



Five mistakes agencies can make when implementing generative AI—and how to avoid them

Many government agencies are thinking about how generative artificial intelligence (AI) is going to impact them, and perhaps how they may be able to use it to their advantage. These generative AI solutions are based on large natural-language models and are capable of creating remarkably sophisticated, humanlike responses with very limited prompting. They've caught the attention of nearly every public- and private-sector leader.

Caught up in the excitement, some agency leaders may be looking only at the upside and failing to look at the pitfalls that could doom their efforts. If you're among those at least toying with the idea of generative AI, it's worth considering what we think are the five biggest mistakes you can make—and what you can do to help avoid them.

1. Waiting too long to implement generative AI internally

Despite what may at times feel like excessive hype, generative AI is real. Some have dismissed it as a fancy autocomplete technology or a high-end chatbot, but it's far more than that, and its capabilities seem to be advancing nearly daily. These models have learned to communicate in languages they were never taught, for example.¹ Given the extraordinary pace of its development, it can be difficult to imagine what it might be capable of even just a year or two from now—or what implications this may have for government agencies.

Even if it's only 10 percent as disruptive and transformative as it's being portrayed, generative AI will change the way we all live and work in significant ways. Therefore, perhaps **the worst thing you can do is to pretend it isn't happening**, or to think that you can make it go away.

Yet it may appear that some organizations are trying to do just that by banning its use. To be fair, many of these organizations have legitimate and serious concerns with

Why modern government is important

Government agencies in the U.S. must modernize in order to keep up with changing user needs, regulations, and health and public safety requirements. Leaders of modern governments rethink business processes and service delivery models to more effectively achieve their mission. This article is one of a series that features how modernizing affects the government workforce and the user experience, improves security and public trust, and accelerates the digital journey. KPMG team members offer insights intended to help guide governments in their modernization efforts to encompass all processes, technologies, policies, and the workforce so each works together to create connected, powered, and trusted organizations.

generative AI that must be addressed, most notably intellectual property leaks and privacy or regulatory violations. However, banning AI functionality and platforms may be as harmful as it is impossible. **Trying to ban or ignore generative AI is like trying to ban or ignore the internet.** Like it or not, it will be used. Government employees at all levels—federal, state and local—are already using ChatGPT or one of its equivalents, even if it's surreptitiously. Indeed, it *should* be used, but for the proper purposes, and with the proper education and guardrails.

Our colleague, Viral Chawda, has written a [great article](#) on generative AI risks specific to government agencies, and perhaps more importantly, what you must do to address them, including establishing an effective responsible AI program.

¹ Source: "Solving The Mystery Of How ChatGPT And Generative AI Can Surprisingly Pick Up Foreign Languages, Says AI Ethics And AI Law," Forbes, Lance Eliot, April 19, 2023





2. Thinking of generative AI in terms of “tasks” instead of “outcomes”

Given the hype around generative AI, which seems to suggest that it can do everything from curing cancer to bringing about world peace, what’s easy to miss is that generative AI models are primarily designed “only” to generate human-readable content—that’s the “generative” part. They can perform remarkable tasks such as providing analysis of massive data sets, generating images and writing software source code, but the output is always a human-readable response to a prompt users provide. **Recognizing its limitations** is an essential first step in defining what you can realistically accomplish with it.

Once you understand its capabilities and limitations, the next step is to **determine how your agency can best exploit it**. In our conversations with clients in the government and private sector alike, we hear that many want to explore how they can use generative AI to perform a task instead of achieving an outcome. While using generative AI for tasks has the potential to create fruitful outcomes such as automation, taking an outcome-based approach instead can help you realize its full potential.


You must avoid implementing generative AI and hoping it will find its purpose. The field is littered with examples of IT rushing to implement the latest and greatest technology, yet zero business results are achieved. You must define clear use cases and realistically attainable goals up front.

It’s also not an “install it and check the box ‘done’ project.” It’s a business transformation, and as daunting as it may sound,

that means there is no end. It means it can’t be thought of in isolation; you must **continually reexamine your goals** and rethink how generative AI should integrate into the end-to-end ecosystem of your technologies and your business processes. Get ready—given the extraordinary pace of generative AI improvements, this cycle may repeat far faster than you’re used to.

Just as important, you must **define clear metrics** to help ensure those goals are being met. All too often, organizations define and implement the wrong measures of success—output rather than outcomes, usage rather than results—especially when the project is largely technology driven. You can’t forget the human input, either—from both your agency’s employees and its customers. Defining mechanisms that capture and help you evaluate their feedback is essential.

The measures of success are an essential part of the feedback loop designed to continually improve and refine the technology, its usage and the experience people have with it. Even here, there’s a common mistake: believing that having such a continuous improvement loop will eventually lead to success. Everything must be constantly questioned and reexamined, and feedback from all sources—including humans—continually reevaluated to help ensure even the metrics themselves are still valid.





3. Thinking that generative AI is a replacement for humans instead of an enabler for humans

It's tempting to think that generative AI might be a viable replacement for people, especially given its very humanlike interactions. It's not. While an algorithm can surely be a helpful tool for an employee or customer, it's just that: a tool, and one of many in the bag that can help augment or enhance an experience. Generative AI is a powerful enabler that can improve worker productivity by handling more routine tasks, freeing them to handle more complex or sensitive issues, but it can't eliminate them from the equation—nor should it.

The recent experience some agencies have had with conversational-AI chatbots is relevant here. These chatbots had been touted as a massive productivity booster by replacing many (if not most) customer service representatives. Many public and private sector organizations have spent millions of dollars chasing this dream, building chatbots using the much-hyped commercially available solutions only to find they didn't come anywhere close to living up to this promise.² Users were often frustrated, disillusioned or worse: misinformed.

This isn't really a failure of the technology—in most cases, it performed as described. The issue has more to do with expectations, and specifically with how organizations expected their users would respond. A primary challenge is that, good as they are, AI-powered solutions aren't yet as capable of understanding human emotion and intent as other humans are. They are not as adept at handling nuance or "edge" cases, which is often what is prompting people to seek assistance and which in government may be more common than not. Most particularly, they are not yet capable of "reading between the lines" and understanding what is *not* being said.

Because of these limitations, **AI-powered solutions aren't a substitute for humans**—they create very different experiences and therefore are best applied to very specific use cases. Programs such as social services, for example, typically involve dealing with people in difficult and complex situations. This requires human empathy, judgment, and assessment, which cannot yet be replaced by machines. Yet some agencies are trying. Some child welfare agencies, for example, are using automated "robocalls" to communicate with parents involved with child welfare cases, informing them that they have missed appointments or hearings. However, critics argue that this approach can be impersonal and insensitive, particularly given the sensitive nature of child welfare proceedings.

The best applications deploy generative AI as an enabling tool—for example, to help front-line workers in your agency on whose expertise you depend, including helping them document issues, summarize lengthy content or highlight red flags. We're already seeing agencies train generative AI models on internal policies and procedures to support case workers, with positive results.

Others have replaced reporting dashboards with generative AI. These implementations can enable customer service managers, for example, to ask questions and get answers in real time, even questions that hadn't been anticipated in their predefined, hard-coded dashboards. Ask, "What are the most frequent reasons for call drops?" and a graph or chart is rendered instantly, without any custom software. Such solutions are enablers for humans to make better decisions and not anything approaching a replacement for them.

With the lack of experience both public- and private-sector organizations have using generative AI to enhance user experiences and given the associated risks tied to potential inaccuracies, it may be best to explore employee-focused initiatives such as these first.

² Source: "AI Chatbots: Reality vs. Hype," Navveen Balani, *Becoming Human: Artificial Intelligence Magazine*, July 3, 2019

4. Focusing on the technology instead of the experience

Because generative AI is a collaborative tool and not merely an automation technology, a change in thinking is required. Typically, humans and automation technologies are siloed, but here, you must think about how they must work together.

The key is to **always be focused on the experience that generative AI creates, enables or enhances—not how that experience is powered.** Once you define the business issue that you believe generative AI can help address, your execution must start with the employee and/or customer experience and work backward. You may find that generative AI isn't the solution after all. Don't let the technology dictate the experience.

KPMG has defined six pillars of customer experience excellence that can serve as a guide as you consider this process. Integrity and trust are the foundation of an excellent experience—and perhaps the one most difficult to achieve with a generative AI solution.

Focusing on the experience rather than the technology may offer an additional benefit unique to government.

Unlike private sector organizations that have the freedom, if not the flexibility and agility, to pivot quickly to meet a new business challenge or seize a new opportunity, the structure and missions of government agencies are immutable. Almost on a dime, Corning, for example, stopped making dishware and started making fiber optics. Imagine housing and urban development making a similar left turn.

Generative AI may represent an opportunity for agencies to create such flexibility. It's not uncommon for customer issues to span multiple agencies, and despite the goal of supporting one-stop customer access, in reality customers are often left on their own to navigate a complex maze of agencies and their unique requirements. By enabling flexible, intelligent, contextually-relevant user experiences that are independent of or blind to agency boundaries (based on the data you train your models on), you may be able to provide one-stop access, or at least provide a roadmap that can help customers navigate the maze rather than solving a siloed need.

The Six Pillars of Customer Experience Excellence



Source: KPMG 2018 Customer Experience Excellence Centre



5. Ignoring the legal, ethical and privacy concerns

By default, generative AI is trained on public data—not your agency’s data or data from other agencies. Unless and until you integrate it with and train it on the appropriate internal systems, it will not have any knowledge related to the data in those systems. Deploying and integrating it with other systems is a nontrivial exercise, often grossly underestimated, especially if you’re in the same boat as many other agencies that are saddled with a slew of aging legacy systems. But the deployment and integration challenges are just the tip of the iceberg.

How much data access do you want to give to your generative AI model? What’s it going to do with that data? Who gets to see the output and what do they get to see? What are the possible implications of decisions made based on its “analysis” of it? How much transparency do you have into how it makes those decisions or generates those responses? Are decisions or output free from bias or discriminatory responses? Does it comply with legislative or executive requirements? **The legal, ethical and privacy concerns are significant** for any organization, but for government they are critical.

Data quality matters, and the classic risk of “garbage in, garbage out” is very real with generative AI. You must have the proper data governance framework in place or you open yourself up to privacy, discrimination or other regulatory violations—or at least to embarrassment and criticism since laws or regulations vis-à-vis AI have not been able to keep pace with such a fast-moving technology. Efforts are underway to more clearly define the rights employees or customers should have, but most are still in relatively early stages.³

Finally, don’t underestimate the emotional reaction users—both employees and customers—may have when “forced” to engage with an AI-powered model. This can extend beyond people’s immediate interactions with the technology to include larger, more philosophical concerns. As Stephen Marche, a noted author and essayist who has written extensively on AI, observed, “Artificial intelligence is an ethical quagmire...If users fear artificial intelligence as a force for dehumanization, they’ll be far less likely to engage with it and accept it.”⁴

Even if the solution is designed for customers, don’t dismiss the experience of your employees. Many will fear that generative AI will replace them. Others will have to deal with customers who’ve used the tool—again, think of those poor customer service agents who are connected to customers left angry or frustrated by a poorly designed chatbot.


³ Source: “How privacy legislation can help address AI,” Cameron F. Kerry, Brookings Institution, July 7, 2023

⁴ Source: “The Chatbot problem,” Stephen Marche, The New Yorker, July 2021

KPMG: an intelligent approach to business transformation

KPMG is an early and enthusiastic advocate for the power of AI. We have broad and deep experience in generative AI technologies, finance functions and processes, and business planning. With our suite of intelligent tools and services spanning project architecture, operating models, data and signals, and highly skilled data science and business talent, we are well positioned to help your organization leverage generative AI to transform operations and outcomes.

We can help guide your organization through strategy development, platform selection and implementation—and then provide ongoing support to help you optimize your investment in this powerful technology. We understand both the promise of generative AI and the process and cultural changes, including the embrace of responsible AI practices, that will be required to realize its full potential. Talk to us about how we can help you turn data into insights, and then into advantage.



About KPMG

KPMG has worked with federal, state, and local governments for more than a century, so we know how agencies work. Our team understands the unique issues, pressures, and challenges you encounter in the journey to modernize. We draw on our government operations knowledge to offer methodologies tailored to help you overcome these challenges and work with you from beginning to end to deliver the results that matter.

The KPMG team starts with the business issue before we determine the solution because we understand the ultimate mission. When the way people work changes, our team brings the leading training practices to make sure your employees have the right knowledge and skills. We also help your people get value out of technology while also assisting with cloud, advanced analytics, intelligent automation, and cybersecurity. Our passion is to create value, inspire trust, and help government clients deliver better experiences to workers, citizens, and communities.



Contact us



Arthur Higbee
Managing Director, Advisory
Transformation Delivery
KPMG LLP
ahigbee@kpmg.com



Christina Melton
Experience Design Director, Advisory
Transformation Delivery
KPMG LLP
cmelton@kpmg.com



Pamela Powers
Managing Director, Advisory
Transformation Delivery
KPMG LLP
ppowers@kpmg.com



Tom Frame
Managing Director, Advisory
Digital Transformation
KPMG LLP
tframe@kpmg.com



Sarah Rodriguez
Director, Advisory
State and Local Government
Solutions
KPMG LLP
sbrodriguez@kpmg.com



David Klimek
Principal, Advisory
Customer Service and Experience
Lead, Health and Government
KPMG LLP
dklimek@kpmg.com



Bharat "Bobby" Gorantla
Managing Director, Data & Analytics
KPMG LLP
bgorantla@kpmg.com

read.kpmg.us/modgov

kpmg.com/socialmedia

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.



The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2023 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization.