication presupposes a range of taken-for-granted competencies. Central among those is the ability to engage in collaborative sense making. The latter is not just a matter of recognizing meanings that are pre-given, but of engaging in an open-ended way in the co-production of mutually accountable (inter)action. Taking this view of communication seriously as a basis for analyzing interactions at the interface reveals the significant and enduring limits to communicative abilities on the machine side and highlights just how difficult a problem the design of human-machine communication really is.

In *Human-Machine Reconfigurations*[17] I revisit these arguments through an annotated version of the original text, reframed by a series of new chapters considering relevant developments both in HCI and AI, and in social studies of science and technology. The annotations allow me to express changes in my ways of thinking about the issues discussed in the original text, while the new chapters provide an opportunity to update both theoretical resources and empirical examples. While it is unquestionably the case that the fields of HCI and Artificial Intelligence have advanced significantly over the decades since the original publication of *Plans and Situated Actions*, I believe that the argument put forward in that book still holds. The primary developments in both fields rely upon a combination of expanding networked infrastructures, massive data capture,

17 Suchman, *Human-Machine Reconfigurations*.

storage capacity and processing speed, and associated elaborations of technologies and techniques of data analysis. Yet as Chapters 12 through 14 of *Human-Machine Reconfigurations* suggest, there has been notably little progress towards the creation of systems capable of engaging in social interaction. The reasons for this lie in the situated qualities of both action and communication, specifically their reliance on capacities of generative co-production of a contingently unfolding and dynamic world. These persistent problems are indicative of enduring differences that matter between persons and machines; they call in turn for design practices that engage with those differences creatively, rather than aiming to obscure or even to erase them.

Most importantly, I hope that the trope of re/configuration can help to underscore the central question of how we conceptualize or *figure* humans and machines respectively; what that means for the ways in which we *con*figure human-machine relations both imaginatively and materially; and how we might *re*figure and *re*configure persons, machines and their relations as part of a wider project of displacing the autonomous Human subject and his (subservient) Others as a previously unexamined foundation for design. This involves continuing attention to differences that matter between humans and machines, as part of the larger project of enabling more just and sustainable futures.

References

Ahmed, Sara, *Differences that Matter: Feminist Theory and Postmodernism* (Cambridge: Cambridge University Press, 1998).

Balsamo, Anne, Reading Cyborgs Writing Feminism, in: *The Gendered Cyborg, ed. Gill Kirkup* (New York: Routledge, 2000), pp. 148–158.

Boden, Deirdre and Don H. Zimmerman, *Talk and Social Structure* (Berkeley: University of California Press, 1991).

Castañeda, Claudia, *Figurations: Child, Bodies, Worlds* (Durham: Duke, 2002)

Garfinkel, Harold, *Studies in ethnomethology* (Englewood Cliffs: Prentice-Hall, 1967).

Garfinkel, H., *Ethnomethodology's program: working out Durkheim`s aphorism* (Lanham: Rowman & Littlefield Publishers, 2002).

Garfinkel, H. and Harvey Sacks, On formal structures of practical actions, in: *Theoretical Sociology: Perspectives and Development*, eds. John C. McKinney and Edward A. Tiryakian (New York: Appleton-Century-Crofts, 1970).

Goodwin, Charles, *Co-Operative Action* (Cambridge, MA: Cambridge University Press, 2017).

Haraway, Donna, *Modest _Witness @Second_Millenium.FemaleMan_Meets_OncoMouse™: Feminism and Technoscience* (New York: Routledge, 1997).

Hayles, N. Katherine, *How we became posthuman: virtual bodies in cybernetics, literature, and informatics* (Chicago: University of Chicago Press, 1999).

Heritage, John, *Garfinkel and ethnomethodology* (Cambridge: Polity Press, 1984).

Rhee, Jennifer, *The Robotic Imaginary: The human and the price of dehumanized labor* (Minneapolis: Minnesota University Press, 2018).

Sacks, Harvey, On Sociological Description. *Berkeley Journal of Sociology* 8 (1963), pp. 1–16.

Schiebinger, Londa, Taxonomy for Human Beings, in: *The Gendered Cyborg*, ed. Gill Kirkup (New York: Routledge, 2000), pp. 11–37.

Shannon, Claude, *The Mathematical Theory of Communication* (Champagne-Urbana: University of Illinois Press, 1949).

Suchman, Lucy, Demystifying the Intelligent Machine, in: *Cyborg Futures: Social and Cultural Studies of Robots and AI*, ed. Teresa Heffernan (Basingstoke: Palgrave Macmillan, 2019), pp. 35–61.

Suchman, L., *Human-Machine Reconfigurations: Plans and Situated Actions (Learning in Doing: Social, Cognitive and Computational Perspectives)* (Cambridge 2007).

Suchman, L., *Plans and Situated Actions: The problem of human-machine communication* (Cambridge 1987).

Weizenbaum, Joseph, ELIZA – a computer program for the study of natural language communication between man and machine. *Communications of the ACM* 9 (1966), pp. 36–45.