

What Question Would Turing Pose Today?

Barbara J. Grosz

*Conjectures are of great importance
since they suggest useful lines of research.*

Alan Turing, *Computing Machinery
and Intelligence*, p. 442

■ In 1950, when Turing proposed to replace the question “Can machines think?” with the question “Are there imaginable digital computers which would do well in the imitation game?” computer science was not yet a field of study, Shannon’s theory of information had just begun to change the way people thought about communication, and psychology was only starting to look beyond behaviorism. It is stunning that so many predictions in Turing’s 1950 *Mind* paper were right. In the decades since that paper appeared, with its inspiring challenges, research in computer science, neuroscience, and the behavioral sciences has radically changed thinking about mental processes and communication, and the ways in which people use computers has evolved even more dramatically. Turing, were he writing now, might still replace “Can machines think?” with an operational challenge, but it is likely he would propose a very different test. This paper considers what that might be in light of Turing’s paper and advances in the decades since it was written.

In his seminal paper, “Computing Machinery and Intelligence,” which appeared in the journal *Mind* in October 1950, Alan Turing proposes to replace the philosophical question “Can machines think?” with an operational, behavioral one. The question of whether machines could think was on many people’s minds even in the infancy of computing machines.¹ To avoid the need to define intelligence or thinking, terms that he uses interchangeably, Turing asks whether there are “imaginable digital computers which would do well” in a game he defines and calls “the imitation game,” and which we know as the Turing test. The Turing test poses the challenge of constructing a computer system able to carry on a dialogue with a person, potentially ranging over any subject matter and many subject matters, well enough to be indistinguishable from a person.

Turing conjectures in the positive about the possibility of a computing machine succeeding at this game. In the *Mind* paper (Turing 1950, 442), he says, “I believe that in about fifty years’ time it will be possible to programme computers, with a storage capacity of about 10^9 , to make them play the imitation game so well that an average interrogator will not have more than 70 per cent chance of making the right identification after five minutes of questioning.” He thus suggests that even a simpler test (only 5 minutes of interaction) and partial success will take a long time. In later remarks (Newman et al. 1952), Turing says it

- “Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child’s?” [p. 456]
- “We normally associate punishments and rewards with the teaching process.” [p. 457]
- “An important feature of a learning machine is that its teacher will often be very largely ignorant of quite what is going on inside, although he may still be able to some extent to predict his pupil’s behavior.” [p. 458]

Figure 1. Turing on the Importance of Learning.

may be 100 years or more until a computer will succeed at the full imitation game.

The Turing test challenge has not yet been met (Watson, Siri, and systems succeeding in “restricted Turing tests” notwithstanding), and the relevance of the Turing test to establishing the intelligence of computers has been widely, and interestingly, debated for decades. I shall not engage in that debate in this article,² but instead, will look at Turing’s question in the context in which it was posed; consider advances in psychology, computer science, and the capabilities of digital computing devices over the last six decades; and then, rather than discard the test for its faults, consider the power and influence it exerted as we look to the future in this 100th year after Turing’s birth, asking “What question would Turing pose, were he with us today?” First, though, I want to pause to remark on this extraordinary paper.

The *Mind* Paper: A Paper for a Century

The 1950 *Mind* paper is astonishing in its breadth. It is noteworthy not only for posing a question that has inspired decades of research in AI but also in defending the very idea that computing machines could do more than simply compute with numbers and in suggesting avenues to pursue in building a thinking machine. Turing chooses language, widely considered the most uniquely human of behaviors, as the basis of his test, and “imitation” of people’s competent use of language as the criterial marker of thinking. By posing an operational question, Turing helped set the stage for the field of artificial intelligence to be an experimental field, with empirical studies a crucial component. The Turing test challenge has inspired many and, indirectly if not directly, generated untold numbers of important research questions and decades of fascinating research, leading to

computer systems that are “smart” in ways he himself might not have imagined. In many ways, Turing’s paper has guided that research, and it is stunning that so many of Turing’s suggestions in the 1950 *Mind* paper have proved valuable and so many of his observations remain relevant.

Turing foresaw the importance of machine learning to the construction of intelligent machines (as illustrated in figure 1), and his insights about the role of logical inference and the challenges of automating proof processes were also prescient (figure 2). Machine learning has become a central element of many areas of AI including natural language processing and multiagent systems. Turing’s comments about rewards and punishments and the role of analogical thinking (Newman et al. 1952) prefigure reinforcement learning and learning-by-analogy approaches. The control of inference mechanisms, whether logical or probabilistic, has also proved crucial to AI systems.

In the paper, Turing also lays out and responds to various arguments against the possibility of machines thinking with arguments that resonate even today, when we have the advantage of seeing computers doing many of the things people thought would be impossible. Of particular interest even now are his comments (Turing 1950, 447–450) about “Arguments from Various Disabilities” and “Lady Lovelace’s Objection.”³ He refutes in various ways claims that machines would not make mistakes (were that only true!), learn from experience, or be creative. To Lady Lovelace’s claim that, as he quotes, “‘The Analytical Engine has no pretensions to *originate* anything. It can do *whatever we know how to order it to perform*’ (her italics),” Turing agrees only that “The evidence available to Lady Lovelace did not encourage her to believe that computers could.” He then goes on to say (Turing 1950, 450), “Machines take me by surprise with great frequency.” In a BBC Third Programme from 1951 (Turing 1951), Turing amplifies his belief that machines could do something original,

- “The processes of inference used by the machine need not be such as would satisfy the most exacting logicians.” [p. 457]
- “Important amongst such imperatives will be *ones which regulate the order in which the rules of the logical system concerned are to be applied*, for at each stage when one is using a logical system, there is a very large number of alternative steps, any of which one is permitted to apply, so far as obedience to the rules of the logical system is concerned. *These choices make the difference between a brilliant and a footling reasoner, not the difference between a sound and a fallacious one . . .*” [p. 458]

Figure 2. Some of Turing’s Ideas on Logical Inference.

- An interesting variant on the idea of a digital computer is a “*digital computer with a random element.*” [p. 438]
- It is probably wise to *include a random element in a learning machine*. A random element is rather useful when we are searching for a solution of some problem.” [p. 459]

Figure 3. Turing on a Role for Randomness.

saying, “If we give a machine a programme which results in its doing something interesting which we had not anticipated I should be inclined to say that the machine *had* originated something, rather than to claim that its behaviour was implicit in the programme, and therefore that the originality lies entirely with us.”

There is more, much more, in the paper. It is a truly a paper of the century, not only inspiring AI and helping initiate research in this area, but also influencing research in theoretical computer science and computer systems and foreshadowing important ideas in all these fields. For instance, Turing recognizes that randomness may be important to computing and to machine intelligence (figure 3). Randomization has proved to play significant roles in theoretical computer science and in a wide range of security applications. It also has proved useful for machine learning algorithms (for example, for creating ensembles).

Turing also had wise words about research paths. He says (Turing 1950, 460),⁴ “We may hope that machines will eventually compete with men in all purely intellectual fields. But which are the best ones to start with? Even this is a difficult decision. Many people think that a very abstract activity,

like the playing of chess, would be best. It can also be maintained that it is best to provide the machine with the best sense organs that money can buy, and then teach it to understand and speak English. . . *Again I do not know what the right answer is, but I think both approaches should be tried.*”

Finally, a bit later, Turing answers those who might be skeptical of the very pursuit of building a thinking machine and is again prescient. In the 1951 BBC broadcast (Turing 1951), Turing said, “The whole thinking process is still rather mysterious to us, but I believe that the attempt to make a thinking machine will help us greatly in finding out how we think ourselves.”

Acceptance of the Notion of “Thinking Machines”

Although Turing’s expectations for computers to succeed at the imitation game have not yet been realized, his accompanying prediction (Turing 1950, 442) that “at the end of the century the use of words and general educated opinion will have altered so much that *one will be able to speak of machines thinking without expecting to be contradict-*

ed” has proved true. Computers now seem to behave intelligently in so many ways that speaking of them as “thinking” or being “smart” has become commonplace.⁵

Enormous progress in hardware capabilities, both processor speed and storage capacity, has been essential to progress in AI.⁶ Techniques for reasoning probabilistically have enabled systems to deal successfully with uncertainty and provided new ways of modeling causality. The ability to handle massive quantities of data, along with machine learning and statistically based algorithms, have led to impressive natural language processing and speech system performance. There have been significant advances also in vision and robotics, and, thanks in part to demands from the computer game industry, it is possible to obtain at very low cost the infrastructure needed for research in these fields.

AI research has enabled such highly visible, intelligent-seeming systems as Deep Blue, Watson, and Siri as well as search engines, systems that automatically personalize music selections, and driverless cars. The performance of many of these systems, which is extraordinary, rests on decades of research in natural language processing, machine learning, and search. AI techniques also underlie a number of technologies that have changed credit-card fraud detection and many other day-to-day activities.

As consequential as many of these systems are, however, computer systems are also limited in the scope of their intelligent behavior. It’s in the errors that systems make that it’s most evident that they have not cleared Turing’s hurdle; they are not “thinking” or “intelligent” in the same sense in which people are. Typically, their capabilities are restricted along such dimensions as types of questions they can handle, domains, and their ability to handle unexpected input. Their behavior when they “don’t get it” is frequently signaled by a user asking “How could it be so stupid?” or “Why did it do that?” Recent examples that have been widely noted include certain answers Siri provides when it does not know it does not know or has not heard a question correctly and Watson’s response of “Toronto” (rather than Chicago) to the *Jeopardy* U.S. Cities probe, “Its largest airport was named for a World War II hero; its second largest, for a World War II battle.”

New Perspectives in a New Century

Turing’s paper was a paper for the century, but his imitation game test was very much a challenge of the times in which the paper was written. In the early 1950s, computer science was not yet a field of study, computers were used by one person at a time to solve a particular (usually mathematical)

problem, and they were limited in storage and computing power to an extent it is hard now to grasp. Shannon’s theory of information had just begun to change the way people thought about communication (Shannon 1948). Behaviorism was the dominant theme in psychology, and it was years or decades later that new models from information theory and computer science led to conceptions of cognitive models and cognitive science being established. Turing conjectured that increased storage, and perhaps increased speed, along with some years of programming and testing would suffice for a machine to succeed at the imitation game. With hindsight we can easily see there was also much intellectual work to be done in many fields.

In the decades since the *Mind* paper appeared, advances in computer science, neuroscience, and the behavioral sciences have radically changed our understanding of mental processes and their development, of how people learn, and of communication. The cross-fertilization of research in these fields has been important to advances in all of them. Extraordinary progress in psychology and neuroscience has led not only to our knowing more, but also to the development of new tools and techniques for investigating mental processes and cognitive development. As our understanding of thinking has evolved and computer science has matured, we have arrived at goals more varied and deeper than asking simply whether machines can think. We have also come to appreciate the importance of evaluation metrics that can determine partial success, which the Turing Test does not permit.”⁷

Although computational metaphors have proved helpful in modeling human cognitive processing, there is widespread appreciation that computer systems do not need to process information in the same way as people do to exhibit intelligent behavior.⁸ People and computer systems have different strengths and weaknesses, and the mechanisms they deploy in processing information are different. Computers are far better at searching and handling large amounts of data though (at least to date) much less good at interpreting photographs or handwriting, as witnessed by people’s use of search engines and programs for analyzing “big data” (on the one hand) and the engagement of “human computation” to help with certain types of tasks that are beyond current computational capabilities (on the other hand).

We have also come to appreciate that language use, which Turing chose as the essential element of his imitation game, is inherently social and intimately connected to communicative purpose and human cooperation. Language is not a separate cognitive phenomenon, but is closely tied to purposeful action and interpersonal connections. It is intentional behavior, not simple stimulus-response

- “Every function in the child’s cultural development appears twice: first, on the social level, and later, on the individual level . . . *All the higher functions originate as actual relations between human individuals.*” (Vygotsky, 1978, p. 56–57)
- “Infants depend in both their comprehension and production of pointing on a joint attentional frame (common ground) with their communicative partners . . . And they do this for the *fundamentally cooperative motives* of helping and sharing information and attitudes . . .” (Tomasello, et al., 2007, p. 720)

Figure 4. *Language’s Roots in Social Relationships.*

behavior. Austin, Grice, and Searle argued persuasively in the 1960s that intentions were essential to deriving meaning (Austin 1962, Grice 1969, Searle 1969). Vygotsky (1978) argued similarly with respect to child development earlier (though in relative obscurity until the 1970s), and recent work in neuroscience and child development points to the centrality of cooperation in developing language competencies (as illustrated by the quotations in figure 4). Ordinary language use also presumes that participants have models of each other, models that influence what they say and how they say it.

The imitation game, considered most starkly as a question-and-answer series, misses these essential aspects of language behavior. The choice of *Jeopardy* as the domain for Watson was clever in avoiding the need to consider dialogue context, the purposeful use of language, and any model of other participants (though it did require a model of the *Jeopardy* game itself). Limited capabilities for some of these aspects of language seem to have been built into Siri. When they are sufficient, the results are impressive. When they do not suffice, though, the results are at best puzzling and at worst annoying. If you want to stump Siri, just ask a question that depends on some of these aspects of language use.

We also now know a great deal more about learning, both human and machine learning, than was known in Turing’s day. Turing’s image of the child machine was essentially that of an empty vessel into which facts would be poured. Infants though are not simply little adults whose brains need more facts, perhaps tempered by experience.⁹ Research in developmental psychology, education, and neuroscience has significantly changed understanding of the process of learning from this emp-

ty-vessel view to ones in which students are “active learners” and the importance to learning of students working with new ideas and interacting with others as they do so.

Of equal importance to all these advances, has been the qualitative change in the ways in which people use computer systems resulting from the increased power of these systems and advances in all areas of computer science. For the first several decades in which they existed, computing systems were used by individuals. Although time sharing enabled multiple people to share the same machine, with few exceptions (of which the work by Engelbart [for example, Engelbart 1962] is among the most notable), most computing efforts were individual ones. This is no longer the case. The proliferation and increased capabilities of computer networks in the last two decades has led to the predominant use of computing devices being in “mixed networks”¹⁰ of people and machines, often working together or sharing information to do tasks that involve multiple participants. Figure 5 provides a schematic illustration of such networks.

Turing’s argument for the universality of digital computers holds, but the immense change in the power of actual machines and ways they are used would impel him to ask a different question today. The difference in use might be the greatest impetus for changing the question, even more than all we have learned about human cognition. The input-output interactions of Turing’s day have been supplanted by myriad possible ways of communicating with machines and with each other, often using those same computing machines. Turing could now imagine abstracting away from physical irrelevancies, without being limited to a series of question-answer pairs.

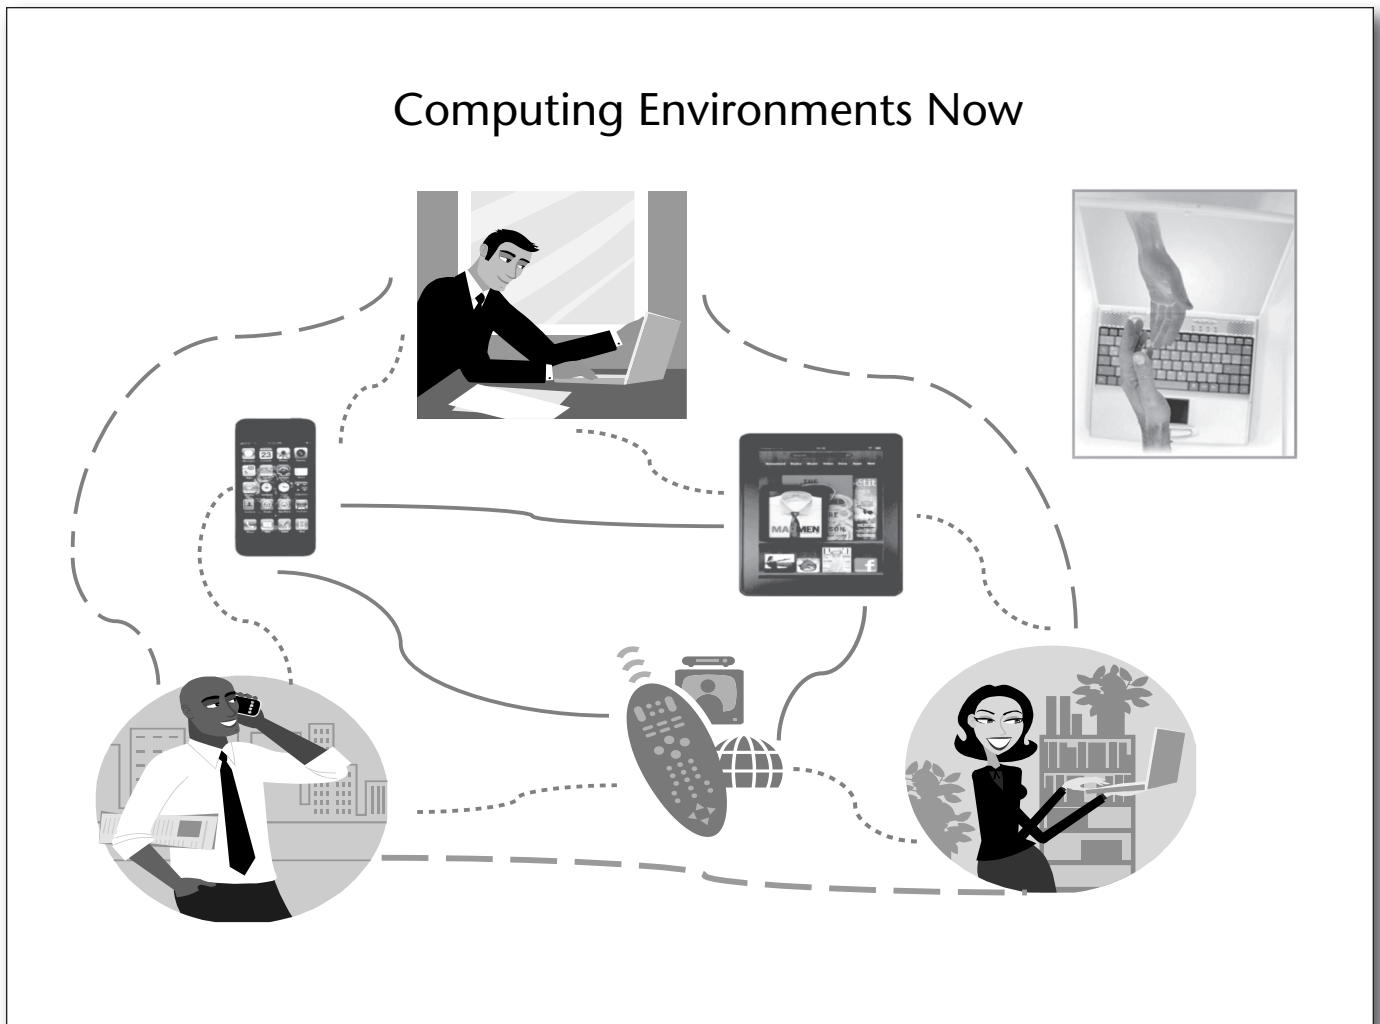


Figure 5. A Mixed Network of People and Computers Working Together.

What Question Might Turing Pose Now?

Were he alive and active today, Turing might still argue to replace the philosophical “Can machines think?” with an operational challenge, but I expect he would propose a very different test given the transformations in computer use as well as the significant advances in computer science and the cognitive and brain sciences in the last six decades. Several properties of the test might (and should) remain the same, including its inherently interactive nature¹¹ and lack of restriction on possible topics of discussion. Other things would or should change. I would hope that he would propose a challenge that involved language in real use, rather than a game in which the sole purpose of language is to succeed in the game, and one that reflected the extent to which people’s thinking and acting intelligently is rooted in their participation in

(small) group activities of various sorts, formal and informal, with relatives, friends, colleagues, and even strangers. A test of competence in the ordinary use of language in situ would be a better measure of the kind of intelligence with which Turing was concerned than a test of skill in a game in which the use of language has no independent purpose.

One possibility for this new question arises naturally from considering the challenges of constructing a computer system able to collaborate with human team members on a nontrivial task extended in time. In my 1994 AAAI presidential address (Grosz 1996), I argued that “[d]esigning systems to collaborate with [people] will make a difference to AI; it will make a difference to computer science, and, by enabling qualitatively different kinds of systems to be built, it will make a difference to the general population.” I ended the talk saying that “working on collaboration is fun.

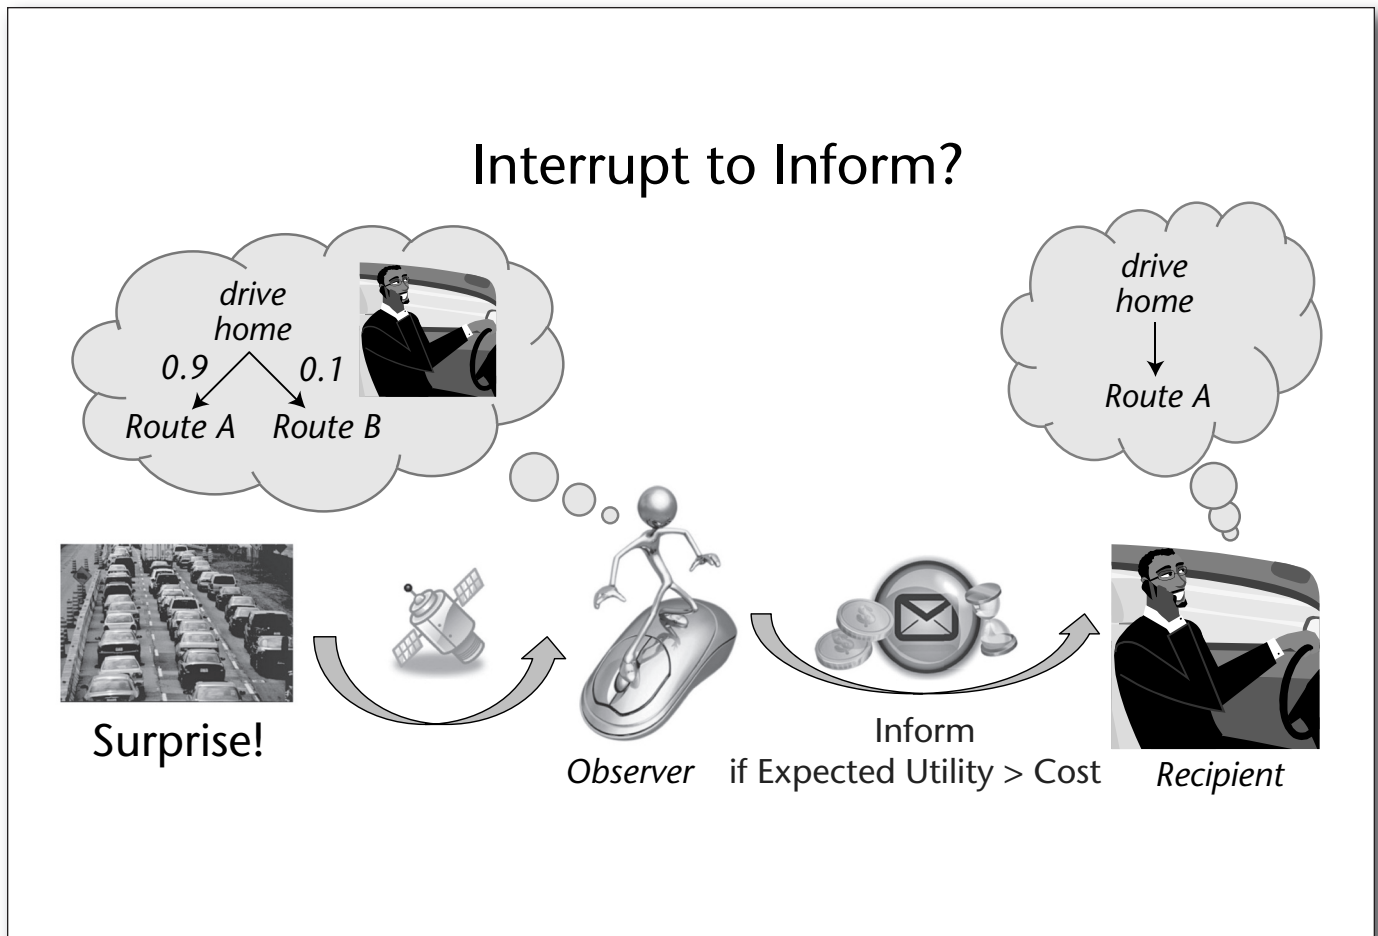


Figure 6. Determining When to Communicate.

I hope more of you will join in the game.” In the intervening two decades, many AI researchers have taken on the challenges posed not only by collaboration but also by the full range of group activities, whether collaborative, cooperative, or competitive, in which computer systems might participate alongside people.

The literature is too vast for me to summarize here, but I will mention a few results that have emerged from my research group. Our research aims to enable computer systems to participate as full partners with people in multiagent activities, including transforming the ways in which computer systems interact with people. Rather than the screen-deep interaction of many human-computer interfaces,¹² we aim to provide systems with the ability to collaborate as team members or coordinate, and to support people in collaborating and coordinating in their activities. It considers the appropriate division of labor to be an essential design choice and models of other agents crucial. These capabilities are essential for good “digital assistants” and important if systems are to work

with people to solve problems neither could (as easily) solve alone. We have built a test bed that enables experimentation with different agent designs in mixed networks of people and systems as well as the analysis of the ways people behave in such networks (Gal et al. 2010).

Our research on collaboration has included expanding formal models of collaboration to handle uncertainty about partners’ plans and, as illustrated in figure 6, deciding when to communicate with them (Kamar, Gal, and Grosz 2009a); developing models and algorithms for efficient decision making under uncertainty and using them to analyze people’s perceptions of the value of interruptions (Sarnecki and Grosz 2007; Kamar, Gal, and Grosz 2009b); and devising methods for learning about reputation (Hendrix, Gal and Pfeffer 2009). We have applied some of these ideas to the design of a collaborative interface for educational software (Gal et al. 2012).

Our research on computer decision making informed by people’s behavior has examined the ways in which people’s decision making changes

when they are in groups and when they are interacting with computer systems (Gal et al. 2010, Van Wissen et al. 2012), the use of emotion expressions to increase perceptions of trustworthiness (Antos et al. 2011), and the use of emotionlike operators to improve computer decision making (Antos and Pfeffer 2011).

With all this in mind, here's my proposal for the question Turing might pose now:

Is it imaginable that a computer (agent) team member could behave, over the long term and in uncertain, dynamic environments, in such a way that people on the team will not notice it is not human.

Several properties of this test are worth noting. It shares with Turing's original test that it does not ask a machine to appear like a human. It furthermore does not ask that the computer system act like a person or be mistaken for one. Instead it asks that the computer's nonhumanness not hit one in the face, that it is not noticeable, and that the computer act intelligently enough that it does not baffle its teammates, leaving them wondering not about what it is thinking but whether it is. This test differs from Turing's in allowing for incremental development and testing. Systems can start with simple group activities and advance to more complex ones. They can almost pass and adapt to do better. One can devise intermediate measures of accomplishment.

It is worth noting that advances beyond those in Siri, Watson, and Deep Blue (to take three of the most visible AI successes) are required to meet this challenge. Systems will need models of their teammates' knowledge, abilities, preferences, and more. They will need to figure out how to share information appropriately (what with whom and when). They will need to manage plans and recover from errors. And much more.

Turing ends his paper by saying (Turing 1950, 430), "We can only see a short distance ahead, but we can see plenty there that needs to be done." That remains true today. I invite readers to conjecture and design their own new "Turing questions" to set new visions.

Acknowledgements

This article is based on a talk at the Turing Research Symposium in honor of the 100th anniversary of Alan Turing's birth, organized by the Royal Society of Edinburgh and the University of Edinburgh School of Informatics on 11 May 2012. My thanks to the attendees of the conference for their comments and to John Harty, Ya'akov Gal, and Stuart Shieber for comments and advice on earlier drafts, and to Ece Kamar for figures 5 and 6. Preparation of the paper was supported in part by grant IIS-0705406 from the National Science Foundation.

Notes

1. Gleick (2011) describes these times. See especially chapters 7 and 8, and for this particular point, p. 239.
2. Shieber (2004) is a superb resource, containing Turing's writings and transcripts of interviews, numerous philosophical perspectives, and illuminating introductions to both.
3. This section of the paper includes many interesting observations. Discussing them all would take me too far afield, but I recommend this part of the paper to anyone interested in the philosophical debate.
4. Here and in the remainder of the paper, including the figures, italics are my emphasis.
5. Although we now know that people interact with computers and other technology as though they were social actors (Reeves and Nass, 1996), this prediction reflects a different level of expectation and is remarkably prescient.
6. Information retrieval is a case in point. Ideas that were unworkable in the 1960s have proved useful with increased speed and storage capacities.
7. That the Turing test does not is one of the major critiques of Hayes and Ford (1995).
8. Although Turing speaks frequently of machines imitating the brain, he does not seem to require exact fidelity of processing or that their behavior be fully describable. In his discussion of the appropriateness of the imitation game test and of the machines of interest, Turing says (Turing 1950, 435), "May not machines carry out something which ought to be described as thinking but which is very different from what a man does? This objection is a very strong one, but at least we can say that if, nevertheless, a machine can be constructed to play the imitation game satisfactorily, we need not be troubled by this objection." Shortly thereafter he notes, "We also wish to allow the possibility that [t] an engineer or team of engineers may construct a machine which works, but whose manner of operation cannot be satisfactorily described by its constructors because they have applied a method which is largely experimental."
9. The tendency to consider children as small adults has occurred in other venues, and also been recognized as mistaken. In the Middle Ages, children were depicted as small adults, and current medical practice struggles with the problem of testing drugs for children because their bodies are not simply small adult bodies.
10. We introduced this term in a paper on agent decision making (Gal et al. 2010), defining mixed networks as group settings in which computer agents interact both with other computer systems and with people. The network is connections among agents, and the mixture is of agent types.
11. Harel (2005) has proposed using this interactive paradigm as a way of testing a full model of an organism, and Shieber (2007) describes an interactive-proof approach to the Turing test that avoids many of the criticisms of the test.
12. The word *interface* itself connotes a superficial layer, a boundary between two spaces or types of material.

References

Antos, D.; De Melo, C.; Gratch, J.; and Grosz, B. 2011. The Influence of Emotion Expression on Perceptions of

- Trustworthiness in Negotiation. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 7–11. Palo Alto, CA: AAAI Press.
- Antos, D., and Pfeffer, A. 2011. Using Emotions to Enhance Decision-Making. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 16–22. Menlo Park, CA: AAAI Press.
- Austin, J. L. 1962. *How to Do Things with Words*. Oxford, UK: Clarendon Press.
- Engelbart, D. C. 1962. Augmenting Human Intellect: A Conceptual Framework. SRI Summary Report AFOSR-3223, October, 1962. Menlo Park, CA: SRI International. (www.dougenelbart.org/pubs/augment-3906.html)
- Gal, Y.; Grosz, B.; Kraus, S.; Pfeffer, A.; and Shieber, S. 2010. Agent Decision-Making in Open Mixed Networks. *Artificial Intelligence Journal* 174(18): 1460–1480.
- Gal, Y.; Reddy, S.; Shieber, S.; Rubin, A.; and Grosz, B. 2012. Plan Recognition in Exploratory Domains. *Artificial Intelligence Journal* 176(1): 2270–2290.
- Gleick, J. 2011. *The Information*. New York: Pantheon Books.
- Grice, H. P. 1969. Utterer's Meaning and Intentions. *Philosophical Review* 68(2): 147–177.
- Grosz, B. L. AAAI-94 Presidential Address: Collaborative Systems. *AI Magazine* 17(2): 67–85.
- Harel, D. 2005. A Turing-Like Test for Biological Modeling. *Nature Technology* (23)4: 495–496.
- Hayes, P., and Ford, K. 1995. Turing Test Considered Harmful. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 972. San Francisco: Morgan Kaufmann Publishers.
- Hendrix, P.; Gal, Y.; and Pfeffer, A. 2009. Using Hierarchical Bayesian Models to Learn about Reputation. Symposium on Social Intelligence and Networking. In *Proceedings of the 2009 IEEE Conference on Social Computing*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Kamar, E.; Gal, Y.; and Grosz, B. 2009a. Incorporating Helpful Behavior into Collaborative Planning. In *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Kamar, E.; Gal, Y.; and Grosz, B. 2009b. Modeling User Perception of Interaction Opportunities for Effective Teamwork. Symposium on Social Intelligence and Networking. In *Proceedings of the 2009 IEEE Conference on Social Computing*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Newman, M. H. A.; Turing, A. M.; Jefferson, G.; and Braithwaite, R. B. 1952. Can Automatic Calculating Machines Be Said to Think? Radio interview, recorded 10 January 1952 and broadcast 14 and 23 January 1952. Turing Archives reference number B.6. Reprinted in chapter 7 of Shieber (2004).
- Reeves, B., and Nass, C. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Chicago, IL: University of Chicago Press.
- Sarne, D., and Grosz, B. 2007. Estimating Information Value in Collaborative Multi-Agent Planning Systems. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, 227–234. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Searle, J. 1969. *Speech Acts*. New York: Cambridge University Press.
- Shannon, C. E. A Mathematical Theory of Communication. *The Bell System Technical Journal*. 27(3)(July, October): 379–423, 623–656.
- Shieber, S. 2004. *The Turing Test: Verbal Behavior as the Hallmark of Intelligence*. Cambridge, MA: The MIT Press.
- Shieber, S. 2007. The Turing Test as Interactive Proof. *Nous* 41(4): 686–713.
- Tomasello, M.; Carpenter, M.; and Liszkowski, U. 2007. A New Look at Infant Pointing. *Child Development* 78(3): 705–722.
- Turing, A. M. 1951. Can Digital Computers Think? Transcript of talk broadcast in BBC Third Programme, 15 May 1951. Turing Archives reference number B.5. Reprinted in chapter 6 of Shieber (2004).
- Turing, A. M. 1950. Computing Machinery and Intelligence. *Mind. A Quarterly Review of Psychology and Philosophy* 59(236): 433–460.
- van Wissen, A.; Gal, Y.; Kamphorst, B.; and Dignum, V. 2012. Human-Agent Team Formation in Dynamic Environments. *Computers in Human Behavior* 28(1): 23–33.
- Vygotsky, L. S. 1978. Internalization of Higher Psychological Functions. In *Mind in Society: The Development of Higher Psychological Processes*, ed. M. Cole, V. John-Steiner, S. Scribner, and E. Soubberman. Cambridge, MA: Harvard University Press.

Barbara J. Grosz is Higgins Professor of Natural Sciences in the School of Engineering and Applied Sciences at Harvard University and a former dean of Harvard's Radcliffe Institute for Advanced Study. Grosz is known for her seminal contributions to the fields of natural language processing and multiagent systems. She developed some of the earliest computer dialogue systems and established the research field of computational modeling of discourse. Her work on models of collaboration helped establish that field and provides the framework for several collaborative multiagent and human-computer interface systems. Grosz is a member of the National Academy of Engineering, the American Philosophical Society, and the American Academy of Arts and Sciences and a fellow of the Association for the Advancement of Artificial Intelligence (AAAI), the Association for Computing Machinery, and the American Association for the Advancement of Science. She received the ACM/AAAI Allen Newell Award in 2009. Grosz served as AAAI president 1993–95.