CPA.com Generative AI Toolkit

A roadmap for accounting and finance professionals to understand and leverage the transformative impact of GenAI

CPA.com™

CPA.com's Multi-Part Generative AI Initiative

Over the past year, we've witnessed the tremendous potential of generative AI (GenAI) to reshape how businesses work. In response, CPA.com is leading a multi-part GenAI initiative to support CPAs' understanding and utility of the transformative technology. This toolkit is part of a broader collection of resources developed to empower professionals to assess the possibilities and navigate the associated risks of implementing GenAI into their practices.

We encourage readers to explore the additional resources, including our AI-focused sessions at the Digital CPA conference, videos featuring expert insights, and more. We will also have an AI-specific cohort as part of the 2024 Startup Accelerator program and will hold our first AI Symposium that will bring together subject matter experts, firm leaders and AI solution providers to discuss the real and potential impacts of the transformative technology on the profession. Learn more about the broader initiative and find the latest resources at CPA.com/Gen-AI.

> **We believe generative AI's impact will be much more profound and happen a lot faster than other technology advancements, such as cloud computing.**
>
> — ***Erik Asgeirsson,*** *CEO, CPA.com*

Table of contents

3 Introduction

6 Driving digital transformation

8 Seven steps to building an AI strategy

10 Use Cases: Practical application of generative AI for accounting and finance

14 Challenges and limitations of generative AI

17 The great equalizer

18 Security and risk considerations: Safeguarding validity, privacy and transparency

Introduction

Very shortly after it was first introduced, ChatGPT became the fastest-growing consumer application in history. In less than five days, the program reached a million users; a technology rollout record. Within two months, it had been used by over 100 million users. By contrast, it took roughly nine months for TikTok to hit 100 million users and over two years for Instagram to do so. That's a testament not only to the excitement about the technology but also to its accessibility and ease of use. Just a few months after ChatGPT's release, 77% of business leaders surveyed cited generative AI (ChatGPT is one of many examples) as tech they expected to have the largest impact on their businesses out of all emerging technologies over the next three to five years. That was nearly double the number for the next highest technology according to the KPMG survey. A total of 71% of leaders plan to implement their first generative AI solution within the next two years.

Artificial intelligence (AI) is an umbrella term that encompasses a range of technologies that are generally intended to perform tasks that might otherwise require human intelligence. It is already very much a part of our everyday lives, powering search engines, self-driving cars, facial recognition and virtual assistants like Siri and Alexa as just a few examples. Accountants use it to automate data entry and reconciliation, perform financial modeling and forecasting, detect anomalies in auditing, risk management and in many other use cases.

A message from ChatGPT, I'm here to help!

Working with generative AI is often a conversation. When we asked ChatGPT to compose a message to accountants about how it could assist them, its first output was much longer than we wanted, so we asked the system to shorten it. The next response was better, but still a little wooden, so we asked it to turn it into a rhyming verse. Here's the final product:

Hey there, accountants, listen up, it's me, your AI
friend, ChatGPT! With generative AI, we've got a
plan to revolutionize your work—and we can!

Automating tasks with a swift hand, tedious
chores, I'll take a stand. Data entry, reconciling,
tax prep too, all in a flash, I'll do for you!

Predictive insights, I'll unveil, financial trends,
I won't fail. With time freed up, you'll soar so
high, advisory roles, you'll reach the sky!

Personalization, I've got it down, customers
with smiles, not a frown. Chatbots aid, support
in sight, answering queries day and night.

Digital transformation, we'll lead the way, a
partnership that's here to stay. Together we'll change
the game, embracing AI, we'll rise to fame!

So, accountants, let's unite and soar, with AI's help,
we'll do much more. In this tech-driven world, we'll
stand tall, making accounting a blast for all!

Generative AI is a subset of the broader artificial intelligence category that can perform more complex steps. It has reshaped the technological landscape by offering advanced capabilities across a variety of domains. Beyond just Large Language Models (LLMs, see Glossary) like ChatGPT, generative AI encompasses a broad range of modalities:

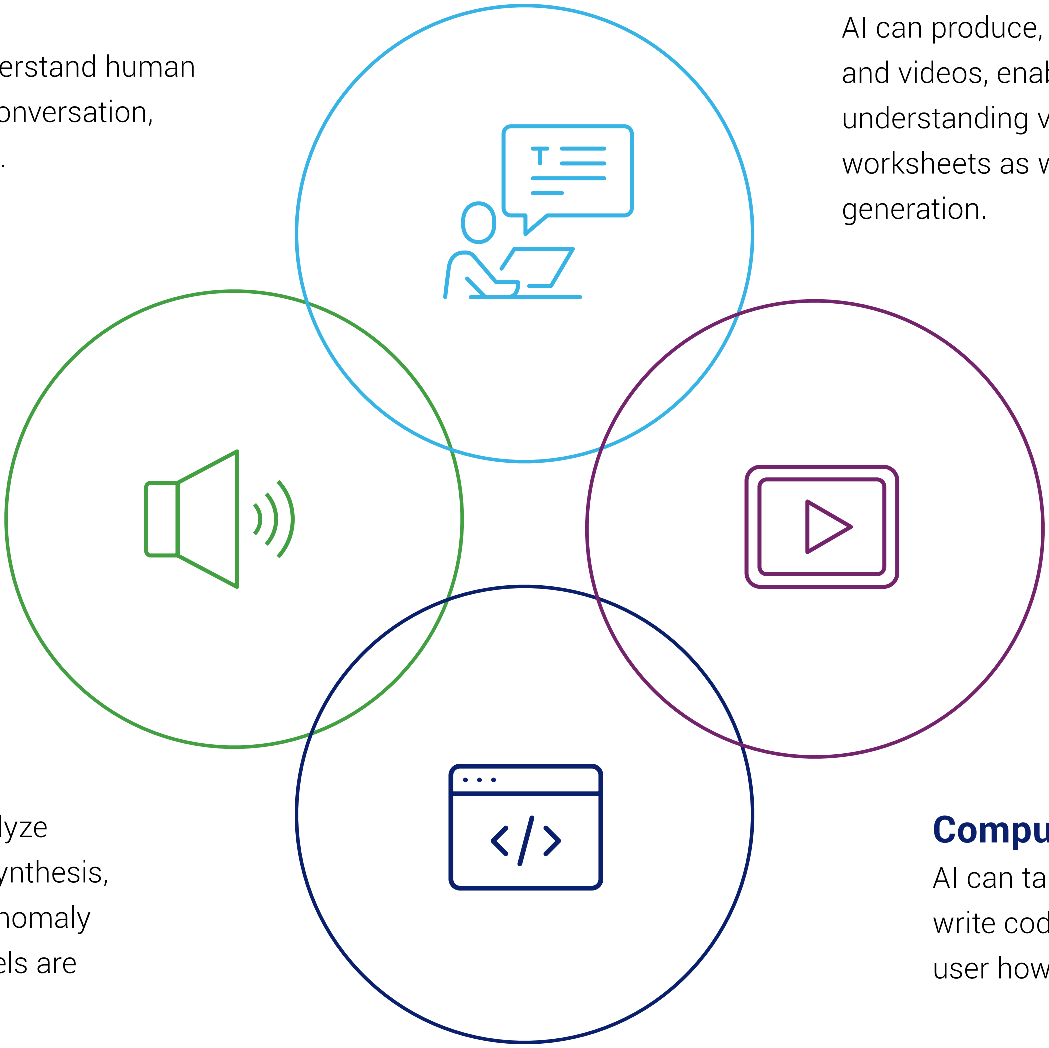

Text

These models generate or understand human language, enabling tasks like conversation, summarization and translation.

Image and Video

AI can produce, modify or interpret images and videos, enabling applications such as understanding visual tables, flowcharts, and worksheets as well as art creation and deepfake generation.

Sound

These models generate or analyze sound waves, useful in voice synthesis, music production, and audio anomaly detection. Speech to text models are also gaining in popularity.

Computer Code

AI can take a description of a task and write code to achieve it, or teach the user how a code example works.

Natural Language Processing has been around for a while, most recognizably in the chatbots that pop up on many websites. However, generative AI has advanced quickly in recent months due to new innovations in the LLMs on which it is based.

Although the content it can generate is important, the real power of generative AI lies in its ability to understand complex instructions and break them down into individual tasks it can perform in response, according to Aaron Harris, CTO of Sage. There are also limits and serious risks associated with generative AI. The security of any data that is shared with this technology is a major concern, particularly when it comes to sensitive client information. At least in its initial iterations, generative AI has also been found to get things wrong and even make things up—or hallucinate. Its confident, articulate and well-reasoned answers may mask this significant failing.

This toolkit is intended to help accounting and finance professionals to understand generative AI along with its risks, challenges, opportunities and limitations. It includes various use cases that demonstrate how generative AI can be put to work for accounting and finance professionals. It also provides a 7-step plan to help build your AI strategy, security & risk considerations to help firms anticipate and minimize related threats, and a glossary of terms.

> **Nearly every accounting technology provider today is exploring if and how generative AI can be integrated into their existing solutions to address firms' top challenges. When harnessed the right way, these capabilities present an enormous opportunity for the accounting profession to deliver higher value services."**
>
> — ***Erik Asgeirsson,*** *CEO, CPA.com*

SCAN ME

Be sure to visit CPA.com/Gen-AI for explainer videos on each use case, additional learning opportunities and more.

Driving digital transformation

Organizations have embraced a variety of forms of digital transformation in recent years, many of those being driven by technological advancements, as well as the disruptions and innovations caused by global events. Generative AI is producing new applications and opportunities that are enabling wider and speedier transformation. The emerging tech has been compared to the internet and smartphones in terms of the impact it will have on daily life and business. In fact, Open AI released a study highlighting occupations with the greatest risk of being disrupted, noting that 100% of tasks for accountants are currently exposed.

When it comes to digital transformation, among the many things that generative AI can enable are:

- Automating workflows from beginning to end.
- Handing off numerous tasks to digital assistants that function as resources within the organization or for clients.
- Hyper-personalizing marketing content, taking on routine sales tasks and conducting better lead identification to optimize business development and growth.
- Allowing companies to design, test and bring new products to market more quickly.
- Speeding uses of tools, such as 3D printing and computer-aided design, and identifying and enhancing design options.
- Minimizing the accounting profession's staffing shortage by taking on a wide range of tasks, even beyond the repetitive steps that many AI tools are already handling.
- Analyzing and identifying potential instances of fraud and other risk anomalies.
- Potentially making it easier to address environmental, social and governance (ESG) issues.

"Part of the challenge is that it's very difficult to see all the ways that generative AI will evolve and emerge," Harris said. "It is so powerful and it's moving really fast. There are already hundreds of companies getting venture funding and doing really impressive things with it." CB Insights reports that investment in generative AI startups reached $14.1 billion in equity funding in 86 deals during the first half of 2023.

How will it work in practice? Within CPA firms, customized generative AI systems could be designed to immediately provide a list of new changes to the tax code or compare the details of a new standard with an old one. Firms could use it to create personas for internal use that have expertise in particular service lines, technologies or other areas, thereby significantly cutting down research time. To prevent inaccuracies or hallucinations, a firm can limit the content that the technology is trained to pull from so that it is referencing only completely reliable sources, such as the Internal Revenue Code or the text of other standards, laws or regulations. Sage's Harris envisions specialized generative AI agents that will be trained specifically to manage the accounts payable workflow or interact with vendors. An AI agent designed for office productivity might help a firm perform analysis and research, then turn its findings into a professional document or build PowerPoint slides. To create a more thorough, personalized and competitive proposal, firms can use generative AI to gather publicly available information on a prospect's industry, market and competitors, making it easier to highlight the firm's understanding of the prospect's business and the many ways that it can meet their needs.

> **Part of the challenge is that it's very difficult to see all the ways that generative AI will evolve and and emerge. It is so powerful, and it's moving really fast."**
>
> **— *Aaron Harris,*** *CTO, Sage*

CPA.com Seven steps to building an AI strategy

Imagine that someone who plays baseball suddenly switches to basketball. Their underlying athletic ability remains the same, but they have to relearn what it takes to succeed in the new arena. When we learn a new sport, we not only learn the rules of the game, but we must also become comfortable with new movement and develop new muscle memory. "That change is similar to the mindset shift necessary with generative AI," said innovation expert and CPA.com collaborator Pascal Finette, co-founder of be radical. It's not just a new way of using technology. It's a new way of thinking about technology and work in general. That's why, in understanding the rules of the game with artificial intelligence—and more specifically generative AI—it's important to avoid trying to slot them into existing expectations. Instead, we need to be open to developing new ideas, uses and capabilities.

The steps below outline a way to shift your mindset and help your organization to develop its own AI strategy.

STEP 1

Experiment

Tools like ChatGPT and Bard are very easily accessible, so pick one, create an account and start experimenting to understand the technology's potential. Considering the rapid pace of change in this space, it's a good idea for firms to get their feet wet as soon as possible, just don't put personal identifiable information in these...yet.

STEP 2

Focus on the near-term

Given the accelerated advancement in this space, attempting to create a five-year plan for generative AI use is futile. At least at first, a two-year strategic plan may be optimal.

STEP 3

Consider client expectations

On the one hand, many companies are likely still in the learning phase themselves, so firms should not shortchange their learning curves because they worry about staying ahead of client expectations. However, firms should expect that generative AI will quickly become clients' user interface of choice. Get a sense of where clients are, then plot out where the firm needs to be and how generative AI can address both firm and client needs.

STEP 4

Research what solution providers are building

Before building your own generative AI tools, ask current and potential vendors about their existing offerings or plans. "I can guarantee that nearly every solution provider has something in the hopper regarding use of AI in the accounting, tax or audit software space," Finette said. Determine how new or emerging tools can be incorporated into the firm's strategic planning.

STEP 5

Define acceptable use

Some organizations have questioned whether it's best to ban generative AI use to limit the associated risks to privacy and confidentiality. Unfortunately, prohibitions may backfire if staff members are accessing the technology on their own, unsecure devices. In addition, firms should want to promote staff capabilities with artificial intelligence. To establish guardrails, circulate an acceptable use policy as early as possible in the process. Among other things, the policy should prohibit uploading or asking questions about client data within a public generative AI tool. Exposing sensitive client or firm information is a risk because anything entered into ChatGPT, for example, becomes public domain and is used in the system's continuous training of its machine learning model. To address staff research questions related to client data, a better ultimate choice might be to develop an internal-use bot that can provide answers based on reliable authoritative or locally stored data. (For more information and advice on acceptable use policies, see the Risk & Security Considerations later in this toolkit.)

STEP 6

Put the technology's best features to work

Since the systems available right now work well with text, identify the ways that the organization is spending time and energy on text-based work, such as responses to clients, business development and marketing emails. Then determine how generative AI can reduce current effort and costs. Revisit this step regularly as the technology's capabilities mature.

STEP 7

Ensure that humans maintain a key role in the process

Because generative AI output may include biases, errors or hallucinations, a human should review and ensure the appropriateness and accuracy of any content that is being used in decision-making or shared with clients. As an example, the system could be asked to identify emails that are marked as urgent and to analyze and report on the content. The user can then ask the system to create responses, stipulating their length and tone. The final, and most important step, would be for someone to review the emails before they are sent to ensure they have the right tone and accurate content for each recipient.

Trust will be vital when offering products or services associated with AI, as clients are handing over responsibility for completing a task that's critical to their businesses. If there's not perfect confidence in an outcome, generative AI projects will require human oversight.

CPA.com Use Cases: Practical application of generative AI for accounting and finance

Organizations can make the most efficient use of generative AI if they understand how to ask the right questions. These use cases illustrate the kinds of requests users might make in different situations and examples of how to phrase them to get the best outputs. **These examples are referred to as PROMPTS.** For each use case, you can watch an explainer video and access swipeable prompts at CPA.com/GEN-ai. Note that these use cases can be used with most any GPT system; we have identified specific solutions like Claude, ChatGPT Code Interpreter, etc. due to functionality. There are hundreds, if not thousands, of generative AI tools on the market. For these use cases, we have leveraged low-cost/no-cost solutions that anyone can easily access.

The use cases have been categorized based on the practice area or business subset that would most likely benefit from its implementation.

SCAN ME

Companion Videos
CPA.com/genai-practical-use-cases

Please note, any examples that reference uploading data will require all identifiable information to be removed prior to use.

Finance Use Case: Structured Data Extraction

In a perfect world, systems would talk to each other and sync data seamlessly. Unfortunately, working with many systems still requires exporting, importing and reformatting data. Language models are particularly adept at taking unstructured data—either from an export or even simply copied and pasted from a PDF report—and turning it into structured data. Here's how to use ChatGPT to take unstructured data in various formats, quickly reformat it into a structured comma-separated values (CSV) file, and validate that output is correct.

Example prompts:

- *The following is an accounts payable detail report. Structure the data into a CSV file with columns for Invoice Date, Vendor, Amount and Currency.*
- *The following is a bank statement. Structure the transactional data into a CSV file with columns for Transaction Date, Description, Amount (which is the sum of Debit minus Credit), and Available Balance.*

- *The following is an audit trail report. Return all the data from John Smith user into a CSV file with the columns Date, Time, Action and IP Address.*

Tax Use Case: Legislation Research Assistant

Navigating large documents like legislation can be difficult when relying solely on text-based search. While Ctrl+F may get you to the right section, it will only work if you use the exact keywords. Large language models, however, enable a more general semantic form of search that can surface relevant sections more flexibly. By prompting the model to return verbatim excerpts, you can mimic the way a human researcher gathers information in order to reach conclusions more efficiently. This leverages an LLM called Claude for its ability to handle large documents up to 75,000 words. This human-in-the-loop approach combines the breadth of AI with human judgment.

Example prompts:

- *Attached is the text of Secure Act 2.0. I'm researching the impact it has specifically on 401(k) plans. In a bullet point list, provide the most relevant verbatim excerpts from the document, with a citation at the end of each quote to the exact section to find the excerpt in the bill.*
- *Attached is proposed legislation HR3937. I'm researching the impact it has specifically on 1099 reporting. In a bullet point list, provide the most relevant verbatim excerpts from the document, with a citation at the end of each quote to the exact section to find the excerpt in the bill.*
- *Attached is the text of IRS Revenue Procedure 2023-8. I'm researching whether there is any guidance surrounding audit protection in the rev proc. In a bullet point list, provide the most relevant verbatim excerpts from the document, with a citation at the end of each quote to the exact section to find the excerpt in the bill.*

> **We just have to start. Find a way to make the use of AI a habit."**
>
> **— *Jason Staats,*** *Founder, Realize*

Advisory Use Case: AI-enriched KPI Analysis

The programming capabilities of ChatGPT's Code Interpreter model make it well-suited for analyzing large volumes of data. This can be particularly useful in advisory contexts when evaluating a dataset against key performance indicators (KPIs). Code Interpreter can help ensure nothing gets overlooked by thoroughly examining the full breadth of the data, even aspects that may not be directly captured by the predetermined KPIs.

Specifically, Code Interpreter could be prompted to review the complete dataset and identify any notable trends, patterns or insights that the existing KPIs fail to highlight. Based on this analysis, Code Interpreter could then suggest potential new KPIs to better encapsulate these additional findings. Taking advantage of ChatGPT in this way enables the development of KPIs that more fully and accurately reflect the meaningful results within a dataset.

Example prompts:

- *Attached are several accounts receivable reports. The company has two accounts receivable KPIs: Days Sales Outstanding, which was 41 this month, and Average Days Delinquent, which was 6. Review the data to identify any advisory insights that these two KPIs may not take into account and, if there are any omissions, how these KPIs could be improved.*
- *Attached is the general ledger detail from last month. I would like you to analyze three KPIs: Salaries as a percentage of revenue, which was 20% this month, Asset Turnover, which was 0.79, and Operating Cash Flow, which was 1.29. Review the data to identify any advisory insights that these three KPIs may not take into account, and if there are any omissions, how these KPIs could be improved.*
- *Attached is a set of sales reports. The sales team has two KPIs: Average Purchase Value, which was $19,582 this month, and Lead-To-Close %, which was 17.52. Review the data to identify any advisory insights that these two KPIs may not take into account, and if there are any omissions, how these KPIs could be improved.*

Audit Use Case: Interview Transcript Assistant

Audit software centralizes workpaper documentation, but important client conversations often are not captured for the workpaper file. Now, AI meeting assistants can provide near real-time, accurate meeting transcriptions. By leveraging interview transcripts and using large language models to search them, audit teams can quickly find specific information across extensive transcripts.

Example prompts:

- *In this transcript, was anything discussed surrounding inventory receiving procedures?*
- *In this transcript, was an update provided on the confirmation letter from National Bank?*
- *In this transcript, were user permissions for the accounts receivable module discussed?*

Internal Business Use Case: Recruiting Agent

AI agents use LLMs to reason through tasks. Some are even capable of navigating applications. Human users prompt the agent to perform a task, and the agent plans out how to complete it, critiques its own plan and continuously course-corrects until the task is complete. This use case leverages Hyperwrite's Personal Assistant agent (a Chrome browser extension) to navigate hiring websites to identify qualified candidates for an open role at the organization.

Example prompts:

- *Navigate this job board and find me 10 applicants whose qualifications are similar to those of Jane Smith. They must be CPAs with a minimum of 5 years experience.*
- *Review John Smith's connections on LinkedIn to identify any potential candidates for the position I'm hiring for. I'm looking for junior- and senior-year students studying accounting within 50 miles of Austin. Prioritize candidates who actively engage on LinkedIn. Once you've found the three best candidates, draft a message to John for me to review, asking John to make an introduction.*
- *Locate the following individuals on LinkedIn: [list]. Identify if they're connected with anyone in my network. If so, provide a couple people in my network who could make an introduction to each individual.*

Join the conversation

These are just a few ways that accounting and finance professionals can use the power of generative AI. What are some ways you would train the platform to be more resourceful for your organization or clients? Join the conversation at **#DigitalCPA**.

Challenges and limitations of generative AI

The wide-open generative AI universe promises a range of possibilities, but with it lies challenges. Here are some of the current limitations and security risks that firms need to consider when looking to integrate the technology into their practices.

Security is paramount

As the security checklist in the next section makes clear, generative AI poses a number of challenges that firms must be prepared to mitigate. The newer and more popular an open-source solution is, the less mature its security is. It's especially important to address security concerns when working with vendors whose products embed generative AI because, in the headlong rush to bring new tools to market, some may not have addressed all security issues or be working in private data pools. Systems can be especially vulnerable because:

- Bad actors can use generative AI to guess passwords, enhance phishing attempts, perform CAPTCHA-cracking and build better malware.
- Hackers can invade a company's system through unsanctioned generative AI tools used by employees.

Questions to consider if the firm is working with a SaaS provider:

- What protections will be in place if staff upload client data?
- Is data identifiable?
- Is the data encrypted?
- Where is the data stored?
- How is data shared with public LLMs? As a general rule, no client or business information should be entered into a public LLM.
- What kind of verification is done on facts and math calculations?

Communication and critical thinking skills will matter

Developing communication and critical thinking skills has long been a priority for firms, and they will be especially important when using generative AI. Since the technology is language-based, firms will have the best outcomes if their people have the necessary skills to review and improve the content that the technology creates and spot incorrect nuances, biases or gaps in logic included in the responses. They should also be able to identify, for example, that a stiffly written communication is not aligned with a firm's tone and client relationship, or that an email or blog is too long or doesn't contain enough detail to be valuable to clients.

> **"Today's LLMs produce generated output, not computed answers, which means you shouldn't trust them with math or financial analysis. Expect this to improve in the next 12-18 months."**
>
> — ***Jeff Seibert,*** *CEO, Digits*

Prompts are important

Generative AI is only as good as the prompt that drives it. With that in mind, some organizations are hiring prompt engineers—people with the skills and backgrounds needed to communicate with generative AI systems, according to the Wall Street Journal. Prompt engineering will come in handy at firms of any size. Zunie Nguyen's seven-person firm, Yogi CPA, has provided a training session on prompt engineering, sharing useful examples and offering guidelines on how to write new ones.

Among the best practices for prompts:

Be specific

Don't just ask the system to write a client proposal. Instead, ask it to set out a specific number of points for the proposal using information from particular sources such as recent presentations your firm has given or available client or industry and market information. Providing the model with examples helps get a better output.

Be more specific

If you don't like the first answer, decide if you need something that's shorter or more detailed or that is written in a different tone. You can use terms that are broad like "more relaxed" or targeted (e.g. "in the tone of my last letter to client X"). A perfect example of this is our own message from ChatGPT at the opening of this toolkit.

Have a dialogue

The systems will keep delivering new iterations based on your instructions, so don't give up if the first couple of tries don't work. You can ask the system to tell you what information or instructions it needs to create content—such as a proposal or cover letter—and you can then supply the details it needs.

The great equalizer

McKinsey predicts that by 2030, activities that required up to 30% of the hours currently worked across the U.S. economy could be automated, a shift that generative AI will accelerate. The consultants expect, however, that generative AI will enhance the way that business professionals work rather than replacing many jobs in this area outright.

> **AI will make you a superhuman, you will become even more valuable to clients. The only problem is if you don't evolve yourself and use it."**
>
> ***— Pascal Finette,*** *co-founder, be radical*

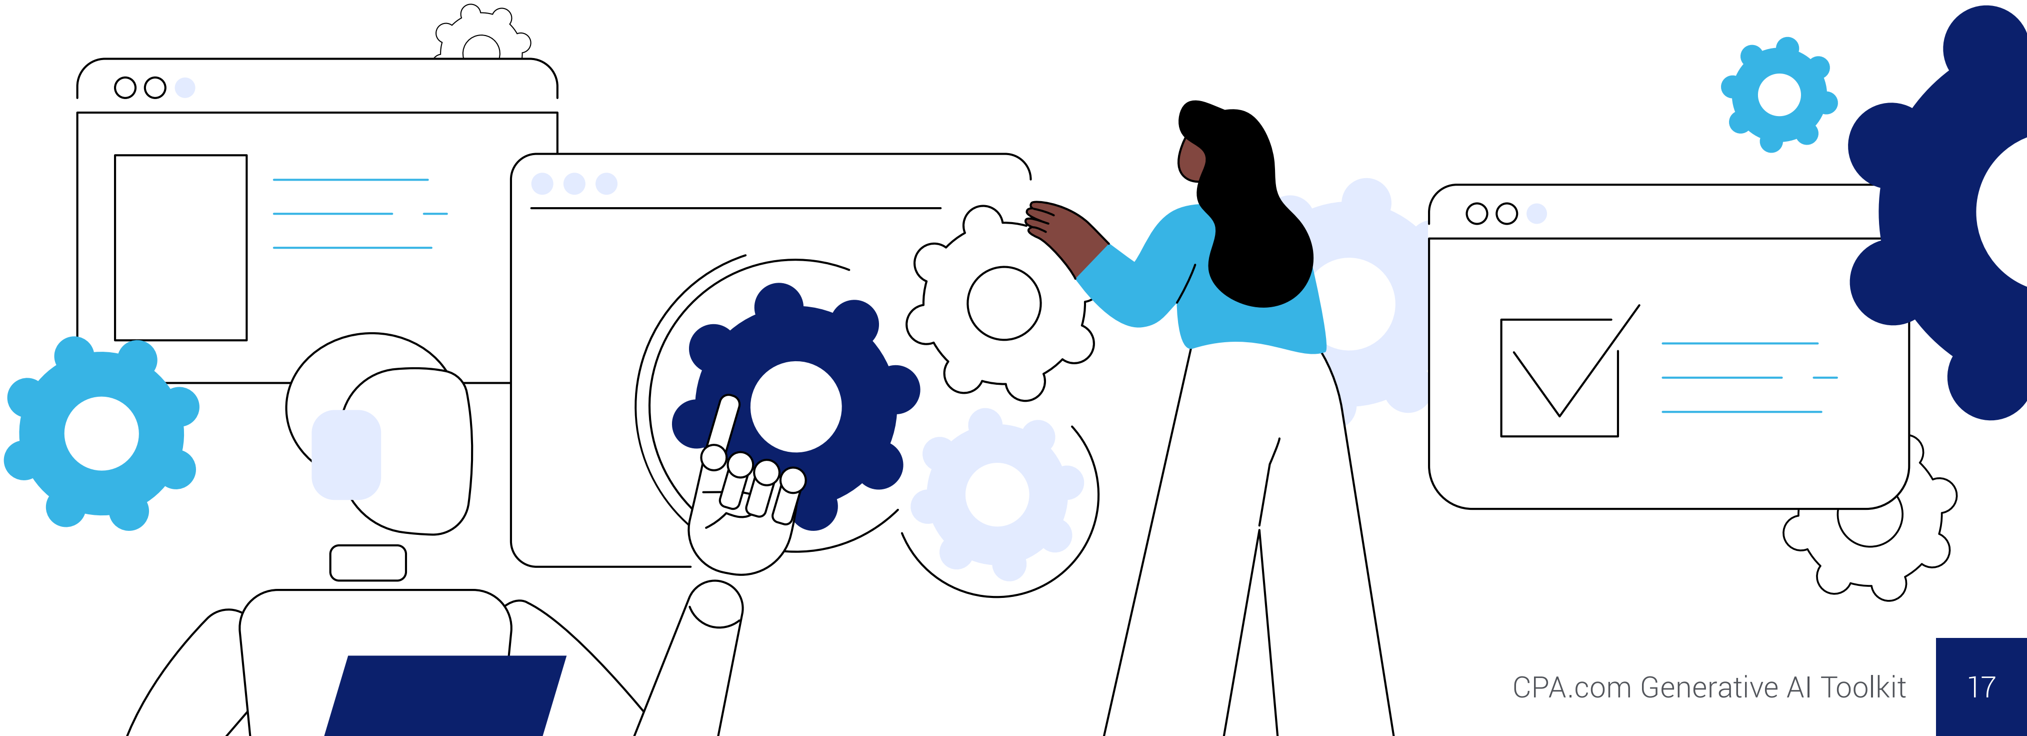

Security and risk considerations:
Safeguarding validity, privacy and transparency

As the utilization of generative AI continues to expand, understanding the associated security implications becomes crucial for businesses across industries.

There are 5 key risk considerations for the accounting profession, both for internal and client use:

1 Privacy

a. Understand the information you're ingesting into AI and its sensitivity. Public tools such as ChatGPT often rely on user input for its own training.

b. "De-identify" (sanitize) personal information before ingesting it into both internal AI systems and public tools. Personal identifiable information is anything that can be used to identify an individual person.

c. Have a conversation with the provider of your enterprise AI platform to exercise control over usage, retention, and disposal of information stored on the platform.

2 Transparency

a. Engage legal counsel to identify any laws that may apply which require disclosure use of AI.

b. If a client's confidential information will be shared with a third party vendor, including AI, provide written disclosure to and obtain specific consent from the client in the appropriate format before the confidential information is shared.

c. If client information is used to train the AI model, then consider explicitly stating such in your disclosure to and consent from the client.

d. If disclosure is not specifically required, consider still disclosing a firm's use of AI in the engagement letter. Include AI in the disclosure if it is anticipated AI tools will be used to render the agreed upon services.

e. An addendum to an engagement letter or other written consent from the client can be created if it is later determined that AI tools will be used.

CPA.com™

3 Avoiding Bias

The way in which we ask questions are often biased, so sometimes an AI can respond to a question in a way which confirms our bias. We need to be thoughtful about both inputs and outputs (responses).

4 Human Review

a. Accounting professionals have a responsibility to monitor answers that are coming out of generative AI to ensure accuracy.

b. Discuss with general counsel necessary documentation of a review process for AI output.

5 Limiting Use Cases

a. Begin with the firm's existing "acceptable use policy," which deals with access and use of company hardware and software, VPN access, client privacy concerns, etc.

b. Define ordinary use for AI tools, such as writing marketing language vs. extraordinary use such as using AI to produce counsel or professional advice.

c. Ask, who in the firm should have access to AI tools.

Additional risk considerations, examples and recommendations for mitigation:

Data Exposure and Leaks

The transformative capabilities of generative AI have many risks. One primary concern is the exposure of sensitive information:

Employees Exposing Confidential Information

Case Study: In March 2023, Samsung employees inadvertently exposed confidential data to OpenAI via ChatGPT. This included source code, internal meeting notes, and product roadmaps. ChatGPT uses these inputs to further train their models, which creates the risk that this specific information might leak to the wider ChatGPT user base. Samsung has since banned the use of ChatGPT and other generative AI systems on company-owned devices and internal networks. The leaks have raised concerns about the security risks of using AI tools in the workplace without proper safeguards.

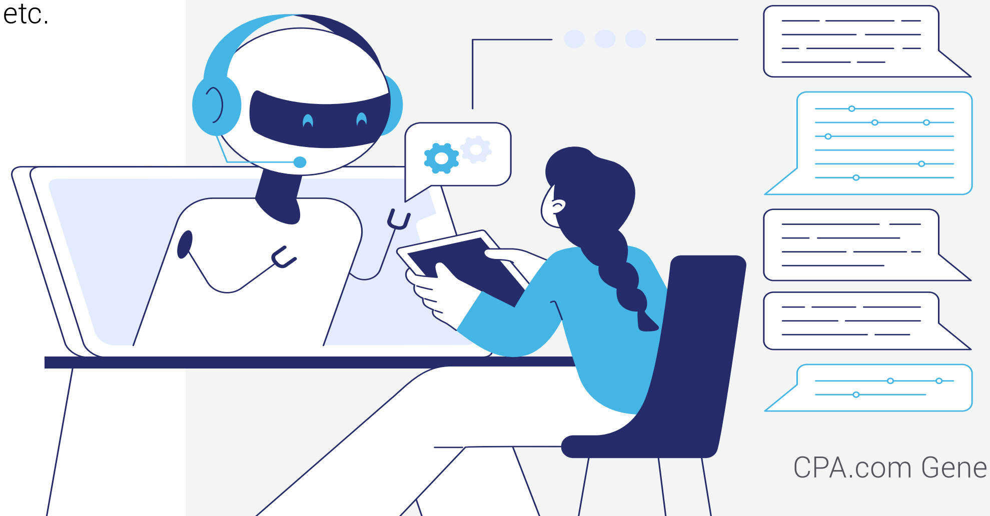

Mitigation: Reviewing the terms of use for the AI systems – some vendors explicitly do not train their models on your data, for others it is part of the T&Cs. Clear policies and procedures in place for the use of AI tools. Regular audits of interactions with external systems, coupled with stringent access controls, to prevent unintentional data sharing.

Data Leakage from LLM Responses

Large Language Models can sometimes disclose sensitive or proprietary information in their outputs. This may result in unauthorized data access, privacy violations, and potential security breaches.

Mitigation: A review layer to vet AI responses, along with the usage of data sanitization protocols and implementing strict user guidelines, can significantly reduce the risk of unintended data disclosure. It's worth noting that some vendors start to indemnify their users from any claims (e.g. Adobe with their visual GenAI).

Insufficient Access Controls

An oversight in the deployment of LLM models can lead to unauthorized users gaining access and model and data exposure.

Mitigation: Multi-factor authentication, role-based access controls, and periodic access reviews ensure only authorized personnel can interact with the AI systems.

Infrastructure Vulnerabilities

The intricacy of enerative AI systems can sometimes introduce specific vulnerabilities that attackers can exploit:

Prompt Injections

Attackers can supply crafty inputs to manipulate LLMs, causing unintended actions or even bypassing filters.

Mitigation: Rigorous input validation and filtering mechanisms. Consider the use of a whitelist approach where only certain predefined types of inputs are permitted.

Inadequate Sandboxing

AI models, when inadequately sandboxed, can allow unauthorized access to the underlying IT environment. Consequences can include:

- Insecure Output Handling
  When an LLM output is accepted without scrutiny, exposing backend systems.
- Unauthorized Code Execution
  Using malformed natural language inputs to execute malicious code or commands.
- Server-side request forgery (SSRF) Vulnerabilities
  Accessing restricted resources, internal services, APIs, or data.
- Exploitable Plugin Design
  Some LLM plugins can have insecure inputs and insufficient access control.

Mitigations:

- Deploy AI models within isolated environments, ensuring they don't have unwarranted access to external systems.
- Regularly update and patch systems to defend against known vulnerabilities.

Improper Error Handling

Errors, when not properly masked, can reveal sensitive data, including personal information.

Mitigation: Customize error messages and debugging output to ensure they remain generic and non-revealing. Conduct thorough testing to identify and rectify information leak points.

Model Denial of Service

Given the resource-intensive nature of generative AI models, attackers might initiate resource-heavy operations causing service degradation or high costs.

Mitigation: Rate-limiting and monitoring systems to detect and alert on unusual activity patterns.

Data Integrity Threats

Generative AI models are only as good as the data they're trained on. Any compromise in data integrity can lead to serious ramifications:

Training Data Poisoning

Deliberate tampering with training data can introduce biases, backdoors, or vulnerabilities.

Mitigation: A secure and controlled environment for data collection and curation. Regularly review and validate datasets to ensure their quality and integrity. Employ anomaly detection to identify unusual patterns in training data.

Model Theft

Theft of proprietary AI models not only results in economic losses but also exposes sensitive information and intellectual property.

Mitigation: Encryption techniques for model storage and transmission. Use secure enclaves or trusted execution environments to protect models during inference. Additionally, utilize model watermarking techniques to trace unauthorized usage.

Compliance and Ethical Concerns

The wide-ranging capabilities of generative AI also introduce potential compliance and ethical challenges:

Over-reliance on Generated Content

Generative AI models can sometimes produce outputs that aren't grounded in their training data, "hallucinations". Excessive trust in LLM-generated content can result in misinformation, miscommunication, legal issues or security vulnerabilities.

Mitigation: Manual review processes for critical AI-generated content. Educate users on the potential risks and encourage a culture of verification.

AI Alignment and Excessive Agency

LLMs can exhibit behaviors misaligned with intended use cases, leading to unintended consequences.

Mitigation: Regularly test and validate AI systems against real-world scenarios to ensure alignment. Restrict excessive permissions and autonomously granted functionalities to LLMs.

Bias and Discrimination Risks

AI responses may sometimes demonstrate biases, risking violations of anti-discrimination laws.

Mitigation: Regularly audit models for bias. Use diverse training datasets and employ fairness-enhancing interventions during model development.

Intellectual Property and Copyright Implications

AI-powered tools are often trained on massive amounts of data and are usually unable to provide sources for their responses. Copyrighted resources, such as books, magazines, and academic journals, may be included in some of the training data. Using AI output based on copyrighted materials without citation could result in legal penalties.

Mitigation: Provide disclaimers when using AI-generated content. Implement fact-check mechanisms to detect potential copyright infringements.

Licensing Restrictions and Content Use

Another consideration involves using content that may be under a license or other agreement. The terms of these agreements might not explicitly permit or could even prohibit the use of such content with AI tools, either fully or without citation.

Mitigation: While this area remains legally ambiguous, it's worth consulting with legal counsel on this matter.

Legalities Surrounding Chatbot Use

Using AI-powered chatbots to answer customer inquiries without proper disclosures may invite legal penalties.

Mitigation: Always disclose AI interactions to users. Stay updated on regional laws and regulations concerning chatbot and AI use.

Open-Source Models Risks

Unvetted Code & Components

Open-source models and their infrastructure are more likely to contain code or components that haven't undergone rigorous security checks. Malicious actors can introduce vulnerabilities or backdoors.

Mitigation: Regularly audit the open-source code for vulnerabilities. Utilize static and dynamic code analysis tools to inspect for potential threats.

Inconsistent Updates & Maintenance

Some open-source projects might not receive regular updates, leaving them exposed to known vulnerabilities.

Mitigation: Ensure you're using actively maintained open-source projects. Regularly check for updates and patches, and integrate them promptly.

Exposure to Pre-trained Data

Open-source models, especially pre-trained ones, might have been exposed to unvetted data sources. This could introduce biases or make the model respond in unpredictable ways.

Mitigation: When possible, fine-tune the open-source models using your own vetted datasets. This can help in overriding potential biases introduced during the initial training.

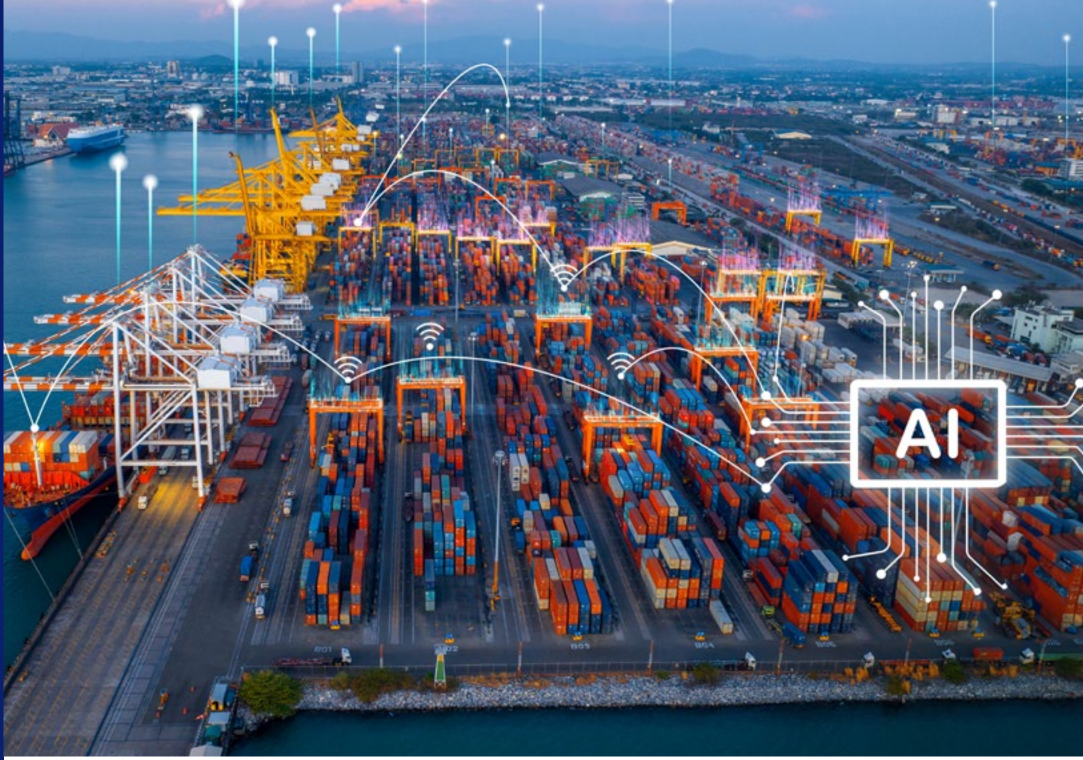

Additional Attack Vectors and Misuse

The vast potential of generative AI opens avenues for misuse, both in terms of external attacks and malicious applications:

Supply Chain Vulnerabilities

The lifecycle of generative AI can be compromised by vulnerable components or services. The AI ecosystem is rapidly evolving, creating multiple opportunities for malicious actors to attack.

Mitigation: A thorough vetting process for third-party datasets, pre-trained models, and plugins. Ensure a regular update and patching schedule. Conduct security audits on all integrated components.

Data Privacy Concerns

For enterprises who develop their own generative AI models, the process requires large amounts of training data, raising concerns if breaches occur and threat actors gain access to sensitive data, including personal information.

Mitigation: Robust data encryption, both in transit and at rest. Implement strict access controls on data storage locations, and ensure compliance with relevant data protection regulations for the usage, retention, and disposal of stored information.

Malicious Use of Deepfakes

AI's capability to create realistic fake video and voice cloning can be used to impersonate key personnel and obtain sensitive information.

Mitigation: Educate employees about the risks and signs of deepfakes. Implement detection tools that leverage AI to recognize manipulated media.

Personalized Social Engineering Attacks

AI can craft highly targeted phishing or scam messages, enhancing their efficacy.

Mitigation: Train employees and stakeholders on the latest social engineering tactics.

Mitigation Strategies

Beyond addressing specific vulnerabilities, a broader, proactive approach is essential to safeguard against evolving threats in the generative AI landscape:

Continuous Security Training

Educate teams about the evolving threat landscape and best practices in handling and deploying AI tools. A well-informed team can be the first line of defense against potential breaches.

Robust Infrastructure Security

While the focus might be on AI, ensuring that the foundational IT infrastructure is secure is paramount. Regularly update, patch, and monitor underlying systems to deter potential breaches.

Third-party Vendor Audits

Regularly review and audit third-party AI vendors. Ensure they adhere to security best practices and are transparent about their data handling and processing methods. Consider obtaining an assurance report from a licensed CPA firm.

Incident Response Plan

Prepare for potential security incidents by having a response plan in place. This plan should detail steps to contain the threat, communicate with stakeholders, and recover operations.

Feedback Loops

Establish feedback mechanisms with users of AI systems. They can often provide valuable insights into potential system misbehaviors or vulnerabilities.

CPA.com™

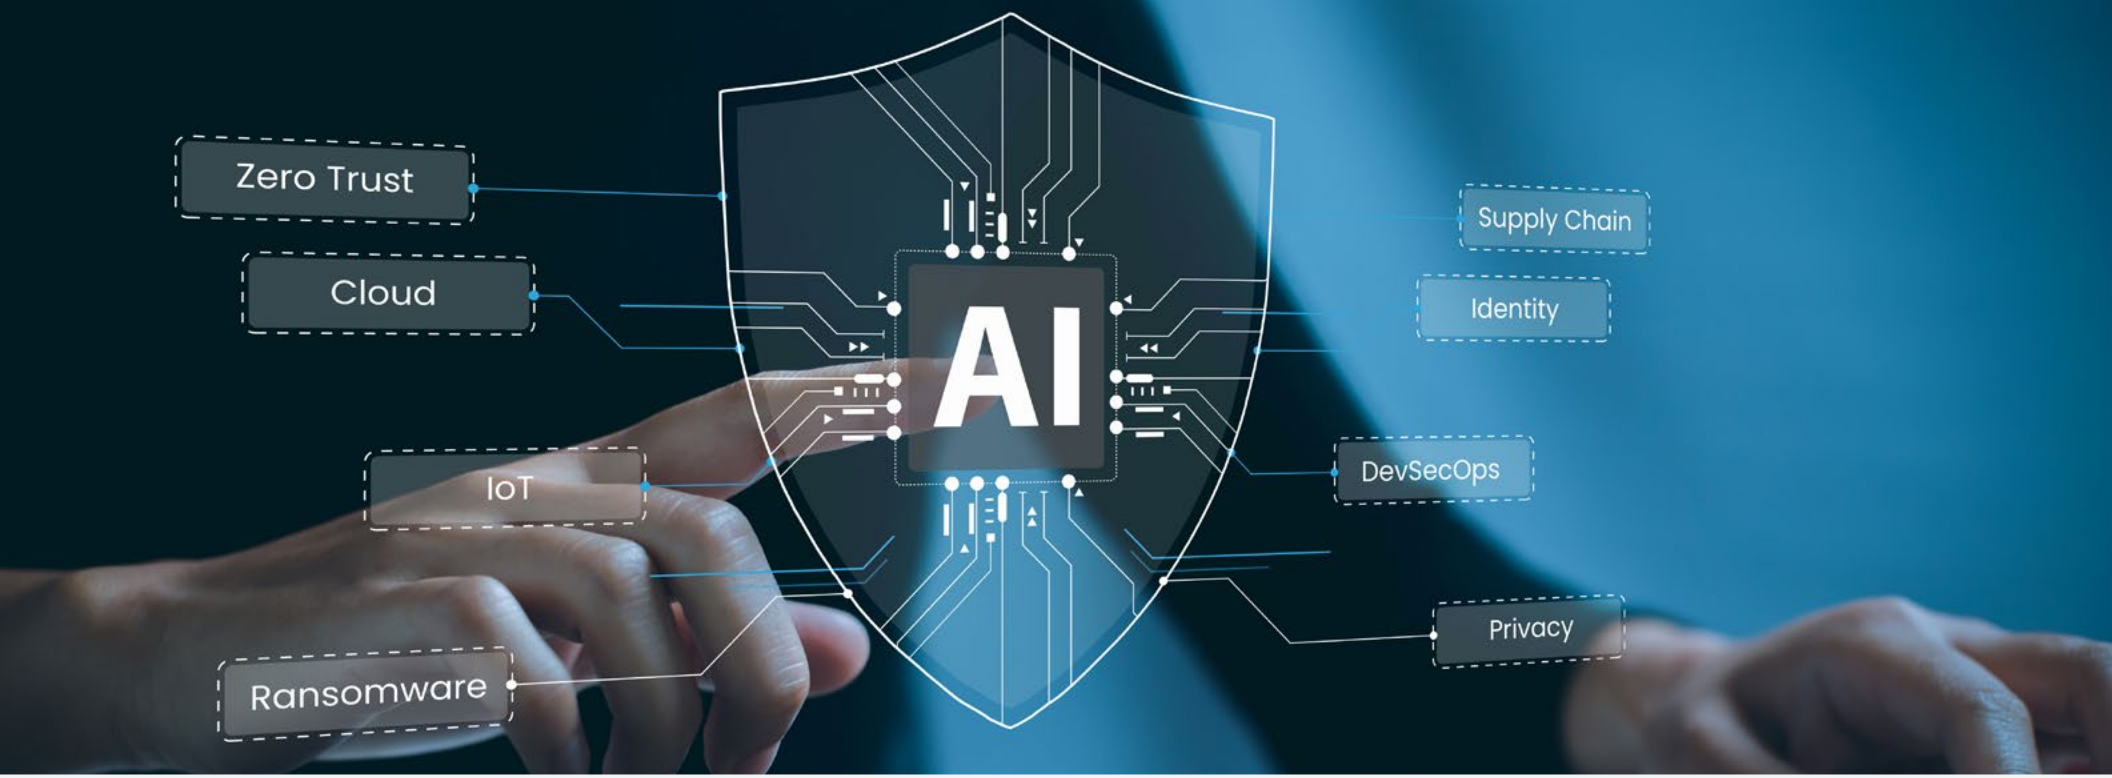

Limit AI Autonomy

For high-risk scenarios, consider limiting the autonomy of AI systems. Manual oversight or approval mechanisms can act as safety checks against unintended AI actions.

AI Adoption and Risk Levels

The risks associated with generative AI vary depending on a firm's engagement with the technology. Different threats emerge based on whether a company merely uses tools like ChatGPT for basic tasks, leverages vendor-specific applications for specialized operations, or actively develops and trains their own models using client data.

The Double-Edged Sword of AI

As generative AI tools enhance productivity and open new possibilities, they also present novel risks. It's imperative to approach them with a blend of enthusiasm and caution.

Hacker Productivity

Just as businesses benefit from advancements in AI, so do malicious actors. The tools and methods that enhance legitimate operations can also streamline illicit activities.

Traditional Cybersecurity

Embracing new AI-driven solutions should not come at the cost of sidelining traditional cybersecurity measures. Both old and new defenses must work in tandem to ensure robust protection.

Glossary of key terms

Here are a few key terms related to generative AI that will help you navigate the emerging tech:

Algorithms
In computer science, this refers to tools that can perform calculations, data processing and automated reasoning tasks.

Application Programming Interface (API)
A way for two or more software programs to talk to each other.

Artificial Intelligence (AI)
Computer systems that can perform tasks that would typically require human intelligence to complete.

Bots
A computer program that can interact with users or other systems. Since they are automated, they can follow specific instructions without continuing human oversight.

Data lakes
A central repository that can store large volumes of raw data in its native format, whether structured, semi-structured or unstructured, and a place to organize large volumes of data from a variety of sources. Data can be ingested faster into a data lake than into a data warehouse.

Deep Learning (DL)
This subset of machine learning learns from multiple layers of neural networks, enabling it to perform more complex tasks.

Generative AI (GenAI)
A type of AI that can create new content. GPT or generative pre-trained transformers. A transformer is a neural network that enables AI tools such as ChatGPT to generate content similar to what a human might create.

Hallucination
This can occur when an AI system is asked a question it can't answer, so it makes up false or nonsensical information instead. In one case, when asked to cite precedents to be used in a court case, the system responded with cases that did not exist. This tendency is one major risk in working with generative AI.

Large Language Models (LLMs)
The foundation of generative AI tools related to language, these models are trained on vast amounts of data, enabling them to create original output in response to prompts, such as text, images and computer code.

Machine Learning (ML)

Describes how machines can perform and refine tasks without specific instructions, using algorithms and statistical models to find patterns in the data.

Natural Language Processing

This makes it possible for LLMs to understand and respond to text or voice data, rather than computer code.

Neural Networks

Systems modeled on the workings of the human brain. Generative AI uses them to recognize data patterns and structures and then create new content based on what it learns.

Robotic process automation (RPA)

RPA can mimic human actions by taking on basic, repetitive tasks, such as data processing. generative AI, on the other hand, can simulate human intelligence, with an ability to learn and reason.

SCAN ME

Additional Resources

Access videos and additional resources at CPA.com/Gen-AI

About CPA.com

CPA.com's mission is to empower accounting professionals in the digital age—to be the advisor's advisor. Technology's impact on business model transformation and practice evolution is undeniable. It's creating new practice areas and driving rapid change across firms. We aim to help practitioners successfully navigate this change. We do this through numerous strategic initiatives focused on driving the transformation of practice areas, advancing the technology ecosystem for the profession and leading technology research and innovation efforts for practitioners.

CPA.com's research and innovation efforts include the CPA.com Startup Accelerator, ecosystem networking events, category symposiums and the curation of future-focused developmental and educational resources like toolkits and special reports intended to demystify emerging technologies within the context of their impact on the accounting profession.

A subsidiary of the American Institute of CPAs, CPA.com is also part of the Association of International Certified Professional Accountants, the world's most influential organization representing the profession.

For more information, visit CPA.com.

CPA.com™

© 2023 CPA.com. All rights reserved. CPA.com and the CPA.com logo are trademarks and service marks of CPA.com. The Globe Design is a trademark of the Association of International Certified Professional Accountants and is licensed to CPA.com. All rights reserved.