Google’s AI pitch is a recipe for email hell

Of all the emotions I expected out of this week’s Google I/O keynote, “dread” was not on my short list. Google has spent a couple of decades creating products that make my life easier. This week, it demonstrated the exact opposite: a way to make communicating with other people worse.

Google I/O, this year, was focused on artificial intelligence, and one of its centerpieces was prompt-based text generation, particularly in Google’s office suite. At several points throughout the keynote, it demoed a system called “Help me write” — which rewrites simple prompts in Gmail and other apps as more polished paragraphs of text. But *more polished* doesn’t automatically mean *better.*

Good writing can delight us with an intriguing turn of phrase. It can tell us something about how a person thinks and feels. Or it can offer some kind of useful detail about the world. And there’s no reason AI can’t play a role. I’ve seen authors use generative tools to create unexpected prose or summarize data in a readable way. I’m excited to see people explore new uses for it.

But Google’s onstage ideas were almost impressive in their ability to serve precisely none of writing’s core purposes. Its examples of Help me write calcify the worst tendencies of “professional” communications. They’re bland, bloated boilerplate that turns a human prompt into something that uses more words to say less.

Let’s look at one of the examples. Early in the keynote, a presenter auto-generates a job description for a textile

design position that they (in the demo's scenario) send into the world. Later, another presenter responds to it, supplying the following prompt:

> *interested! 8 years exp with intl brands. Call on Tuesday.*

Help me write works its magic and produces the following:

> *Hi Lori,*
>
> *Thank you so much for reaching out! I am very interested in the Senior Textile Designer position at Cymbal. I have 8 years of experience working with international brands, and I am confident that I have the skills and experience that you are looking for.*
>
> *I am available for a call on Tuesday at your earliest convenience. Please let me know what time works.*

The latter version is obviously the way you're expected to apply to a job because, written by a human, it suggests you're committed enough to type out a multi-paragraph reply and are capable of hitting certain class and cultural signifiers. An AI-generated response doesn't serve either of those purposes, though. It was created almost instantly based on a predictive text system that requires minimal English-language competency, and so far, that system can only handle fairly formulaic writing; it won't substitute for human work in many real-world situations. So all it does is construct a longer and more stilted version of the original prompt — one that probably only has value until everyone expects it was written with AI.

And even worse, the AI generation reinforces the idea that overenthusiastic US business speak is the *required* way to write, regardless of whether it's a necessary skill for the job. I've seen thoughtful stories about people with dyslexia using ChatGPT to produce text that is — as a *Washington Post* article puts it — "unfailingly professional and polite." But

there’s an unspoken, simpler alternative: being willing to accept wider variations in how people communicate. I don’t begrudge anyone who uses AI writing to meet largely arbitrary standards, but at a society-wide level, it’s a linguistic arms race toward a more boring future.

For another example, try this friendly message of congratulations from Google Android engineering VP Dave Burke to devices and services SVP Rick Osterloh:

> *rick congrats on pixel fold launch amazing with Android dave*

Here’s what Help me write churns out:

> *Dear Rick,*
>
> *Congratulations on the successful launch of the Pixel Fold! I’m really impressed with the device and how well it integrates with Android. The foldable screen is a game-changer, and I can’t wait to see what you do with it next.*
>
> *Best,*
> *Dave*

I’m going to assume this is tongue-in-cheek because I know the way executives send personal emails, and they seem entirely comfortable with one-sentence punctuation-free missives. But again, it’s worse! It tells me less about what the sender sounds like. It takes longer to read. It makes the sender sound like an interchangeable robot.

To understand the pragmatic value of this process right now, you have to look at an email from the very beginning of the keynote. The core of the generation is actually pretty cool — it takes a very simple prompt for an airline refund request and fills it in with concrete details drawn from your inbox, like the flight name and date. But after that, it embellishes the facts with a few paragraphs of irritated prose. Here’s the

original:

> *ask for a full refund for this canceled flight*

And the AI-generated version:

> *Dear Customer Service,*
>
> *I am writing to request a full refund for my canceled flight, DF 4542, on February 4, 2023. I understand that you are offering a voucher as a gesture of goodwill, but I would prefer to be reimbursed for the cost of my ticket.*
>
> *I have been a loyal customer of your airline for many years and have always been satisfied with your service. However, I am very disappointed with the way my recent flight was handled. I was not given any advance notice of the cancellation, and I was forced to make alternative arrangements at the last minute. This caused me a great deal of inconvenience and expense.*
>
> *I believe that a full refund is the only fair way to compensate me for the problems I experienced. I have enclosed a copy of my original ticket for your reference.*
>
> *Thank you for your time and consideration.*
>
> *Sincerely,*
> *Rowan*

There’s no indication these furious claims are accurate. It’s a bureaucratic deimatic display: convincing some hapless customer service representative that you’re mad enough to type out all that text and will probably be a real nuisance until you get your money back. I’ve seen this idea pitched better on Tumblr of all places. In a popular post about AI, someone described using ChatGPT to pull off a sort of high-tech version of *Brazil*’s Form 27B / 6 gambit, generating a

threatening faux-legal letter to a landlord who was breaking housing laws. As a fellow longtime tenant, I applaud them.

But this stuff is effective in a brief transition period while generative text isn't in widespread use, while readers are likely to assume it's connected to a human writer. If you know it's a machine, the illusion evaporates. You're left with a world full of communications that are longer, less thoughtfully created, and more mind-numbing to read. I'd rather hire someone based on an honest "8 years exp" than a cover letter full of empty automated prose.

By contrast, Google's most useful example of Help me write involved simply conveying information. In an email about a potluck, its AI was able to look at a document with a list of dishes people had signed up to bring, then summarize that list as a line in an email. It saves writers the step of pasting in a series of items and readers the inconvenience of clicking through to another tab. Most importantly, its value doesn't rely on pretending that a human wrote it — and if Google has its way, that's a trick that won't last for long.

https://www.theverge.com/2023/5/13/23719115/google-ai-help-me-write-communications-email