



After a Stroke, AI Helped Me Learn to Write Again

Mukul Pandya

The Wharton School, University of Pennsylvania

Every year, fifteen million people around the globe suffer strokes. Mukul Pandya, a lifelong writer and editor, describes how recent developments in artificial intelligence helped him to recover his abilities and sense of self after a debilitating stroke changed his life overnight.

Globally, 15 million people suffer a stroke every year. If they survive, stroke victims often find themselves in a dark place once they realize their physical impairments may be long lasting. The mental anguish that usually follows can be just as debilitating. That was

my state of being after suffering from a stroke in September 2021. I had just retired as the founding editor of Knowledge@Wharton, the Wharton School's management journal, where I was used to being in the midst of some bustling activity around campus.

Overnight, the stroke caused me to lose the use of the left side of my body. I could no longer walk so I had to use a wheelchair. I began to slur my words. Since my hand sometimes threw food around while I was eating, I wore a bib around my neck at mealtimes. I was in shock. I could hardly believe how dramatically my life had changed in a matter of hours.

In her 1997 novel, “The God of Small Things,” Arundhati Roy, the Booker Prize-winning author, writes about how life can change in a day. No one knows the truth of that statement better than a stroke survivor.

What bothered me most was the hellishness of being helpless.

What bothered me most was the hellishness of being helpless. I became overdependent on my family, friends, and caregivers, who were saintly in their kindness, patience, and support. Despite my best efforts at trying to stay positive, I had miserably dark days. The main reason why life felt so bleak was the conviction that three months after my retirement from the Wharton School, my professional life was over.

After more than forty years as an editor and writer, I could neither write nor edit. If I could not be a writer or editor, who was I? It was a catastrophic crisis of identity, which every stroke survivor goes through.

Tap, don't type

I had mostly regained the ability to walk, talk without slurring, and use my hands, although it was still difficult to use a keyboard to write and edit. Relearning these motor skills was a painfully slow pro-

cess, but made possible by new technology, particularly artificial intelligence, or AI.

My first baby steps took the form of using tools from Google and Apple on my laptop and iPhone; the apps would look at the words I had typed and try to anticipate and suggest the next word. For example, if I wrote “I am in the ...” the AI algorithm would ask if the next word ought to be “hospital.” If that was correct, all I had to do was to tap that word rather than type it. This process worked well, I discovered, for text messages and short emails.

Often, if I mistyped a word, the algorithm would underline it and suggest the correct spelling. I could construct short messages to keep in touch with family and friends around the world, although there was a sort of sameness that crept into these texts. Still, it meant that I did not have to wait for anyone else to type emails for me. Although my wife Hema and daughter Tara had kindly done this for me in the first few days after my stroke, the fact that I could do it myself gave me a small measure of freedom. Some of my agency returned. A ray of light broke through the darkness.

The technical term for the AI technology that makes this possible is predictive analytics. As my friend and former Wharton colleague Kartik Hosanagar, author of a wonderful book titled, *A Human's Guide to Machine Intelligence: How Algorithms Are Shaping Our Lives and How We Can Stay in Control*, has helped me understand, “AI increases the accuracy and reduces the cost of making data-driven predictions.”

After more than forty years as an editor and writer, I could neither write nor edit. If I could not be a writer or editor, who was I?

In my experience, this tool got better with use. As the AI algorithm learned which words I liked to use – based on the frequency with which I used them – its predictions improved. Gradually, I was able to write longer email messages, even though each message took an excruciatingly long time to compose.

Many friends, including Kartik, suggested trying out speech-to-text software programs, and I did. As the name implies, these are programs to which you can dictate your messages, and the software turns them into text. This AI technology, called automatic speech recognition or ASR, is similar to the one that many companies use to allow you to speak your preferences during a phone call rather than pressing a number on a touch-tone phone. In addition to recognizing commands given in natural speech, these algorithms are able to convert them into text.

Many people have had positive experiences with this software, but my early efforts ended in disaster. I am not sure if it was because I was slurring my words or if the AI algorithm did not understand my Indian accent, but each text message I got was riddled with errors. This was frustrating; it took longer to retype an error-filled message than it did to compose it slowly but correctly letter-by-letter in the first place. There appeared to be no efficient way to write long-form text. I felt doomed to do double work if I were to depend on this technology. I got tired of it quickly and gave up.

Then I found an intermediate solution. WhatsApp, which Facebook (now Meta) acquired in 2014 for more than \$19 billion, had a recording feature on its mobile app: I could press a button and speak into the app to create a voice message that I could send to family and friends on my contact list. This was effective because I could now send voice messages that were several minutes long. Often these were complex explanations of medical issues I was dealing with, and I did

not have to face the hassle of the AI algorithm miscommunicating what I wanted to say.

WhatsApp's privacy features also meant that I could speak freely about my health. My friend Rohan Murty, founder of Soroco, a U.K.-based startup, says: "Until you said it, I never thought that WhatsApp could help somebody who's gone through a medical condition like this. If I were a product manager, I would have never realized that maybe one day someone will use it like this."

Another advantage was that my communication could be asynchronous. In other words, I could leave messages for friends in different time zones, and they could respond whenever they had the time. This voice technology allowed me to progress beyond short, terse texts and emails, but I still could not write or edit articles.

Just as the frustration was beginning to build again and the darkness threatened to return, unexpectedly, I had a breakthrough. Before my stroke, I had agreed to interview Google's Neil Hoyne about his book, *Converted*, which is about how companies use data to win customers' hearts. I emailed Neil a list of questions I wrote by typing and tapping on my iPhone. He was kind enough to send back his answers as audio messages. I sent those on to my friend Deborah Yao, editor of AI Business, who had them transcribed, and then edited and published the interview.

Someone reading the article in its final form could hardly have imagined how the process had worked. Thanks to kind and compassionate friends, I was able to produce a long article eight months after my stroke. That gave me an immense boost of positive energy. It was therapeutic and helped me keep healing.

The following month, I was able to do a second story about

ransomware and cybersecurity using the same technique, featuring David Lawrence and Kevin Zerrusen, experts from the Risk Assistance Network + Exchange. The glimmer of hope grew brighter.

Deborah, who previously was my colleague at Wharton, told me about the AI software she had used to produce these transcripts. It was made by a company in Los Altos, CA, called Otter.ai. "Have you tried it?" she asked. "It's good." I downloaded it, and that transformed my life.

Speech-to-text on steroids

My use of Otter.ai was initially a bit complicated. Let us say I had to write a 1,500-word article. I would start by hand-writing a short outline of the story, mapping its structure paragraph by paragraph. (If the article was longer, say 3,000 words, I would map out groups of paragraphs.) After that, I used the iPhone's Voice Memos app, which turns the phone into a recorder, to dictate the entire article. I ended up with an audio file that I could upload to the Otter.ai website.

In a few minutes, Otter.ai's algorithms would create an almost perfectly accurate transcript of what I had said and email it to me. I could now copy and paste the transcript into Google Docs, Microsoft Word, or any other word processing program, clean up the text, and have the final draft ready. While I was impressed that the Otter.ai algorithm got most of the text right, what was truly amazing was the speed with which the AI converted the audio file into text. It could turn even a 60-minute interview into an editable transcript in a few minutes.

What made this magic possible? According to my friend Apoorv Saxena, who once worked for Google and now works for Silver Lake, a private equity firm, it was advances in automatic

speech recognition. An influential paper in 2016 titled, "WaveNet: A Generative Model for Raw Audio," radically redefined the way that algorithms turn speech into text. "We have seen next generation speech-to-text being produced in the last three to four years," he says. This is what makes companies such as Otter.ai as effective as they are, thanks to deep learning.

These days, I use a somewhat different process. Otter.ai lets me create my own digital assistant who 'attends' my Zoom or Google Meet meetings. I introduce 'her,' my AI assistant, as a participant in the meeting to my interviewees, asking if they mind if she joins the meeting to take notes. A few minutes after the meeting ends, 'she' emails me a transcript.

I have taken to practicing typing for an hour every day, so that I can edit the transcribed text. It is important to me to use – but not overuse – the AI technology. If I were to use AI to do everything, I would have no incentive to keep working at strengthening my hand and the neural connections between my brain and fingers. It would simply transfer my dependence from humans to digital technology.

While the AI algorithm that Otter.ai has developed is impressive, it isn't perfect. It gets many things right, but occasionally it gets things spectacularly and hilariously wrong.

While the AI algorithm that Otter.ai has developed is impressive, it is not perfect. It gets many things right, but occasionally it gets things spectacularly and hilariously wrong. For example,

I was recently working on a document in which I had to quote my former Wharton colleague Raghu Iyengar. Otter’s transcript turned his last name from Iyengar to “anger” and his first name from Raghu to “Rachel” – getting the name, gender, and nationality wrong. So it still has some ways to go. Fundamentally, though, it has given me a tool with which to resume my writing and editing, and in many ways, to reclaim my identity.

Human-AI collaboration

As I think about the process that has made this transformation possible, I realize that it has to do with structuring human and AI collaboration the right way.

The work begins with a human process (I think of the interview topic, select the right expert, and come up with the questions to ask). Next, I turn over to AI the relatively narrow task of capturing the conversation in audio format and turning it into text. It does

this at a speed that is unimaginable for even the world’s fastest human transcribers. Finally, I take the task back from the AI algorithm so I can edit and eliminate the laughable “Rachel Anger” kinds of errors and complete the work using human expertise.

I focus on doing what I can do better than the AI and leave it to the AI to do what it does best. This human-AI-human workflow process has allowed me to rebuild my professional life.¹ ■

© Copyright Mukul Pandya, 2022

Author Bio



Mukul Pandya is the founding editor in chief and executive director of Knowledge@Wharton, the online research and business journal of the Wharton School. After retiring from K@W, Mr. Pandya was a senior fellow with Wharton Customer Analytics and AI for Business. A four-time award winner for investigative journalism, Mr. Pandya has published articles in *The New York Times*, *The Wall Street Journal*, *the Economist*, *Time*, *The Philadelphia Inquirer*, and more. He has written or coauthored four books.

Endnote

-
1. To read the entire original essay from which this excerpt is taken, go to: <https://www.linkedin.com/feed/update/urn:li:activity:6991399104807346176/>