Advising on Agentic AI: A Simple Framework for In-house Counsel

Technology, Privacy, and eCommerce

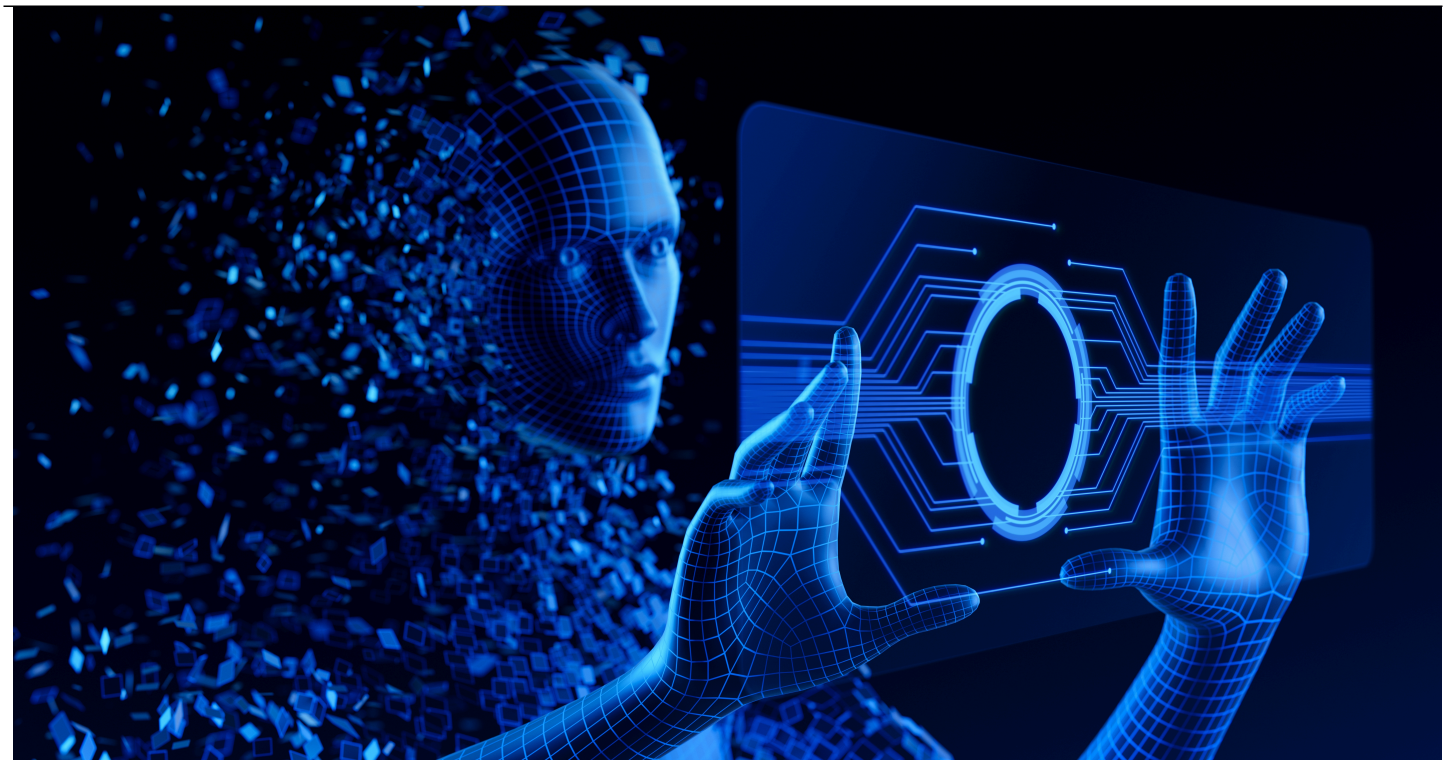

Banner artwork by Alexander Supertramp / *Shutterstock.com*

The Billows Feeding Machine.

Modern Times — Charlie Chaplin's 1936 classic — lampoons our reliance on technology in the workplace. In the opening scene, we see the Billows Feeding Machine in action, an elaborate gizmo that automates meal-eating for workers on the assembly line. "Eliminate the lunch hour, increase your production, and decrease your overhead," the sales pitch promises the factory boss.

The pitch doesn't pan out. When tested on Chaplin, the machine goes haywire.
It blindly speeds through its operations, shoveling steak into the Tramp's still-full mouth, smashing cake in his face, and sawing at his teeth with corn on the cob.

The gag still works, though the gadget looks old-fashioned. The once-futuristic robot
seems primitive next to the robots of *our* modern times — those driven by artificial intelligence.

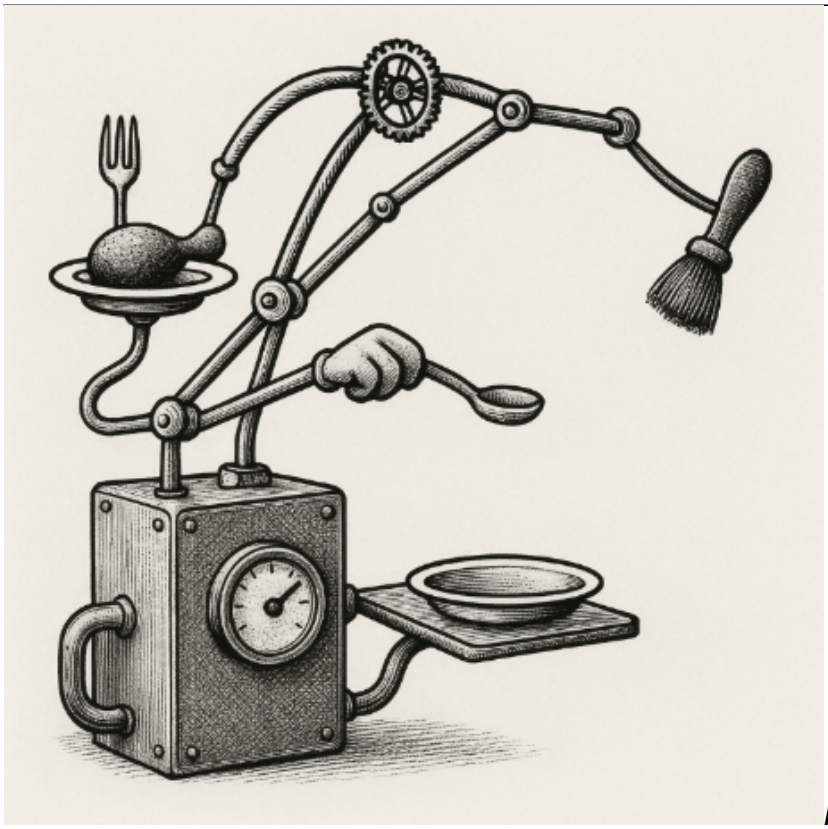

Image created by author / ChatGPT

Today's AI, at its most robot-like, takes the form of "AI agents." Agents, says IBM, are "machine-learning models that mimic human decision-making to solve problems in real time." They can autonomously reach out to sales targets, research and summarize information, screen and contact job applicants, you name it.

It makes an in-house lawyer want to throw their computer out the window. Just when you've grasped the legal issues with "regular" AI, an endless variety of agentic AI applications can seem overwhelming to analyze.

The Billows Feeding Machine offers a framework.

Wondering how AI will affect your legal profession?
Check out the **ACC AI Center of Excellence for In-house Counsel** today.

Agents = automation.

"AI agents" are recent innovations for performing complex tasks. But what agents do, fundamentally, is neither new nor complicated. In fact, it's captured by a word whose origin dates back to the Gilded Age: "automation."

Sure, agentic AI is far more complex than the metal arms and doodads in the Billows Feeding Machine. But step back a bit and you'll see a similar phenomenon: "Human did task. Now machine does task."

Recognizing that similarity represents the first step in analyzing the legal issues that agents

pose. Those issues aren’t unique to AI — as an example show.

Automation isn’t automatic.

Imagine you need to give a robot instructions for pluralizing English nouns.

Your initial direction may be to just affix an “s” at the end of each noun. But of course, that doesn’t do. “The class watches the teachers and child” becomes “The classs watches the teacherss and childs.”

The more you think through the task, the more nuance you see. Existing plurals don’t need the “s.” Singular nouns ending in “s” require “es” instead. Verbs change when nouns are pluralized. Some nouns have irregular plurals. And so on.

Eventually you draft the perfect instructions; the robot generates the right results. But, importantly, you know that’s the case only because you know what “right results" look like.

> Eventually you draft the perfect instructions; the robot generates the right results. But importantly, you know that's the case only because you know what "right results" look like.

If you didn’t know the English plural rules, the robot could spit out gobbledygook and you'd be none the wiser.

AI is smart. And that’s not enough.

Agentic automation doesn’t require the granular, rule-based instruction demanded by the pluralizing robot. AI is smarter; a plain-English prompt to pluralize inputs may produce the results you want.

But the example’s point still holds: It’s dangerous to automate a task unless you first appreciate the task’s complexity.

It’s a lesson well known to those automating driverless cars and surgical robots. But it’s not necessarily realized by designers of AI agents. That’s because anyone with internet access can design an AI agent.

> It’s dangerous to automate a task unless you first appreciate the task’s complexity.

For an organization, the ease and accessibility of creating AI agents raises unique legal risks. Here’s how to tackle them.

First questions.

When your non-legal colleague asks for advice on an AI agent they built, the critical threshold questions aren't about AI — for instance, the model used — or even about the data at issue. The first questions are simply: Is this the sort of task that should be automated? And if so, is this the best human being to oversee that automation?

These aren't "legal" questions. They force in-house counsel to take a step into operations, understanding how each part of the organization works. The lawyer needs a firm grasp of the organization's departments and roles, its decision-makers and subject matter experts. And the lawyer must be prepared to loop those people to analyze the agent.

Assume, for instance, your company hires a new, AI-savvy intern. Eager to help, the intern crafts an agent to draft and publish the company's pesky quarterly reports.

Your first thought won't be about Section 13 of the Exchange Act.

Instead you'll think: This isn't a job to automate senselessly. And if we are automating it, this intern isn't the person to manage it. You'll work with your corporate team to assess.

Disassemble the machine and look inside.

Depending on the agent's responsibilities, your org's stakeholders and experts may decline automating the task, or at least limit which steps to automate. But if they're on board with incorporating agents, then it's time to dig into the legal issues.

There's no one-size-fits-all approach. The variety of agentic tasks is limitless, and so too are the legal issues and the questions you can ask.

Still, common to all legal analysis is to first lay out the tasks that the agent will accomplish. That doesn't mean a line-by-line technical review of code; that's not feasible. But it does mean examining each step in the automated process. For each one, find out what data the agent may access, what systems and persons it may interact with, what actions it will take, and what decisions it will make.

Consider an agent that assesses and sets up interviews with select job candidates. You'll want to ask what info the agent processes (resumes? emails? material found online?); how it evaluates that information (does it rank the candidates?); what it then does (proposes candidates to interviewers? reaches out directly?); how it schedules interviews (email? calendar invite?); and whatever other details you need.

Once you unpack the steps, you can then position to analyze the legal matters involved.

Whatever those matters are, ensure that someone in your business closely monitors the agent, especially before it makes important decisions — like booking interviews with candidates based solely on a robot's read of resumes. This "human in the loop" principle is reflected in emerging AI laws and in Article 22 of the GDPR, which protects data subjects from purely automated decision-making. States and other regulators are also increasingly regulating AI use in employment, so legal compliance is a must.

You'll also need a process to audit agents' work — making sure they're performing as expected. In *Modern Times*, we know that the Billows Feeding Machine malfunctions because a crowd of people witness its failure. If no one is watching an AI agent plug away, it may be smashing cake into faces without anyone knowing.

> You'll also need a process to audit agents' work — making sure they're performing as expected.

Conclusion

Many businesses see agentic AI as the ultimate tech for organizational efficiency — you can automate everything. Their instinct then is to democratize agent creation across the workforce, letting all employees spin up and share agents.

Businesses should think very carefully before doing so. Just because technology is powerful doesn't mean everyone should access it. Often it means the opposite.

Companies should sooner consider access controls that limit who can create AI agents and who can disseminate those agents to others in the organization. Even if an agent's originator understands the automation and employs the agent prudently, other users may be less informed and less careful.

You may quickly have an organization full of Billows Feeding Machines.

Join ACC for more AI insights!

Christopher Wlach

Senior Director, Legal

Acxiom

Chris Wlach is the senior director, legal of Acxiom. Before moving in-house he focused on complex commercial litigation at Arnold & Porter. He is a certified information privacy professional (CIPP/US) through the International Association of Privacy Professionals. He also chairs the board of HEART, a humane education nonprofit.