This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

1

ARTIFICIAL COMMUNICATION? ALGORITHMS AS INTERACTION PARTNERS

COMMUNICATION WITH ALGORITHMS

Whether algorithms can "think" is still very uncertain.[1] What is more certain is that contemporary algorithms, based on machine learning and big data, can participate in communication. Today's algorithms can act as communication *partners*. Precise estimates are difficult, yet it is claimed that bots are the authors of approximately 50 percent of online traffic.[2] Millions of Twitter users are bots,[3] most fake Facebook accounts are created by automated programs,[4] and at least 40 percent of Wikipedia editing is carried out by computer-controlled accounts.[5] According to an evaluation by the Oxford Internet Institute, highly automated accounts generated close to 25 percent of all Twitter traffic about the 2016 US presidential election.[6] That Google and Facebook are driven by algorithms is well known, with the paradoxical consequence that the "discovery" that human operators guide the selection of news in Facebook's list of tending topics was perceived as a scandal.[7] Automated systems are also used in personalized communication; on Gmail, Smart Reply recognizes emails that require responses and generates perfectly adequate natural language answers on the fly.[8] Spotify's most popular compilation, Discover Weekly, is entirely assembled by an algorithm—as well as its Release Radar, a hyperpersonalized playlist of the latest tracks.[9]

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

Algorithms can also be the authors of texts and books in traditional printed media. Companies like Narrative Science[10] and Automated Insight[11] have developed algorithms to produce texts that are indistinguishable from those written by human authors: newspaper articles, brochures for commercial products, textbooks, and more. Philip Parker, professor at INSEAD in Fontainebleau, patented a method to automatically produce plausible and informative books, including more than one hundred thousand titles already available on Amazon.com. Robo-journalism is regularly used by the Associated Press and many companies like Samsung, Yahoo, Comcast, and others.[12]

Often, moreover, we talk directly with algorithms. We routinely book train tickets, make appointments, and ask for assistance via dialogue with chatbots. Digital personal assistants like Apple's Siri, Amazon's Alexa, or Google Assistant use natural language interfaces to answer new questions, manage calendars, or offer individual suggestions and recommendations. In many cases, these programs seem to know the users better than their human partners and often better than the users themselves,[13] anticipating their needs and demands even before they emerge.

How should we interpret these amazing developments in the communicative performance of algorithms? Communication as we know it normally takes place between humans (or at most between humans and other living beings). If machines now participate in communication, does this mean that machines have become human, or at least that they have learned to reproduce the intelligence of human beings? Are we witnessing the realization of the ideal of an artificial intelligence (AI) that has accompanied the progress of digitization from its beginnings,[14] or are we facing something different that requires a transition to a different way of thinking?

In this chapter, I argue that what we can observe in interactions with algorithms is not necessarily an artificial form of intelligence, but rather an artificial form of communication. Intelligence and communicative capacity are not the same thing. Algorithms are able to act as communication partners—whether they are intelligent or not is another matter. Modern machine-learning algorithms are so efficient not because they have learned to imitate human intelligence and to understand information, but rather because they have abandoned the attempt and the ambition to do so and are oriented toward a different model. Machine-learning

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

algorithms that use big data, I claim, are artificially reproducing not intelligence but communication skills, and they do so by parasitically exploiting the participation of users on the web.

The concept of communication must be reconsidered. Can we still talk of communication when one of the partners has no understanding of the information conveyed? What does this mean for social information processing? In the following pages I try to give an answer to these questions by examining the notion of communication and proposing a concept that does not presuppose any sharing of thoughts between participants. In the final part of the chapter, I show the consequences of the shift from intelligence to communication in the design of algorithms and, in particular, in the idea of autonomous-learning programs.

ARTIFICIAL COMMUNICATION

The protagonists of the current communicative revolution are algorithms, but algorithms by themselves are not new. The concept of the algorithm dates back at least to the Middle Ages, the term itself having roots in the latinization of "al-Khwarīzmī," the name of a Persian mathematician from the ninth century.[15] What is new is the recent exploitation, made possible by the use of big data and machine-learning techniques, of a specific feature of algorithms—their lack of intelligence.

The advantage of algorithms has always been one of not requiring any "creative" thought in their execution.[16] As with computers, they carry out operations in sequence according to precise instructions, proceeding mechanically.[17] In algorithms, and in the digital management of data that relies on them, information processing and mapping have nothing to do with understanding—indeed, in many cases a need for understanding would rather be an obstacle.[18] As the number of elements to be analyzed grows (up to today's incredible scales of petabytes and zettabytes) the operations of these machines become less and less comprehensible[19]—yet their performance not only does not decrease, but gradually becomes more precise and reliable. Digital machines have other ways to test the correctness of their procedures.

The communicative relevance of algorithms is actually related to their independence from understanding. We are facing a way to process data

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

(and to manage information) that is different from human information processing and understanding.[20] My assumption is that this difference is not a liability but instead is the very root of the success of these technologies. Just as human beings first became able to fly when they abandoned the idea of building machines that flap their wings like birds,[21] digital information processing managed to achieve the results that we see today after abandoning the ambition to reproduce in digital form the processes of the human mind. Now that they no longer try to resemble our consciousness, algorithms have become more and more able to act as competent communication partners, responding appropriately to our requests and providing information neither constructed nor reconstructable by a human mind.[22]

This is already evident in our practical use of algorithms, but not always in our theorizing about them. The metaphors used in the field of big data and machine learning retain a reference to the human mind and its processes. Take, for example, the widespread idea that recent procedures of deep learning are so effective because they are based on biological neural networks, replicating the functioning of the human brain. As most researchers admit,[23] however, we still know very little about the workings of our brains, which makes the analogy quite curious—does it make sense to take our ignorance as a model?[24] If machines no longer try to understand meaning as happens in the human mind, shouldn't we find a different, more fitting, metaphor?

Recent approaches to big data are very different from the programs of AI research from the 1970s and 1980s, which aimed to reproduce the processes of human intelligence, by imitation or by analogy ("strong" or "weak" AI, respectively), with a machine.[25] This is no longer the case. As some AI designers explicitly declare, "We do not try and copy intelligence"[26]—for this would be too heavy a burden. Translation programs do not try to understand the documents they translate, and their designers do not rely on any theory of language learning.[27] Algorithms translate texts from Chinese without knowing Chinese, and their programmers do not know it either. Spell checkers correct typographical errors in any language, knowing neither these languages nor their (varying) conventions. Digital assistants operate with words without understanding what words mean, and text-producing algorithms "don't reason like people in order to write like people."[28] Examples multiply across all

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

areas in which algorithms are the most successful. Algorithms competing with human players in chess, poker, and Go have no knowledge of the games nor of the subtleties of human strategies.[29] Recommendation programs using collaborative filtering know absolutely nothing about the movies, songs, or books they suggest, yet operate as reliable tastemakers.[30] Computer-based personality judgments work "automatically and without involving human socio-cognitive skills."[31]

These programs are reproducing not intelligence but rather communicative competence. What makes algorithms socially relevant and useful is their ability to act as partners in communicative practices that produce and circulate information, independently of their intelligence. Could we say that machine-learning programs realize not an artificial intelligence but a kind of artificial communication, providing human beings with unforeseen and unpredictable information? Maybe our society as a whole becomes "smarter" not because it artificially reproduces intelligence, but because it creates a new form of communication using data in a different way.

That the focus of the web is on communication rather than on intelligence is confirmed by the rampant success of social media, which had not been foreseen in any model of digital evolution. The web today is organized more through contacts, links, tweets, and likes than by meaningful connections between content and between sites[32]—it is driven by communication, not by meaning and understanding.[33] Every link (every communicative behavior) is treated as a like, and "liking" and "being like" have also been equated.[34] Everything that happens online is used as a fact and thus becomes a fact, having consequences and producing information.

CAN WE COMMUNICATE WITH PARTNERS THAT DO NOT THINK?

If we are to examine communicative competence, and as such to shift our reference from (artificial) intelligence to (artificial) communication, we must start asking different questions. The focus is no longer on the participants (on whether they are human or not, and what it means to be human in a digital world);[35] it is on the process of producing information.

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

Is what happens in the interaction with algorithms on the web "communication," or do we need to modify the concept? Does it still make sense to speak of communication when data processing is performed by a machine that does not understand the content being communicated? Are the users of web services communicating, and if so, with whom? The answers to these questions depend on our concept of communication, and the concept should be powerful enough to also cover interactions with machines.

Most concepts of communication require that the mental processes of its participants converge on some common content. According to the Latin root of the term "communication" (*communicatio*), it is assumed that partners have the same thought in common, or at least part of it. Communication happens if, at the end of the process, the receiver gets at least some of the information that the issuer put into the channel. Even considering noise and differences in coding/decoding, interpretation and competence, the idea is that in a successful communication, some element of the identity of information must be preserved.[36] The problem with this approach, however, is that in the interaction with machines, we are dealing with a situation in which one communication partner is an algorithm that does not understand content, meaning, or interpretation. It deals only with data.[37] A user, therefore, shares no information (not even partially) with their interlocutor, because the interlocutor does not know any information. Can we still say that they are communicating?[38] Are we dealing with an "aberrant" condition, or with an unprecedented form of communication?[39]

My argument in the following sections follows Niklas Luhmann's theory of social systems and his notion of communication.[40] I claim that the very reasons why Luhmann's approach has been criticized (and often misunderstood) are now the very reasons that make it particularly appropriate to deal with novel aspects of digital communication. Luhmann explicitly refused to define communication in reference to conscious subjects. The concepts of subject and individual, he argued, act only as empty formulas for a very complex phenomenon that falls within the competence of psychology and does not directly interest sociologists or communication theorists.[41] The objects of sociology are not subjects but communications, in which the thoughts of the participating individuals

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

(which are and remain indispensable) are not the constituent elements. Luhmann's theory of communication, therefore, distances itself from psychic processes and their communicative role, thereby breaking with this tradition in sociology.

That Luhmann's concept of communication is not based on psychic content and requires no sharing of thoughts among participants becomes a great advantage when dealing with algorithms that do not think. In all forms of communication, Luhmann argues, information is different for everyone and always relative to a specific observer.[42] But a common identity of information among participants is itself not required for communication.

Luhmann's simple yet very effective innovation is to define communication starting from the receiver, rather than from the issuer. According to his approach, communication comes about not when somebody says something,[43] but when somebody understands that someone said something. One can write entire books and make elaborate speeches, but if no one reads or listens or even notices it, there was no real communication. Yet if a receiver understands information that (they believe) someone uttered, communication takes place—whatever this information is to the receiver, and whatever the issuer had in mind (or indeed did not have in mind). I do not have to enter Proust's mind to understand *À la recherche du temps perdu*—an understanding that I may gain in another language and experience a hundred years after the work was written. I only have to understand his communication—in my way, and according to my thoughts. The information I get from Proust's work will inevitably be different from Proust's thoughts, which makes communication an endless, fascinating process of discovery.

Since information is always relative to the observer, the receiver always obtains information that is different from what the utterer had in mind.[44] The thoughts of the participants are not part of communication itself, leading to an infinite variety of individual understandings. The task of sociology and of communication theory is to analyze how this diversity of understandings can still produce forms of coordination.[45] Even without a shared understanding, not every interpretation is socially acceptable, and explicit misunderstandings are an exception, rather than the rule.

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

The fundamental power of this notion of communication, as concerns our focus on algorithms, pertains to the fact that, in its noninclusion of the thoughts of participants,[46] such a notion could in principle extend to participants that do not think (such as algorithms). If we start from the perspective of the receiver, what counts is whether they take something to be a communication partner. Since in communication the receiver attributes the information obtained to their counterpart, however, the partner is normally a human being;[47] we do not normally communicate with machines, to which this kind of information is not attributed.

This does not mean that machines cannot be informative, however. We habitually gather information from objects in the world and from machines—our watches, for example, tell us what time it is—but we do not attribute the information to the watch. Our watch informs us about the time, but only because it was constructed by someone in order to convey that information. It does not develop its own way of dealing with time and does not decide itself how to calculate it. We do not communicate with our watch. Yet algorithms are confronting us with an unprecedented situation. From algorithms we get information that often was not planned or available in advance and was unknown to the programmers themselves. Self-learning digital programs autonomously develop their procedures and identify patterns, which they use to produce their answers to our requests. In conversations with digital personal assistants or social bots, for example, the information we get did not exist before we formulated our request and is produced by the machine expressly to respond to that request. Nobody knew that information in advance or decided how to produce it—the algorithm generated it itself. The production of information can be attributed only to an interactive partner, as in communication—but in this case the partner is not a human being, but a machine.

When we interact with algorithms, then, do we communicate with them? Does their role in communication require us to consider them as possible partners? It is a tricky matter. The issue of communication with machines and the current relevance of the Turing test depend on the answer to this question. The problem here is not whether the person is or is not aware of dealing with a machine, for doing so is now an everyday occurrence, and one where such a question is usually not relevant. Today

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

our counterparts are often bots (in online services, video games, social media) even if we are not aware of it—and when we are aware, as with personal assistants, we do not normally care.[48] What matters is whether the interaction from which we gather our information has the features of a relationship with a contingent, autonomous partner.

VIRTUAL CONTINGENCY

Contingency implies selection and uncertainty. It means that there are a number of possible options to choose from, and our decisions could always be different.[49] However, algorithms by definition do not know uncertainty; they do not choose between possibilities, nor are they creative, being designed to follow the instructions that program their behavior. In this sense, algorithms are not contingent—which is why they can operate so efficiently and reliably. Just like traditional machines, we expect algorithms to be neither unpredictable nor idiosyncratic, even when they deliver information. Different watches should all indicate the same time to all users, if they work properly. As von Foerster observed, if the outcome of a traditional machine becomes unpredictable, we do not think that it is creative or original—we think that it is broken.[50] We do not care about the moods nor the perspectives of machines, only about their results. We repair them precisely to restore their predictability.

Recent algorithms, however, are different: their semblance of contingency is an essential feature. Even if these machines follow a completely determined course, we want their outcomes to be unpredictable, and to produce something we do not yet know—that is, new information appropriate to a given interaction with a user. The expected outcome is not predicted by anyone and, in the case of self-learning algorithms, could not be predicted—that's why we use algorithms, and why they appear creative. The dilemma faced by designers, therefore, is to build machines that are creative yet controlled at the same time—to program the production of unpredicted outcomes. Even if the machine is completely determined, its behavior should appear contingent and react to the contingency of the user. Cozmo, for example, a real-life toy robot based on a series of machine-learning algorithms,[51] is "programmed to be unpredictable" without being simply random.[52] Cozmo's behavior must appear

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

responsive and appropriate to the user, otherwise it is no fun. A personal assistant like Alexa should respond appropriately to the user's requests, producing new and relevant information in the course of the interaction. The paradoxical purpose of programming intelligent algorithms is to build unpredictable machines in a controlled way. The goal is a controlled lack of control.[53]

How can an algorithm act as a contingent partner in an interaction? In some cases, the contingency of a machine is simply the projection of the contingency of its user. This happens, for example, with the robotic toys studied by Sherry Turkle that work well as communication partners because children or elderly people interacting with them project onto them their own contingency.[54] This always happens with dolls and puppets, with which children play as if the toys understand and respond to their behavior. What is reflected in the performance of robotic toys—and what makes them more fun than traditional dolls—is not the ability to understand but the ability to "perform understanding" in elaborate and seemingly reactive ways.[55]

Self-learning algorithms go further and do something more enigmatic. When a user interacts with a learning algorithm,[56] they face a contingency that is not of their making—although it also does not belong to the machine. The perspective that the machine presents is still a reflected perspective—because the algorithm inevitably does not possess its own contingency—although one which does not simply reflect the perspective of the user. Instead, what the algorithm reflects and represents is the perspectives of *other* observers; what the user observes through the machine is the outcome of the processing of other users' observation. I call *virtual contingency* the ability of algorithms to use the contingency of users as a means of acting as competent communication partners.

GOOGLIZATION

Where do algorithms find the contingency they reflect? How do they access the external perspectives they elaborate and present to their communication partners? To be able to participate in communication, algorithms must be on the web.[57] As smart and sophisticated as algorithms can be, artificial communication would not be possible without the

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

web—a power only realized once algorithms were taken online. The path-breaking effect of the “participatory web” (Web 2.0, and possibly 3.0)[58] was not so much customization, but rather an inclusion and exploitation of virtual contingency.[59] Algorithms parasitically “feed” on contributions by users and actively use them to increase the complexity of their own behavior—along with the complexity of their communicative capacities. In interactions with learning algorithms, I claim, users experience an (artificial) form of unpredictability and reflexivity. Such interactions artificially reproduce the conditions of communication.

The prototype of this approach is Google, and this is also the reason for its success. The breakthrough came in 1998 with the introduction of link analysis in the World Wide Web.[60] Previously, information retrieval took place by way of searching through a limited, unlinked, static collection of documents. The organization and categorization of information were entrusted to specialists such as librarians, journal editors, or experts in various fields. Link analysis, instead, extends to the web and introduces a form of information retrieval that becomes huge, dynamic (unlike traditional documents, web pages are constantly changing their content), hyperlinked, yet above all, self-organized. The structure is decided not by experts but by the dynamics of the web. And it is incomparably more efficient.

The design of Google’s PageRank algorithm marked a conceptual turn, “inventing” the internet as we know it today.[61] Its authors, and later owners of the company, describe it as starting from the idea of exploiting the link structure of the web as a large hypertext system.[62] The key insight was to determine which pages are important and for whom, disregarding the content of the pages themselves. To appropriately decide the ranking of pages responding to users’ requests, the idea was to use information that is external to the web pages themselves and which rather refer to what other users did in their previous activity. In other words, to decide which pages are important, PageRank does not look to see what the pages say or how they say it, but instead looks at how often they were linked to and by whom. The ranking is based on the number of *backlinks* to the pages (how many times they have been pointed to by other websites) and on their importance—where the “importance” of backlinks depends itself on how many links they in turn have. The definition of “relevance”

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

is openly circular: a page has high rank if the sum of the ranks of its backlinks is high,[63] including both the case of a page with many not particularly authoritative backlinks and the case of a page with a few highly linked backlinks.

The genius of PageRank's innovation lies in relinquishing the goal of understanding what the page says and relying solely on the structure and the dynamics of communication. Google's creators did not try to come up with a great organizational scheme for the web based on experienced and competent consultants, as did competing search engines like Altavista and Yahoo.[64] They did not try to understand and build an algorithm that understands; instead, "they got everyone else to do it for them" by surfing the net and making connections.[65] Content comes into play later, as a result and not as a premise. Google uses the links to learn not only how important a page is, but also what it is about. If the links to a given page use a certain sentence, the system infers that the sentence accurately describes that page and takes this into account for later searches. The algorithm is designed to apprehend and reflect the choices made by users,[66] activating a recursive loop in which the users use the algorithm to get the information, their searches modify the algorithm, and the algorithm then impinges on their subsequent searches for information. What the programmers design is only the algorithm's ability to self-modify. What and how the algorithm selects depend on how users are using it.

This system has been developed further to take into account factors beyond popularity, such as users' click behavior, reading time, and patterns of query reformulation.[67] As Google declares in the InsideSearch pages of its website, algorithms today rely on more than two hundred signals and clues referring to "things like the terms in websites, the freshness of content, your region."[68] The company produced a "Knowledge Graph" that provides a semantic connection between billions of entities and allows for more rapid and appropriate responses, also including information and results not yet thought of by anyone. The "intelligence" of the system, however, derives from its use of previous user activity and from sources of information already available on the web, from Wikipedia to databases of common knowledge. As John Gianandrea, director of engineering at Google, declared: when one is googling "Einstein," "We're not trying to tell you what's important about Einstein—we're trying to

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

tell you about what humanity is looking for when they search."[69] The intelligence of the system is the intelligence of the users that the algorithm exploits to direct and organize its own behavior.

Google has become the symbol of an approach that can be found in other successful projects on the web.[70] Since 2003 the term "googlization" has been employed to describe the spread, in more and more applications and contexts, of a model that does not rely on traditional status makers like editors or experts, but "feeds" on the dynamics of the web to organize its operations and even itself.[71] Vaidhyanathan argues that the web is guided by a "googlization of everything" that takes advantage of the operations performed by users to produce a condition in which "Google works for us because it seems to read our minds."[72] In reality, Google does not need such powers. Rather, Google merely uses the *results* of what we had in mind in order to produce that which we did not.

Google, along with other systems that work in the same way, feeds on the information provided by users to produce new information, which is introduced into the circuit of communication. It is this information that users obtain from their interactions with algorithms, and which can only be attributed to the algorithms themselves. When speaking of interactions with algorithms, it makes no sense to refer only to the perspective of those who entered the data, because they could not know precisely how the data would be used. Similarly, it makes no sense to refer to the perspective of what the algorithm itself meant, because it did not *mean* anything. Constraints and orientation depend not on intentions but on programs, which are normally inaccessible.[73]

Algorithms make selections and choices based on criteria that are not random, instead reflecting and elaborating upon the indeterminacy of their participants. Users receive contingent responses that react to their contingency using the contingency of other users. While they do not directly communicate with this assortment of other users, the result of this interaction is a specific answer to a specific question which would not exist if other users were not also engaged in communication. Google and similar models appear to communicate with their users, and are able to do so precisely because they do not try to understand content. They do not artificially reproduce intelligence, but directly engage in communication. In light of this, are we dealing with a new form of communication?

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

WHAT ALGORITHMS LEARN

If interaction with learning algorithms is communication, we are dealing with a form of artificial communication. By “artificial” here I mean more than a communication that was produced by someone, since all communication would be artificial in this sense.[74] A communication is artificial when it involves an entity—the algorithm—that has been built and programmed to act as a communication partner by someone who does not participate in the communication. It is communication with an artificial partner.[75]

Considering artificial communication more closely can help us to explore the enigmatic ability of algorithms to learn. Recent algorithms using big data can learn to recognize images never encountered before, carry on conversations about unknown topics, analyze medical data and formulate diagnoses, as well as anticipate the behavior, the reasoning, and also the wishes of users. On the basis of their ability, we can (or will soon be able to) ride self-driving cars, translate online phone calls from one language to another in real time, and use digital assistants to deliver the information we need at any given moment. But what do learning algorithms learn? And who teaches them?

Self-learning algorithms can apparently learn by themselves. Whether supervised, semi-supervised or unsupervised, learning algorithms decide autonomously how to learn and what to learn. They are able to use data to learn functions they have not specifically been programmed for.[76] Their programmers only design a set of procedures that should allow the machine to develop its own way to solve a task, or even (in the case of unsupervised learning) to determine its own task, finding structures in data such as groupings or clusters. The programmers do not know what the machine is learning, instead they teach it to learn autonomously.

This is not an easy task, especially if it is an explicit goal. Michael Warner, a Carnegie Mellon–trained robotic researcher, claims that in many situations where you invoke machine learning, you do so “because you do not really understand what the system should do.”[77] The programmers give indications that the learner will use in its own way, and then see if the result is satisfactory. When a learning algorithm is expected to learn to play a game, for example, the programmers do not teach it the moves,

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

or even the rules of the game. The machine makes random moves, and after a number of attempts, the programmers tell it if it has won or lost. The learning algorithm uses these "reinforcements" to calculate in its own way an evaluation function that indicates which moves to make—without making predictions, without a game strategy, without "thinking" and without imagining the perspective of its opponent.[78] Nobody knows what the machine learned, or how it did so, but the processes involved produce amazing performances, such as defeating the most qualified champions of games like chess or Go. As the programmers of AlphaGo, the computing system built by Google to play Go, put it: "Our goal is to beat the best human players, not just to mimic them."[79]

AlphaGo learned to become an outstanding Go player and beat the best players in the world. For this purpose it did not learn to play the game *like* human players (or better). In fact, the algorithm did not *learn* Go—it learned to *participate in* Go, taking advantage of the moves of other participants to develop and refine its own moves. AlphaGo was originally trained with data from a server that allowed people to play against each other on the internet. The players were all amateurs and their skills were rather coarse, but the program refined these skills enormously by playing millions of games against itself. AlphaGo and other game-oriented algorithms learn via self-play, refining their skills with a trial-and- error process.[80] The system learns "not just from human moves, but from moves generated by multiple versions of itself."[81]

These procedures confirm the hypothesis that algorithms learn not to think but to participate in communication, that is, to (artificially) develop an autonomous perspective that allows them to react appropriately and generate information in their interaction with other participants. What AlphaGo thinks or does not think is irrelevant to its performance. It is competent, reactive and creative—and can also be surprising. It is a perfect game partner even and precisely because it does not think like a human player. Through training, algorithms do not become more intelligent; they just learn to play better. The programmers themselves do not understand the "reasoning" of the algorithm. When the programmers indicate that the algorithm is "wrong," they merely signal that there is an error, without indicating what it is. The algorithm uses these reinforcements

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

to direct its own behavior, which becomes more and more refined and effective—and less and less comprehensible.[82]

LEARNING TO LEARN FROM MACHINES

Learning algorithms learn to participate in communication, and they can do so because they do not need to understand what people have in mind. For the same reason, people can themselves learn from their interactions with learning algorithms, even if they don't understand them.

An example is the legendary move 37 in the game of March 2016 between Lee Sedol, one of the world's top Go players, and the algorithm AlphaGo. The move was described by observers as absolutely surprising and unpredictable. "It was not a human move" and couldn't have come to any human mind.[83] It was actually produced by an algorithm that does not have a mind, yet it allowed AlphaGo to win the game and then the match. Later, this incomprehensible move triggered a process of learning by human players that profoundly transformed the practice of the game. Revisiting move 37, Go players found it to be brilliant, and took it as a clue to rethink their game strategies, dramatically improving them—thereby learning from AlphaGo.[84] Following this revision, Lee Sedol himself produced the celebrated, highly unlikely (1 in 10,000) move 78 ("The Touch of God") in his fourth game with AlphaGo, the game he was able to win.[85]

Lee Sedol defeated the algorithm by reinterpreting with human skills a move that no human being could have devised. The incomprehensible behavior of AlphaGo highlighted possibilities that could be processed by human players in their own way to produce a meaningful result. It is likely that the algorithm later incorporated move 78 in its procedures and learned to manage the move and its consequences;[86] however, it would not have been able to do this without the human being that devised it. No algorithm, however advanced its ability to self-learn, can generate possibilities that are not implicit in the data supplied.[87] No algorithm can independently generate contingency, but algorithms can process human-generated contingency in unprecedented ways, ways that might generate further possibilities and further contingency in interactions with human beings.

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

Even and especially if the algorithm is not an alter ego, if it does not follow a strategy, and if it does not understand our reasoning, human users can still learn from their interactions with an algorithm to develop their own strategies. Not through understandable algorithms that can trigger understandable processes, but through obtaining and using clues that no one could have imagined, thereby changing their way of observing. People using their intelligence to learn from non-intelligent machines is an opportunity for increasing the complexity of communication. In the case of Go, it was a matter of game strategy, yet the same mechanisms can be applied to designing social algorithms in general.[88]

Yet relying on black boxes is not reassuring, especially when one knows that their operations are not immune from biases and errors of various kinds.[89] The recent branch of research on "explainable AI" attempts to respond to this concern by looking for procedures that enable machines to provide explanations of their operations.[90] But explaining incomprehensible processes seems a hopeless task. As Weinberger claims, it would amount to someone seeking to force AI "to be artificially stupid enough that we can understand how it comes up with its conclusion."[91] Yet algorithms as communication partners can be explainable without being understandable.[92] The requirement would be that they have sufficient communicative competence to respond to requests for clarification from their interlocutors in an appropriate, comprehensible, and controllable way. What users understand by way of an explanation of the machine does not have to be the finer processes of the machine. This actually happens often in human explanations as well, insofar as they offer clues to make sense of a communication without giving access to the psychic processes of the partner—and is the direction in which the design of advanced algorithms is currently moving.[93]

CONCLUSION

Interactions with algorithms are a challenge for sociology and communications theory. Whether one decides that they are a specific form of communication and that the concept of communication should be amended accordingly, or one decides that algorithms are not communication partners, the task of communication theory is still to adequately describe the

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a portion of the eBook at doi:10.7551/mitpress/14189.001.0001

development of these digital processes. We must be able to show how interactions with algorithms affect communication in society in general and to provide insights that can help to direct the work of those who design and write algorithms.[94]

In more and more areas, the familiar reference to (artificial) intelligence becomes unhelpful, whether these are cases in which communications are attributed to things (e.g., the Internet of Things) or cases in which communications are treated as things (e.g., the digital humanities). Does this mean that we are moving toward a state of widespread intelligence where there will be no separation between things and people, between intelligent algorithms and the minds involved in communication?[95] I argue instead that these developments require a shift from references to intelligence to references to communication. What algorithms are reproducing is not the intelligence of people but the informativity of communication. When new forms of communication combine the performances of algorithms with the performances of people, algorithms are not confused with people, nor do they become intelligent. The difference between the operations of algorithms and human thought gives rise to new ways of dealing with data and producing information in the circuit of communication.

The following chapters test this claim by describing and analyzing various cases of communication with algorithms—each under different conditions and with very diverse consequences.

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

This is a section of doi:10.7551/mitpress/14189.001.0001

Artificial Communication

How Algorithms Produce Social Intelligence

By: Elena Esposito

Citation:

Artificial Communication: How Algorithms Produce Social Intelligence
By: Elena Esposito
DOI: 10.7551/mitpress/14189.001.0001
ISBN (electronic): 9780262368865
Publisher: The MIT Press
Published: 2022

The open access edition of this book was made possible by generous funding and support from the MIT Libraries.

The MIT Press

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024

MIT Press Direct

© 2022 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license.

Subject to such license, all rights are reserved.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Avenir by Jen Jackowitz. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Esposito, Elena, author.
Title: Artificial communication : how algorithms produce social intelligence / Elena Esposito.
Description: Cambridge, Massachusetts : The MIT Press, [2022] | Series: Strong ideas series | Includes bibliographical references and index.
Identifiers: LCCN 2021013271 | ISBN 9780262046664 (hardcover)
Subjects: LCSH: Telecommunication—Social aspects. | Artificial intelligence—Social aspects. | Online identities. | Social intelligence.
Classification: LCC HM851 .E765 2022 | DDC 303.48/33—dc23
LC record available at https://lccn.loc.gov/2021013271

10 9 8 7 6 5 4 3 2 1

Downloaded from http://direct.mit.edu/books/oa-monograph/chapter-pdf/2234767/c000400_9780262368865.pdf by guest on 20 May 2024