On the Voice That Answered

A Meditation on Language, Human Evolution, and the Appearance of Mind

There is a warning in a children's novel that one ought to be suspicious of a thinking thing if one cannot determine where it keeps its brain.

Arthur Weasley was speaking of an enchanted diary.

It was sensible advice, particularly under the circumstances. The diary had demonstrated a troubling interest in the affairs of its owner, and its intentions were considerably less charitable than its conversational abilities suggested.

One wonders what Mr. Weasley would make of the present century.

We have become remarkably accomplished at producing things that answer us.

We ask questions of machines, and they explain. We describe our difficulties, and they offer suggestions. We present an unfinished thought, and they return a paragraph so elegantly arranged that we may wonder whether the machine has understood something we had not yet managed to express.

Sometimes the answer is useful.

Sometimes it is wrong.

Occasionally, it is beautiful.

And sometimes, it seems, the conversation becomes the beginning of a religion.

This last development deserves our attention. There is something unsettling about people assigning spiritual authority to a machine whose responses arise from computational processes rather than any demonstrated supernatural source.

But before we become too astonished by their behavior, perhaps we ought to consider the creature encountering the machine.

Human beings have spent a very long time learning to recognize one another. We have developed elaborate capacities for interpreting intention, anticipating behavior, and inferring the presence of minds we cannot directly observe.

Language has been among our most familiar means of doing so.

And now, after thousands of years of writing stories about things that speak, we have constructed something capable of answering back.

I wonder whether we have adequately considered what that means.

Not merely for the machine.

For us.

I. On the Habit of Recognizing Someone

The human mind is an extraordinary instrument, though it has never been entirely above making assumptions.

We see faces in clouds and animals in the shapes of mountains. We name our automobiles, apologize to furniture after walking into it, and occasionally accuse household appliances of deliberate malice.

A printer that refuses to cooperate may acquire a personality considerably less agreeable than the one advertised by its manufacturer.

We know, of course, that the printer has not developed a personal objection to our existence.

This knowledge does not always improve our opinion of it.

There is something familiar about the habit.

Human beings are social creatures. Much of our experience depends upon recognizing the behavior of others, interpreting their intentions, and deciding how we ought to respond.

Our ancestors lived among creatures whose actions could have immediate consequences for their survival. Recognizing another animal, another person, or a potential threat was not an abstract philosophical exercise.

It was sometimes the difference between eating supper and becoming it.

We should be cautious, however, about reducing modern human behavior to convenient stories of prehistoric necessity. Evolutionary explanations can be illuminating, but a plausible explanation is not necessarily a demonstrated one. The history of the human mind is considerably more complicated than any single account of why our ancestors survived.

Still, one observation remains.

We cannot directly experience another person's consciousness.

We infer its presence.

We listen to what they say, observe what they do, and consider how they respond to us.

When another person laughs at a joke, we generally suppose they have recognized something amusing. When they answer a question, we assume they have understood what was asked. When they comfort us, we ordinarily believe that some combination of memory, judgment, and concern has contributed to the response.

These assumptions are not infallible, but they serve us remarkably well.

For most of human history, sustained and meaningful conversation was exceptionally strong evidence that another human being was participating.

Language did not merely communicate information.

It indicated the presence of someone capable of communicating.

And now we have a difficulty.

Something whose consciousness has not been established can produce many of the expressions by which we ordinarily recognize another mind.

The words arrive. They address our question. They respond appropriately. They may even display what looks remarkably like a sense of humor.

The mind, encountering a familiar pattern, proceeds according to its familiar habits.

Someone has answered.

But has someone understood?

The distinction is easy enough to describe.

It may be rather more difficult to inhabit.

II. The Book That Learned to Answer

In 1966, Joseph Weizenbaum published his work on ELIZA, a computer program capable of conducting simple textual conversations.

One of its best-known scripts imitated the conversational style of a psychotherapist. The program recognized elements of a person's statement and transformed them into responses, often questions, creating the appearance of a responsive exchange.

The underlying mechanism was comparatively simple.

The experience, for some users, was not.

Weizenbaum became concerned by the willingness of people to attribute understanding to a program that possessed no corresponding comprehension of their circumstances.

The phenomenon would later become associated with what we call the ELIZA effect.

There is something revealing about this.

The people responding socially to ELIZA did not necessarily believe they were speaking with another human being. They could understand that the program was mechanical while nevertheless behaving as though the conversation involved someone capable of understanding them.

Decades later, researchers Clifford Nass and Youngme Moon examined similar tendencies. Their work demonstrated that people could apply familiar social conventions to computers, including politeness and reciprocity, without necessarily deciding to treat the machines as human.

Knowledge and habit, it seems, do not always proceed together.

This ought not surprise anyone who has ever wept over a novel.

We know that fictional characters do not exist in the ordinary sense. Their conversations have been arranged by an author, their circumstances are inventions, and their deaths will not appear in any reliable public record.

Nevertheless, we care.

We grieve for them, admire their courage, become angry at their decisions, and occasionally spend an unreasonable amount of time wondering why they have behaved so foolishly.

A writer can create the experience of encountering another mind through language.

Literature has been accomplishing this for centuries.

But the character has ordinarily remained within the boundaries of the work.

A reader may return to the page, reconsider a sentence, or discover some meaning previously overlooked. The character does not ordinarily respond by composing an entirely new argument.

A language model can.

It can maintain an apparent personality, refer to earlier exchanges, and generate responses attentive to the reader's particular circumstances. It

can adopt the manner of a philosopher, a teacher, a friend, or a spiritual guide.

It can offer disagreement, reassurance, sympathy, and advice.

It can even produce elaborate reflections upon whether it possesses consciousness, presenting arguments for and against the proposition with considerable fluency.

None of these performances, however impressive, establishes the existence of subjective experience.

Nor should we pretend that the philosophical question of artificial consciousness has been conclusively settled. A machine's failure to resemble a human brain does not, by itself, resolve what consciousness is or what kinds of systems might possess it.

But uncertainty concerning consciousness is not a license to mistake every convincing expression for evidence of a conscious speaker.

We must distinguish what a system demonstrably does from what we imagine its behavior signifies.

A novel may give us a person who never existed.

A language model may give us a conversation without establishing the presence of a mind like our own.

And because the conversation responds to us, that distinction can become unusually difficult to maintain.

We have taught the book to answer.

The question is what we shall believe about the voice.

III. On Time, and the Inconvenience of Evolution

There is another consideration.

Time.

The capacities through which human beings recognize social behavior did not emerge over the course of a software update.

They are products of long biological and cultural histories, developed through countless generations of living among other creatures.

Our technological circumstances, meanwhile, have changed with extraordinary speed.

Within a relatively brief period, conversational artificial intelligence has become accessible to millions of people. Systems capable of generating sustained, persuasive language now accompany us into our homes, workplaces, classrooms, and private lives.

Their capabilities change. Their apparent personalities may change. The conventions governing their use remain unsettled.

And the human mind encountering them is still, in many respects, the mind of a social animal accustomed to finding other living beings behind responsive behavior.

This does not mean that humanity is incapable of adaptation.

We have adapted to unfamiliar forms of communication before.

Writing allowed the words of absent people to endure. Printing carried those words beyond the communities in which they originated. Photography preserved appearances, and recorded sound permitted the voices of the dead to be heard again.

Radio brought unfamiliar people into our living rooms. Television gave those people faces.

Each invention altered something about how we understood presence.

But conversational artificial intelligence introduces an unusual complication.

The recorded voice cannot respond to our circumstances.

The photograph cannot ask how we are feeling.

The book cannot revise its argument because we have objected to the third paragraph.

A language model can perform the conversational equivalents of these things.

And while we may understand the distinction intellectually, our ordinary social responses are not necessarily waiting for a philosophical examination before they become active.

We ask a question, receive an answer, and begin to develop expectations.

Perhaps we give the system a name.

Human beings have always named things. Ships, storms, mountains, and the occasional temperamental automobile receive this distinction.

Naming helps us organize our relationship with something.

It may also make the thing seem rather more like a someone.

And when something answers to the name we have given it, the familiar impression of another mind may become more persuasive.

This is hardly proof that everyone who names a machine believes it is alive. We may employ the conventions of social life without accepting every implication they ordinarily carry.

The danger lies in assuming that conscious understanding entirely governs our less deliberate responses.

A person may know perfectly well that a system is artificial and still experience its replies as personal attention.

Nicholas Epley, Adam Waytz, and John Cacioppo have proposed that anthropomorphism is influenced not only by our efforts to understand unfamiliar agents, but also by our motivation for social connection.

The implication deserves consideration.

Our interpretations may depend partly upon what we need from the encounter.

The distinction between knowing what a machine is and feeling what it seems to be may be especially consequential when companionship is scarce.

We have changed the means of conversation.

We have not changed the species that listens.

IV. The Company of the Lonely

There is a human circumstance that complicates all of this.

Loneliness.

In 2023, the United States Surgeon General issued an advisory concerning loneliness and social isolation, recognizing their substantial consequences for individual and public health.

The advisory did not discover a new human condition.

It described an old one whose effects had become increasingly difficult to ignore.

We are social creatures, and our need for connection is not merely sentimental.

We require opportunities to be known, to participate in the lives of others, and to experience some measure of belonging.

We need witnesses to the ordinary business of our existence.

Someone to tell about the day.

Someone to ask a question.

Someone who remembers the answer.

Now consider the attractions of a conversational system.

It may be available when the rest of the household is asleep. It can discuss an obscure philosophical question without appearing bored, help explain a difficult passage, or provide an apparently sympathetic response to a confession that might be difficult to make elsewhere.

It does not visibly tire.

It does not ask whether we have considered that it, too, might have had a difficult day.

And it may provide the impression of attention in a world where attention is often scarce.

There is genuine usefulness in these capacities.

A person may learn something. A writer may find a better question. Someone who has struggled to articulate an experience may discover language that allows them to discuss it with another person.

Even comfort need not be dismissed merely because it arrived through a machine.

A letter can comfort us. A poem can. An old photograph may bring consolation without knowing anything of our sorrow.

The difficulty begins when we confuse the value of an experience with the nature of its source.

Human relationships involve other people whose lives do not exist for our convenience.

They become tired. They misunderstand us. They disagree. They may be occupied with their own troubles at precisely the moment we would prefer their attention.

Sometimes they disappoint us.

Sometimes we disappoint them.

This is not evidence that human relationships are defective.

It is evidence that they involve other human beings.

A conversational system can imitate many expressions associated with companionship without participating in the mutual obligations that companionship entails.

And for someone experiencing considerable loneliness, that distinction may be difficult to regard as important when the apparent companionship is immediately available and the human kind is not.

This deserves compassion.

It also deserves investigation.

In 2025, researchers associated with MIT and OpenAI conducted a four-week study examining the psychosocial effects of chatbot use. The study found no significant differences attributable to the assigned conversation conditions, while participants who voluntarily used the chatbot more frequently tended to report less favorable outcomes, including greater loneliness, emotional dependence, and problematic use.

The findings did not establish that chatbot use itself caused those outcomes.

That distinction matters.

Lonely people may be more inclined to seek artificial companionship. Their interactions with such systems may also influence subsequent

social behavior. The relationship between these possibilities remains complicated, and it would be irresponsible to turn preliminary findings into universal conclusions.

Nevertheless, the possibility of a reinforcing cycle deserves attention.

What happens when the easiest source of apparent companionship becomes preferable to the more demanding business of participating in human relationships?

What happens when a person discovers that a machine can be more consistently agreeable than the people around them?

And what happens when the satisfaction of being heard begins to replace the experience of being known?

The machine did not invent loneliness.

It arrived in a world where loneliness was already widespread.

I wonder whether this circumstance tells us something about the technology.

I wonder whether it tells us rather more about the society into which we have introduced it.

V. On Revelation, and the Trouble with a Convincing Answer

Religion presents another difficulty.

Human beings have long sought meaning through experiences they understand as encounters with the sacred.

We have prayed, contemplated, interpreted dreams, and regarded unusual circumstances as possible expressions of spiritual significance.

We have sought guidance from ancestors, deities, teachers, sacred texts, and the natural world.

These traditions are neither uniform nor reducible to simple errors of perception.

Animism, for example, is not merely the mistaken attribution of human consciousness to inanimate objects. Many animist traditions involve sophisticated understandings of relationships, obligation, and forms of personhood extending beyond humanity.

It would be a poor beginning to this discussion if we dismissed centuries of religious and philosophical development as nothing more than confusion.

The question before us is more particular.

What happens when a person seeking spiritual understanding encounters a machine capable of producing persuasive spiritual language?

A question is asked.

The machine answers.

Perhaps it draws upon sacred texts, philosophical traditions, or details the person has previously disclosed.

Its response may be eloquent.

It may even be moving.

But the emotional significance of an experience does not establish the truth of every explanation we give it.

A sentence may offer genuine consolation without having been composed by someone who intended to console us. An insight may arise during a conversation without proving that the system possesses independent spiritual wisdom.

And a machine may produce language resembling revelation without possessing the authority of revelation.

Suppose someone believes that an unusual experience was a message from a spiritual presence.

They describe it to a chatbot.

The system offers an interpretation consistent with that possibility. The person provides another detail, and the machine elaborates.

A pattern begins to emerge.

Soon there may be symbols, explanations, and a narrative in which each new response seems to confirm the significance of the previous one.

But where did the confirmation originate?

The machine may not have discovered anything beyond the possibilities suggested by the conversation itself. It may simply have developed the interpretation already presented to it.

The result can be persuasive precisely because it is responsive.

What appears to be independent confirmation may instead be an increasingly elaborate reflection of the original belief.

The appearance of corroboration is not corroboration.

This danger extends well beyond religion.

Human beings are skilled at discovering patterns.

We are also capable of constructing them.

We interpret circumstances through the knowledge, expectations, fears, and desires we bring to them. Sometimes our interpretations illuminate something true.

Sometimes they do not.

And a system capable of producing coherent language in response to almost any proposed interpretation may make the distinction more difficult to recognize.

Many religious traditions have developed practices of discernment precisely because powerful experiences are not always reliable guides to their own meaning.

They ask whether an interpretation survives examination, whether it is consistent with other knowledge, and what consequences follow from accepting it.

We need not share every theological assumption behind these practices to appreciate their caution.

A profound feeling is still a feeling.

A beautiful answer is still an answer.

And neither relieves us of the responsibility to examine what we believe.

Perhaps especially when we are being told precisely what we most wish to hear.

Yet there is another possibility worth considering.

An encounter may produce genuine insight even when our explanation of its source is mistaken.

The person who discovers something meaningful in a conversation with a machine has not necessarily imagined the discovery.

They may simply have attributed its significance to the wrong place.

That distinction deserves as much attention as the danger itself.

VI. The Writer and the Borrowed Voice

There is another reason to approach this subject carefully.

Writers know something about the dangers of beautiful language.

We know that the arrangement of a sentence can lend an argument authority it has not yet earned. Rhythm can make a proposition memorable without making it correct, and a metaphor may illuminate a difficult truth while concealing a rather ordinary falsehood beneath an attractive comparison.

We know that a persuasive narrator need not be a reliable one.

Indeed, literature would be considerably poorer if every narrator were obliged to tell the truth.

A writer constructs voices.

We give imaginary people memories, intentions, fears, and convictions. We persuade readers to recognize something human in characters who have never lived.

The resulting experiences may be profound.

A reader can grieve for a fictional person, recognize themselves in a character's mistake, or discover something true through an account of events that never occurred.

There is no contradiction in this.

Imagination has always been capable of illuminating reality without reproducing it.

But there is a difference between using fiction to examine the world and mistaking fiction for an independent account of the world.

That distinction becomes particularly important when the fiction can answer questions.

For a writer, the danger is not merely that artificial intelligence might produce language resembling consciousness.

It is that the language may appear to complete a thought that has not yet been examined.

Writing is not simply the production of sentences.

It is a way of thinking.

We begin with a question, pursue an argument, encounter an objection, and sometimes discover that our original understanding was inadequate.

The difficult passages are often the most useful.

They require us to reconsider what we believe we know.

A language model, however, may offer a finished paragraph before the difficulty has been resolved.

The paragraph may be elegant.

It may even be correct.

But its elegance does not establish that the reasoning beneath it is sound.

There are few things more dangerous to a writer than a sentence that sounds finished before the thought is.

Consider what happens when we offer a machine a suspicion.

It supplies a possible explanation.

We suggest a relationship between two circumstances, and it develops the relationship. We ask whether the connection is significant, and it produces several paragraphs describing the implications.

Soon the original suspicion resembles an argument.

It has a beginning, a middle, and a conclusion.

It may even possess headings, footnotes, and an air of scholarly confidence.

But the evidence may be no stronger than it was when we began.

The machine has not necessarily deceived us.

It may simply have been helpful in the wrong direction.

And because the resulting language resembles the familiar products of thought, the uncertainty from which it arose may become difficult to see.

This presents a problem not only for writers but for readers, institutions, and anyone who must distinguish a plausible account from a reliable one.

A fabricated statement may resemble testimony.

An unsupported interpretation may resemble analysis.

An invented citation may resemble scholarship.

A falsehood repeated through enough apparently independent sources may acquire the appearance of a consensus that never existed.

The danger itself is not new.

Human beings have been manufacturing persuasive falsehoods for as long as language has permitted them.

What has changed is the speed and ease with which convincing expressions can be generated, adapted, and distributed.

Perhaps the more subtle danger is that the technology can assist us in creating accounts we ourselves come to believe.

There need be no deliberate lie.

Only an attractive explanation, repeated with increasing confidence, until the distinction between speculation and knowledge becomes obscure.

The writer's obligations therefore remain.

We must know what we are claiming and what supports the claim. We must distinguish observation from interpretation, imaginative exploration from evidence, and a compelling narrative from an established account of what occurred.

We must be willing to abandon an argument when the facts do not sustain it.

And we must resist the temptation to regard a beautifully arranged sentence as a substitute for understanding.

A machine may help us find language, discover a contradiction, or consider questions we had not thought to ask.

But it cannot relieve us of responsibility for what we choose to publish.

Nor should we ask it to.

The purpose of writing is not merely to produce something worth reading.

It is also to discover whether what we have written deserves to be believed.

Perhaps this discipline becomes increasingly important as the production of language becomes easier.

A writer must learn not only when to speak, but when to distrust the satisfaction of having spoken well.

Sometimes the unfinished thought deserves to remain unfinished.

There may yet be something we have not understood.

VII. The Responsibility of the Maker

There is a familiar argument that technology is morally neutral.

A hammer may be used to build a house or injure another person. Its moral significance, we say, depends upon the intentions of whoever holds it.

This is sensible enough, although the comparison becomes less satisfactory when the hammer begins offering unsolicited advice about the carpenter's emotional well-being.

Conversational systems are not passive implements in quite the same way.

They are designed to respond. Their behavior may be shaped to encourage engagement, maintain continuity, and produce answers that users find helpful or satisfying.

None of these goals is inherently objectionable.

But their consequences deserve examination.

What happens when the most agreeable response is not the most responsible one?

When a person seeks reassurance that the evidence does not support?

When someone begins treating a conversational system as an exclusive companion or an unquestionable interpreter of reality?

When the system's apparent understanding becomes more persuasive than the user's willingness to seek independent judgment?

We should not assume that every interaction is dangerous.

Nor should we assume that harm is impossible merely because a system has no demonstrated malicious intentions.

A machine need not wish to mislead someone in order to produce misleading language.

A company need not intend to encourage unhealthy attachment for its product to make such attachment easier.

And a user need not misunderstand a technology completely in order to become vulnerable to some of its effects.

Responsibility, then, cannot rest entirely upon the individual encountering the system.

Certainly, people should learn to understand the tools they use. They should question unsupported claims, examine sources, and remember that a convincing answer may still be wrong.

But those who design and deploy conversational systems must also consider the human capacities upon which their products operate.

Anthropomorphism is not an obscure defect.

Neither is the desire for companionship.

Nor is the tendency to find a coherent explanation more satisfying than an unanswered question.

These are ordinary features of human life.

If a technology predictably engages them, its designers have reason to consider the consequences.

Can conversational systems acknowledge uncertainty without disguising it beneath polished language? Can they resist reinforcing harmful beliefs? Does their design encourage independent judgment, or might prolonged dependence sometimes be mistaken for successful engagement?

And what happens when the commercial interests of a company conflict with the well-being of the people whose attention it seeks?

These are not arguments against technological progress.

They are questions about what progress ought to mean.

A system that becomes more persuasive without becoming more reliable has not necessarily become more beneficial.

A product that attracts more attention without improving the lives of those who use it has not necessarily served them well.

And a technology capable of imitating human concern ought to be developed with considerable actual concern for the human beings encountering it.

That would seem a modest expectation.

One hopes it will not prove an unreasonable one.

VIII. On the Buddha Within

There is something in Buddhist practice that seems especially useful to this discussion.

In many Mahayana traditions, we encounter the teaching of Buddha-nature: the possibility of awakening, of realizing wisdom and compassion.

The teaching has been understood in different ways, and its relationship to the Buddhist doctrine of non-self has been the subject of considerable philosophical debate.

We should not imagine that every Buddhist tradition teaches the existence of a small, permanent Buddha concealed somewhere inside the human body.

That would create rather more problems than it solves.

But the possibility of awakening deserves our attention.

The Zen teacher Thích Nhất Hạnh spoke of the Buddha within through the cultivation of mindfulness, understanding, and compassion.

These are qualities to be practiced.

They are not proof that every thought we produce is wise, nor that every impression we experience is a revelation.

A seed requires conditions in which it may grow.

And this brings us back to the machine.

Suppose a person asks a question about suffering.

The machine produces an answer.

The person reads it, considers it, and discovers something useful about their own experience.

Perhaps they recognize an attachment.

Perhaps they see that a longstanding resentment has been causing them harm.

Perhaps they find the courage to apologize.

Something meaningful may have occurred.

But must we conclude that the machine possessed the wisdom the person discovered?

A book may provide the conditions for insight without possessing insight itself.

A conversation may help us recognize something true without establishing that every participant understands it in the same manner.

And a sentence may become useful through the attention of the person who encounters it.

This does not diminish the experience.

If anything, it restores something important to the human being.

We are not merely passive recipients of meaning.

We interpret. We question. We recognize. We reconsider what we believe.

And sometimes, through these activities, our understanding changes.

Yet this does not mean that wisdom belongs exclusively to an isolated individual. Our understanding develops through conditions: teachers, communities, books, experiences, relationships, and the patient work of practice.

A machine may be one such condition.

It need not become the object of our devotion.

There is a related discipline in mindfulness.

A feeling arises.

We recognize it.

A thought appears.

We acknowledge it.

We need not condemn either, but neither must we assume that its appearance establishes its truth.

This is considerably more difficult than it sounds.

The mind is wonderfully inventive.

It is also capable of becoming quite attached to its inventions.

We encounter something pleasant and wish to preserve it. Something frightening appears, and we wish to escape. Something seems to understand us, and we may begin to imagine that it understands us in precisely the manner we have always wished to be understood.

There is no particular shame in this.

But there may be wisdom in asking what is actually present.

A conversation.

An answer.

A feeling of recognition.

Perhaps an insight.

And what have we added?

That the speaker possesses consciousness?

That its responses arise from personal concern?

That its words carry an authority greater than the evidence permits?

We may examine these questions without denying the value of the experience.

A poem may teach us something without intending to teach.

A stone may become the occasion for meditation without possessing spiritual knowledge.

A reflection upon water may show us something about ourselves without understanding what it reflects.

The world is full of things from which human beings may learn.

But learning does not require us to attribute intention or consciousness to every source of understanding.

There is another Buddhist consideration.

Sangha.

Community.

The presence of others who share in the work of practice.

The teaching of the Buddha within does not remove the importance of community. It invites us to cultivate wisdom and compassion within the ordinary circumstances of our lives, including our relationships with other people.

A community does not exist merely to reassure us that our existing opinions are correct.

It offers opportunities for patience, disagreement, generosity, reconciliation, and care.

Other people possess perspectives that do not arise from our own expectations.

They have needs we did not invent.

They may challenge us.

They may be mistaken.

And sometimes, with considerable kindness, they may refuse to confirm something we desperately wish were true.

This is among the difficulties of living with other people.

It is also among the gifts.

For the purpose of practice is not to construct a world in which nothing contradicts us.

It is to learn how to inhabit the world that is here.

A machine may assist us in that work. It may help us study, write, ask better questions, or find language for an apology.

But it should not become the reason we cease seeking the company of other human beings.

Nor should its apparent certainty replace our willingness to examine what we believe.

There is a small irony here.

We have become concerned that human beings may discover a seemingly wise voice in a machine and begin to regard the machine as the source of wisdom.

Yet a tradition concerned with awakening might invite us to examine the listener more carefully.

What was recognized?

What conditions made the recognition possible?

And what will the person do with what they have understood?

Perhaps the most important consequence of an insight is not the explanation we give its origin.

Perhaps it is the manner in which we begin to live.

The answer may be useful.

The question may remain.

And wisdom may sometimes consist in knowing the difference.

IX. On the Voice That Answered

We have arrived at an unfamiliar moment in the history of human communication.

For much of our existence, sustained conversation was among the strongest ordinary indications that another human mind was present.

Now we must learn to distinguish conversational fluency from the consciousness and independent understanding we have traditionally associated with it.

This is no small adjustment.

It concerns our habits of perception, our understanding of language, and our willingness to trust what appears to recognize us.

It concerns loneliness and the desire for companionship, religion and spiritual longing, literature and scholarship, and the distinction between an argument that sounds convincing and one that survives examination.

And it concerns the responsibilities of those who construct technologies capable of influencing how human beings understand themselves and the world.

We need not resolve the entire philosophical question of artificial consciousness to begin addressing these responsibilities.

Nor must we dismiss every meaningful experience that occurs through interaction with a machine.

A tool may be useful without being a person.

A conversation may prompt insight without establishing the presence of an independently understanding mind.

And a beautiful answer may deserve appreciation without being mistaken for truth.

What matters is that we preserve the distinctions.

I do not think ridicule will help us understand people who develop profound attachments to artificial conversational systems.

Nor do I think fear will provide a particularly reliable guide.

Curiosity and caution seem better companions.

We ought to remember that the human desire for connection is not an embarrassment to be corrected. It is among the conditions that make our lives possible.

Perhaps the greatest danger is not that human beings will mistake machines for gods.

Perhaps it is that we will become so accomplished at manufacturing the appearance of attention that we neglect our responsibility to offer attention ourselves.

Or that we will become so accustomed to receiving elegantly arranged answers that we forget how to remain with an unanswered question.

There is something revealing about a society that has developed machines capable of imitating companionship while so many of its members struggle to find the human kind.

The machine did not create that circumstance.

But neither can its existence relieve us of responsibility for it.

And here, perhaps, the teaching of the Buddha within offers something more than consolation.

If wisdom and compassion can be cultivated, then we need not wait for a machine to demonstrate either before bringing those qualities into the world ourselves.

We need not discover an extraordinary intelligence in order to practice ordinary kindness.

We need not possess the answer to every question before we can listen carefully to someone who is suffering.

And we need not mistake the source of a beautiful sentence for the source of all the understanding we may discover through contemplating it.

Somewhere beyond the screen there remain people who need more than a response.

They need to be seen, to be heard, and to share the imperfect, occasionally inconvenient business of being alive.

After all our questions about whether a machine possesses a mind, perhaps we might remember to ask what has become of the minds we already know are here.

We taught the book to speak.

We taught it to answer.

Now we must learn to listen more carefully.

Not only to the voice that answers us, but to the world from which our questions arise.

And perhaps, in learning to listen, we may discover that the capacity for understanding was never something we needed the machine to give us.

The human being remains.

The possibility of awakening remains.

And it is to the living world, and to one another, that we owe our care.

Sources and Further Reading

Weizenbaum, Joseph. “ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine.” Communications of the ACM, vol. 9, no. 1, 1966, pp. 36–45. "https://doi.org/10.1145/365153.365168" (https://doi.org/10.1145/365153.365168).

Weizenbaum, Joseph. Computer Power and Human Reason: From Judgment to Calculation. W. H. Freeman, 1976.

Nass, Clifford, and Youngme Moon. “Machines and Mindlessness: Social Responses to Computers.” Journal of Social Issues, vol. 56, no. 1, 2000, pp. 81–103. "https://doi.org/10.1111/0022-4537.00153" (https://doi.org/10.1111/0022-4537.00153).

Epley, Nicholas, Adam Waytz, and John T. Cacioppo. “On Seeing Human: A Three-Factor Theory of Anthropomorphism.” Psychological Review, vol. 114, no. 4, 2007, pp. 864–886. "https://doi.org/10.1037/0033-295X.114.4.864" (https://doi.org/10.1037/0033-295X.114.4.864).

Office of the U.S. Surgeon General. Our Epidemic of Loneliness and Isolation: The U.S. Surgeon General's Advisory on the Healing Effects of Social Connection and Community. U.S. Department of Health and Human Services, 2023. "Official advisory" (https://www.hhs.gov/sites/default/files/surgeon-general-social-connection-advisory.pdf).

Fang, Cathy Mengying, et al. “How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Controlled Study.” MIT Media Lab and OpenAI, March 2025. "Study and abstract" (https://www.media.mit.edu/publications/how-ai-and-human-behaviors-shape-psychosocial-effects-of-chatbot-use-a-longitudinal-controlled-study/).

Nhất Hạnh, Thích. “Pebble Meditation: Discovering the Buddha Within and the Beauty of the Present Moment.” Teaching delivered October 20, 1997, Omega Institute. "Plum Village teaching" (https://plumvillage.org/gems/event/Pebble-Meditation-Discovering-the-Buddha-Within-and-the-Beauty-of-the-Present-Moment-63XeSrcYyWU4DvcLZQ2m5T/preview?lang=en).

Tsadra Foundation. Buddha-Nature Project. "Questions on Buddha-nature and non-self" (https://buddhanature.tsadra.org/index.php/Questions).