

Let There Be

For most of the history of computing, the strange thing about talking to a machine was that the machine didn't listen. You typed. You clicked. You filled in forms. The keyboard was a translator between what you meant and what the computer could accept, and most of the friction of using a computer was really the friction of operating that translator. If you've ever watched someone who types poorly try to do knowledge work, you've seen the cost of it.

The quiet news of the last few years is that the translator is going away. You can now say what you want, out loud, in ordinary English, and something happens. Not always the right thing. But something, and often the right thing, and increasingly not just on the screen.

This is stranger than it sounds. Speech was an early technology humans had, and for a long time it was the most powerful one we owned. Writing added to speech. Computing demoted writing, which had in a sense demoted voice, at least for anything you wanted a machine to do. Speech, for all its expressive range, had no path into the digital world except through a keyboard operated by a person with better-than-average fingers. That was the whole problem. We had this extraordinary, high-bandwidth, emotionally laden communication channel, and the computers of the world could hear approximately none of it.

Voice plus an LLM closes that loop. And because an LLM can turn intent into code, and code can turn into actions, and actions can now reach the physical world through APIs and actuators and 3D printers and cars that drive themselves, the loop doesn't close at the screen. It closes at the world.

You can, in a literal sense, speak things into existence now.

The Biblical echo is on purpose, and I don't mean it grandly. I mean it as a technical description. "Let there be light" is a sentence whose form we've been unconsciously imitating ever since: an imperative spoken aloud, addressed to an agent capable of responding, producing a change in the world. What's changed is only who the agent is. For most of history the agent was God, a human mimicking God in a sense, or no one. Now the agent can be a system that understands you well enough to do what you asked, or at least what it thinks you asked, which as we'll see is not always the same thing.

Consider what this actually lets you do. You can describe a program and get a program. You can describe an image and get an image. You can describe a piece of music, or the mood of a piece of music, and get something plausibly musical back. You can describe a room and get the lights to change. You can describe a part you need and have a machine cut it. The

list keeps growing, and the growth is accelerating, and the interesting question isn't whether the list is impressive — it is — but what kind of skill the new situation rewards.

Here I think the answer is surprising. The new skill isn't technical. It's something older. It's the ability to say what you actually mean.

This turns out to be harder than it sounds. Most of us, most of the time, speak in approximations. We say “make it pop” when we mean “I don't know what I want, but I'll recognize it when I see it.” We say “just ship something simple” when we mean “I haven't thought about this enough to know what simple is.” We say “build me a dashboard” without saying what questions the dashboard is supposed to answer. With a human collaborator this is survivable because the human does a silent, mostly-unconscious job of translating your approximation into something specific. They ask follow-up questions. They guess well. They bring taste and context that they acquired by being alive in the same world as you.

An LLM can do some of this. It can ask follow-up questions; it can guess. But it doesn't have your life, and its guesses, especially when the stakes are high, are only as good as the words you gave it to guess from. The machine takes you at your word in a way a human collaborator usually doesn't. The upside is that your word now carries further than it used to. The downside is the same.

The old fairy-tale archetype is the wish. Monkey's paw, genie, leprechaun — the ancient intuition, embedded across nearly every folklore I know of, is that getting what you asked for is dangerous if you didn't know what you were asking for. That intuition is suddenly practical. We are building a world in which the gap between “said” and “meant” has real and immediate consequences. Clarity of intent — the thing your English teacher was trying to teach you when she made you rewrite the paragraph — is becoming load-bearing in a way it hasn't been for most of your professional life.

This is, I think, good news, though not in the way the marketing decks frame it. The marketing decks tell you that voice plus AI will let anyone do anything, that the democratization is the story. And there's some truth to that. But the deeper story is that voice plus AI rewards people who know what they want. It punishes vagueness severely. And the skill of knowing what you want — being able to describe a thing specifically enough that someone, or something, could build it without asking you forty clarifying questions — is one of the oldest and most undertrained human capabilities there is. Logical thinking and clearly articulating intention are skills we take for granted, just as we do the forgiveness of the listeners to whom we speak.

Writing used to teach you this. You sat with a blank page and discovered, often painfully, that what felt like a clear idea in your head was a fog when you tried to lay it out in

sentences. Writing was thinking, as a man named Paul Graham likes to say, and a lot of what it taught you was the gap between how much you thought you understood and how much you actually did.

Voice plus AI is writing's strange successor. You still have to know what you mean. But you find out faster, because the thing you said produces a result in front of you almost immediately, and you can see whether the result matches what was in your head. If it doesn't, you say it again, differently. The feedback loop that used to take a week now takes thirty seconds. This is the real gift. It isn't that you can avoid thinking; it's that you can't, because the system will show you very quickly when you haven't.

There's something else worth noting. The imperfections of a tool can activate rather than frustrate the creative process, because we discover what we meant by noticing what we didn't mean. Voice plus AI has exactly this property, and at high frequency. Every small misunderstanding is a little mirror held up to the fuzziness of your original intent. If you pay attention, you become a clearer thinker almost by accident, just from the friction of being misheard productively.

And then there is the physical world, which is the part that still feels like magic. For most of my life the computer and the world were separated by a pane of glass. Now the pane is porous. A spoken sentence can start a 3D printer. A spoken sentence can move a robot. A spoken sentence can commit a transaction, route a truck, open a valve, book a flight, draft a contract, turn on a light in a room on another continent. Each of those sentences is a tiny act of creation, and the accumulated weight of all of them is something new in the world — not because any individual act is impressive, but because the distance between thought and effect has collapsed.

Which returns us to the oldest sentence in the tradition. "Let there be light" isn't impressive because of the grammar. It's impressive because there was a gap between the saying and the being, and the saying closed it. For most of human history only poets and prophets pretended to speak with that kind of authority, and most of the time they were bluffing. We are not bluffing anymore. Not entirely. The gap between saying and being is, for a growing class of sayings, genuinely closed or close to it.

The responsibility that comes with that is older than the technology, and it hasn't changed. If your words build things, you had better know what you're trying to build. If your voice can reach into the physical world, you had better mean what you say. The tools are new. The burden is ancient. Let there be — but let there be something worth making.