



Generative AI Professional Usage and Perception Survey

Amid an emerging gap, knowledge is power



Executive Summary

Generative AI has made quite an impact in the year since OpenAI's ChatGPT gained prominence. But, to paraphrase Jim Collins in *Good to Great*, if generative AI is the bus, business leaders still need to make sure the right people are on it, people are in the right seats, and leaders know where to take it.



At Contentful, our goal is to put our customers in the best possible position to make use of digital content and any technology that helps them engage and communicate with the audiences they care about. It's clear that generative AI (or genAI, as we'll refer to it) falls into that category. As we continue to anticipate its capabilities and build them into our own product roadmap, we want to better understand the context and attitudes that shape our customers' priorities and usage.

In our first survey of 820 professionals in a range of technical and non-technical roles around the world, we endeavored to understand how those would-be bus riders see the opportunities and potential drawbacks of genAI.

For us, understanding how people across a business feel about genAI provides essential context. Our customer base comprises a wide range of roles and departments: marketers and other business users, designers of various descriptions, digital strategists, developers, and engineers.

We surveyed a diverse range of roles to understand how differences and similarities among them might influence the ways they use genAI. We sought, in particular, the views of people who are experienced enough to have a perspective beyond just their own roles, but who are not so high-ranking that they were removed from the details of day-to-day work going on.

What we found: More than geographic differences or distinctions between those in technical versus non-technical roles, **perhaps the significant characteristic that marked the most meaningful differences between respondents was the level of self-reported knowledge of genAI.**



There is a significant gap between the people who consider themselves highly knowledgeable on generative AI and everyone else.

Regardless of whether they work in [technical](#) or [non-technical roles](#), the people who consider themselves highly knowledgeable about genAI are the ones working the most with it. They are experimenting and identifying how and where it makes a productive impact on their work and target objectives. This is where the real value of genAI is being uncovered.

Despite the differences, the knows and know-nots strongly agree on several subjects, including the need to disclose use of genAI, that these new tools will require learning new skills, and that they want the ability to turn genAI capabilities on and off.

Over three-quarters of respondents have paid access to genAI tools at work. Only 24% of people in our survey don't pay to use genAI tools at work – and nearly as many pay out of their own pockets (either entirely or on top of what their employers fund) to do so.

Most people – especially those with less knowledge of genAI – want more guidance on how to use it responsibly.

We're on the cusp of a wave of tailored genAI tools: two-thirds of respondents indicated that their businesses either already have plans for some kind of tailored large language model (LLM) or are considering them. Of those who have or are considering plans, the trend is slightly more in favor of applying an existing LLM rather than training their own.

Sticking your head in the sand and hoping genAI will go away is not an effective strategy – for individuals or businesses. The ways in which genAI is already changing how many people work point to a potentially fast-growing divide between the businesses that empower their employees to make use of these tools and the ones that do not.

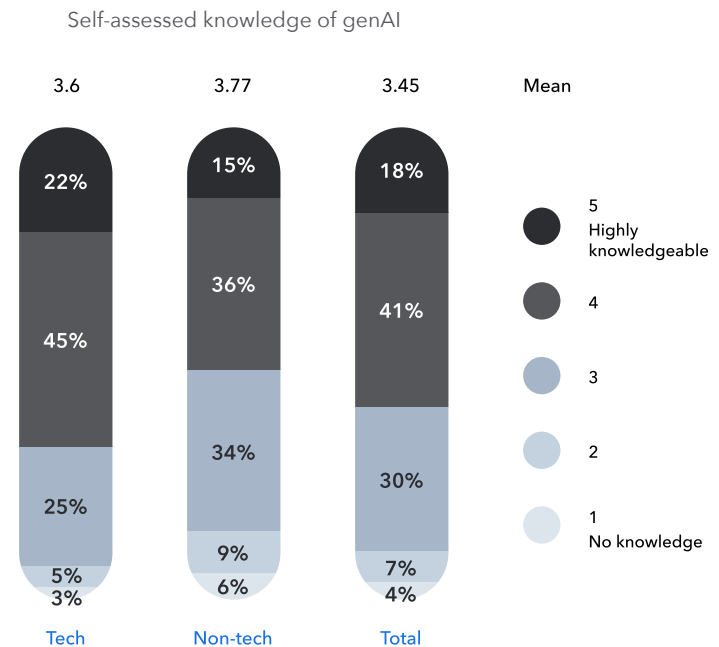


Experience, attitudes, and perceptions

Everyone's a genAI expert

There's an inherent bias in our data set – people who feel they understand a subject are more likely to respond to a survey about it – but it's clear that most of our survey respondents feel they have some level of expertise in genAI. Already! Nearly a fifth (18%) rated themselves a "5" in their knowledge (on a scale of 1 to 5) of generative AI and 71% gave themselves a 3 or 4.

Technical respondents are more likely than non-technical respondents to have rated themselves a 4 or 5 in terms of genAI expertise. (Also, Americans and males were more likely to do so – in line with general trends across all kinds of surveys.)




Total n=820, Tech n=390, Non-tech n=430

By role


Technical respondents are more likely to use genAI tools more frequently, specifically on a daily basis, both professionally and personally. Non-technical respondents are more likely not to have used genAI, to have only tried it a few times for personal use, and to use it several times a month professionally.

Frequency of professional genAI use



Total n=820, Tech n=390, Non-tech n=430

Frequency of personal genAI use




Total n=820, Tech n=390, Non-tech n=430

By knowledge


Perhaps unsurprisingly, those who consider themselves extremely knowledgeable about genAI also tend to use it the most often. Similarly, those who profess to have little knowledge of genAI are very likely to say they have either never used it or have only tried it a few times.

Frequency of professional genAI use



High AI Knowledge n=148, 4 n=332, 3 n=244, Little AI Knowledge n=96

Frequency of personal genAI use



High AI Knowledge n=148, 4 n=332, 3 n=244, Little AI Knowledge n=96

Of those who have not used genAI professionally or personally, we see a mix of skepticism, concern, and lack of knowledge or opportunity

Respondents who said they hadn't used genAI either professionally or personally were asked to explain why in an open-ended response. Of the 159 respondents who indicated they hadn't used genAI either professionally, personally, or at all, the most common reasons include no interest or need, lack of knowledge, concern or fear, and lack of opportunity (see data on page 7).

Another 12% indicated that they had either used genAI in one context or another or that they were about to start using it, typically in a professional context. In this latter group, several indicated that they were waiting for their companies to develop guidelines or policies on how to use genAI. Here's one such example: "We are still in [the] process of finding a proper way of using AI in our work model so that it will be helpful rather than a liability."

A similar proportion of respondents either didn't have a reason or weren't sure why they hadn't yet used genAI. Other reasons included criticism of the capabilities or output, data protection or privacy issues, and lack of guidance.



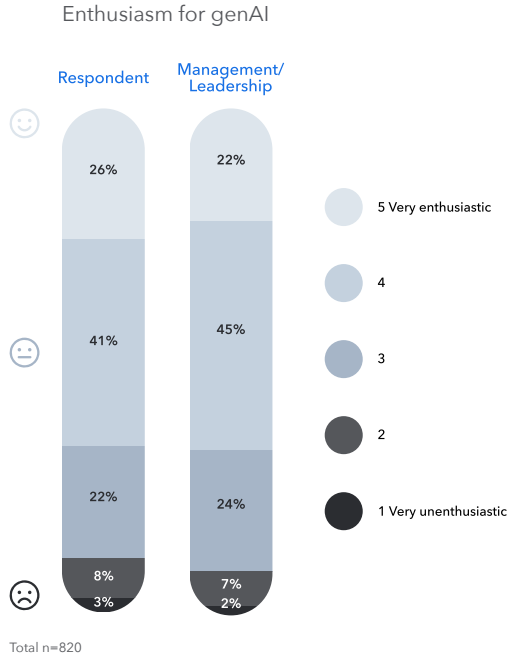
"We are still in [the] process of finding a proper way of using AI in our work model so that it will be helpful rather than a liability."

Most frequent reasons for not using genAI



Overall, enthusiasm for genAI is strong

Across the entire survey population, our respondents are enthusiastic about genAI (3.79 mean and 4.0 median, SD 1.0) and perceive the enthusiasm of their employers' leadership to be on par with their own (3.77 mean and 4.0 median, SD .94). These top-level results surprised us somewhat – we'd expected that there might be a discernible gap between how individuals feel about genAI and how they rate their company's enthusiasm, but in general there really wasn't one.



We hypothesized that there might be significant variations in enthusiasm for genAI across different geographic regions, but none emerged. Across regions, there were only minor differences between respondents' levels of enthusiasm and perceptions of management's enthusiasm for genAI. Even differences between regions were quite small and could largely be attributed to variations in sample sizes, concentrations of respondent roles, and cultural effects on rating scales.

How enthusiastic are you about genAI?

	USA	Canada	Australia	Europe	Total
Mean	3.83	3.73	3.64	3.75	3.79
Median	4.00	4.00	4.00	4.00	4.00
SD	1.01	0.95	0.90	1.01	1.00

Total n=820, USA n=203, Canada n=104, Australia n=102, Europe n=360

How enthusiastic do you believe management and/or business leaders are about genAI?

	USA	Canada	Australia	Europe	Total
Mean	3.86	3.64	3.72	3.71	3.77
Median	4.00	4.00	4.00	4.00	4.00
SD	0.95	0.94	0.92	0.89	0.94

Total n=820, USA n=203, Canada n=104, Australia n=102, Europe n=360

By far, the most meaningful differences depend on levels of self-professed genAI knowledge. Those who are most knowledgeable also tend to be most enthusiastic – and they are somewhat more enthusiastic themselves than they perceive their managers or business leaders to be. That relationship flips at the other end of the spectrum. Respondents who indicated a lower level of knowledge perceive management to be more enthusiastic about genAI than they are themselves.

This highlights one of the most important findings throughout this survey: there are significant differences between those who consider themselves more knowledgeable about genAI – especially the highly knowledgeable – and everyone else. As we’ll examine further, this emerging gap is one businesses should identify and act on.

How enthusiastic are you about genAI?

	5 High AI knowledge	4	3	1/2 Low AI knowledge	Total
Mean	4.57	4.05	3.36	2.80	3.79
Median	5.00	4.00	3.00	3.00	4.00
SD	0.70	0.80	0.90	0.96	1.00

Total n=820, High AI Knowledge n=148, 4 n=332, 3 n=244, Little AI Knowledge n=96

How enthusiastic do you believe management and/or business leaders are about genAI?

	5 High AI knowledge	4	3	1/2 Low AI knowledge	Total
Mean	4.41	3.92	3.38	3.23	3.77
Median	5.00	4.00	3.00	3.00	4.00
SD	0.73	0.79	0.91	1.01	0.94

Total n=820, High AI Knowledge n=148, 4 n=332, 3 n=244, Little AI Knowledge n=96

Perspective from the Prof: Sam Maglio's take



Sam Maglio
Professor of Marketing and
Psychology, University of
Toronto Scarborough

Expertise: The rift between haves and have-nots

It hardly comes as a surprise that technical professionals have higher knowledge about generative AI than their non-technical peers. Their job might demand it. Their personality type might draw them to it. Things get interesting when this gap in knowledge takes them to very different places.

When they think about the future, the same technical professionals who already know a lot are the ones who think they'll need to learn even more skills in the future. It might be tempting to see this as counterintuitive: They already know so much, so the non-technical professionals should be the ones expecting a higher need in order to play catch-up. But that overlooks a snowball effect in the mind. One of the best predictors of future commitment is having made a [smaller commitment in the past](#). Having knowledge fuels the desire to learn more and more.

Technical professionals are a bit more enthusiastic about genAI, non-technical professionals a bit less, and everyone expects that others agree with them.

The enthusiastic experts think that business leaders are enthusiastic; those with less enthusiasm believe business leaders share their lukewarm attitudes. This is a perfect example of what psychologists call [naïve realism](#): People think there's an objective world out there, that they themselves see it clearly, and that anyone else with any sense sees things the same way.

This tendency makes technical and non-technical professionals think different things about other people. But who's right? The evidence favors the technical professionals. For proof, look no further than the fact that two-thirds of businesses have a vision for an [LLM at their company](#) in the works. We could congratulate the technical professionals for their edge in prognosticating. But we'd be better served as managers in getting the non-technical professionals up to speed. Failing to anticipate the future accurately could find them left behind. Which would be a shame, because [new research](#) suggests that they're the ones best poised to reap the benefits of generative AI at work.

The higher the knowledge, the greater the time saved with genAI

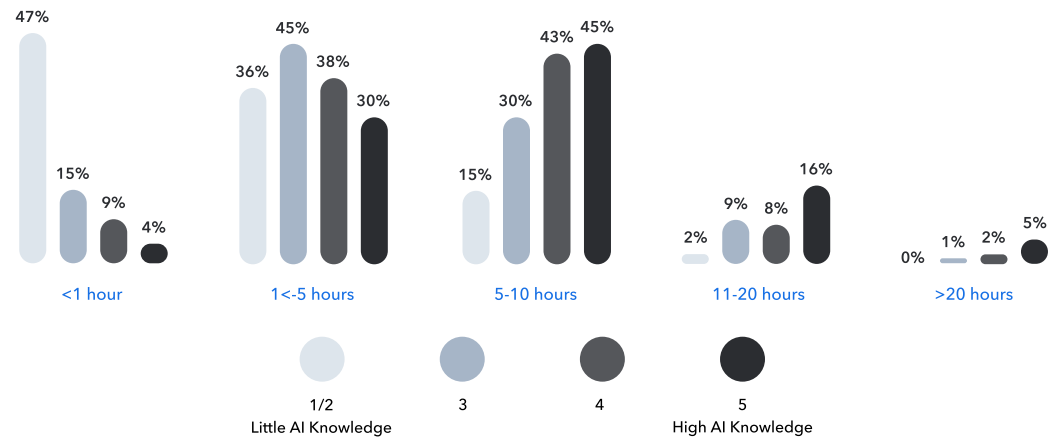
Thirty-eight percent of respondents say they save from one to almost five hours of time a week using genAI tools. An impressive 37% save between five and 10 hours per week and 11% save more than 10 hours per week.



By knowledge

Respondents with higher levels of genAI knowledge are more likely to save a greater amount of time per week using these tools, with 21% of the most knowledgeable saving more than 10 hours per week. By contrast, those who have the least knowledge save the least time.

Number of hours saved by using genAI




High AI Knowledge n=148, 4 n=332, 3 n=244, Little AI Knowledge n=96

By role

Levels of genAI knowledge were more significant in determining who saves more time using genAI than technical vs. non-technical job roles, though there are some differences there as well. Notably, respondents in technical roles are more likely to save between five and 10 hours per week and those in non-technical roles less than an hour.

Number of hours saved by using genAI



Total n=820, Tech n=390, Non-tech=430

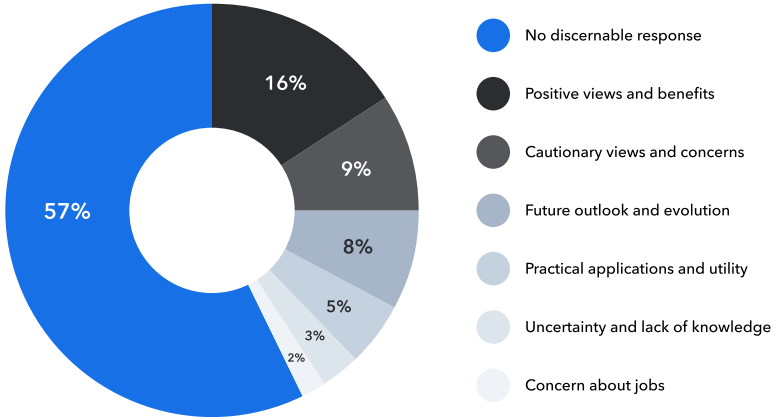
Positive views outnumber cautious, mixed, or negative views

When given the opportunity to share their perspectives on genAI or AI more broadly, most of our respondents demurred, but the overall sentiment ranged from cautious optimism to strong enthusiasm. Fifty-seven percent provided no response, said no comment, told us they didn't know, or provided an incoherent answer.

Of the 43%, or 356, who did comment, the majority (61%) were positive about genAI. Twenty-eight percent expressed concerns or had mixed views, with an overall sense that genAI is happening and is going to make a significant impact regardless of potential drawbacks. Another 8% expressed neutral views. Only 3% were explicitly negative in their views, ranging from it not being necessary in their organizations to genAI being dangerous or outright bad.

We classified the responses into positive views and benefits (16%), cautionary views and concerns (10%), future outlook and evolution (8%), practical applications and utility (5%), uncertainty and lack of knowledge (3%), and concern about jobs (2%). We also received two responses that had clearly been produced by ChatGPT. It's tough to say what's more interesting: the fact that two respondents took this approach, or the fact that we could so easily recognize that's what they did.

Perspectives on genAI





Access and appeal



When it comes to using genAI professionally, the vast majority of our respondents willingly pay to play (sometimes out of their own pockets)

Only 24% of respondents don't pay to use genAI tools at work. For the largest portion of our respondents, 37%, their companies pay for their professional use. Another 16% pay for it themselves and expense it.

But the most interesting group is the 18% of respondents who pay out of their own pockets – without expensing it – to use genAI tools for work. Add the other 5% who pay out of pocket for additional professional use beyond what is company-funded and **nearly a quarter of all respondents find these tools so valuable in a work context that they seem happy to put in their own money to access them.**


We struggle to identify a comparable technology development that individuals have been so eager to access that they've been willing to fund professional usage themselves. Mobile phones might come close but most of us with firsthand experience carried a work phone and a personal phone until employers started rolling out bring-your-own-device policies. Cloud computing may have had some parallels in the early days, but most of that use was billed on credit cards and expensed back to employers.

As in other parts of our survey, patterns vary depending on respondent segments. Respondents in technical roles are more likely to say their companies pay for their genAI usage; non-technical respondents are more likely to say they don't pay for usage. But there's no significant difference between the groups when it comes to paying and expensing their use, paying themselves, or paying for additional usage.



The more meaningful distinctions here seem to be determined by the level of genAI expertise. The people who rated their knowledge a 4 or 5 are the most likely to pay their own money to use these tools at work, whether or not they expense it, with the "5s" more likely to fund over and above employer-paid access. Given that these are the same groups that are most likely to be saving the greatest amount of time per week, it stands to reason that they would also be more willing to pay their own money for access to genAI tools.

The proportion of respondents who fund their own access to genAI tools for work purposes raises several questions: Do employers know that employees are using these tools? Is this use sanctioned? Does it follow established corporate guidelines or policies? These answers are beyond the scope of our survey – but something businesses should seek out for themselves.



Perspective from the Prof: Sam Maglio's take



Sam Maglio
Professor of Marketing and
Psychology, University of
Toronto Scarborough

Give me some of that - whatever it is

How much are people buying the hype around generative AI? Enough to put their money where their mouth is, that's for sure. Close to one in five pay out of their own pocket for access to these tools. The opportunity to use these tools at work makes an employer more enticing to over 60% of workers. If companies don't already have their own LLM, odds are they're working on one.

Nothing riles people up quite like **potential**, even if that potential might take a while to be realized. Between now and then, we see people behaving like they always do when faced with a hazy opportunity. They're terrified by FOMO. They want to make sure they get in on the ground floor. And this frenzied rush guides their investment strategy.

The overwhelming **share of usage for ChatGPT** makes it look like an index fund, seen as a catch-all that can benefit even novices. But people are also diversifying their assets. The average person uses **approximately three**, but people overall sample from a dozen or more different options.


This behavior – dipping a toe into multiple genAI pools – reveals that people want to explore, to learn, to keep their finger on the pulse of this new technology, or at least to hedge their bets. There's still plenty left to sort out among these providers. In the meantime, people navigate this uncertainty by doing what they've always done: **keeping their options open**.

Access to genAI tools at work is a net positive

Overall, most respondents view a potential employer’s decision to provide access to genAI tools favorably in choosing whether or not to take a job, with far more ambivalence than any negative impact.

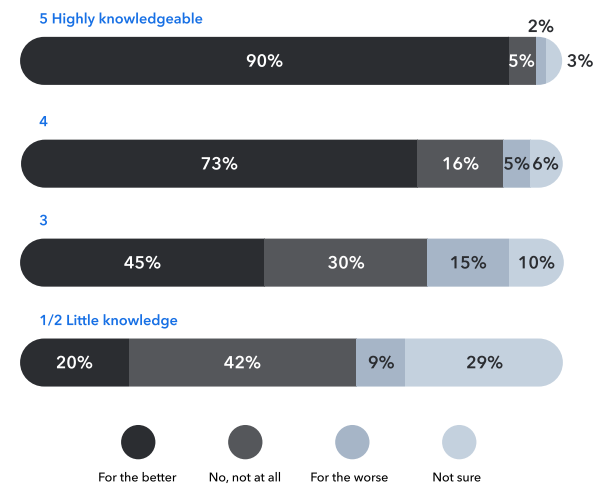
This is even more strongly the case among those who consider themselves genAI experts. Those who rated their knowledge of genAI higher (a 4 or 5) were more likely to say it would positively influence their likelihood of taking a job (73% and 89%, respectively). By contrast, access to genAI tools doesn’t seem to make much of a difference at all to respondents who don’t consider themselves to be among the genAI cognoscenti.

Influence of access to genAI tools on choice to work for an employer



Total n=820

Impact on the likelihood to choose to work for an employer that uses genAI tools

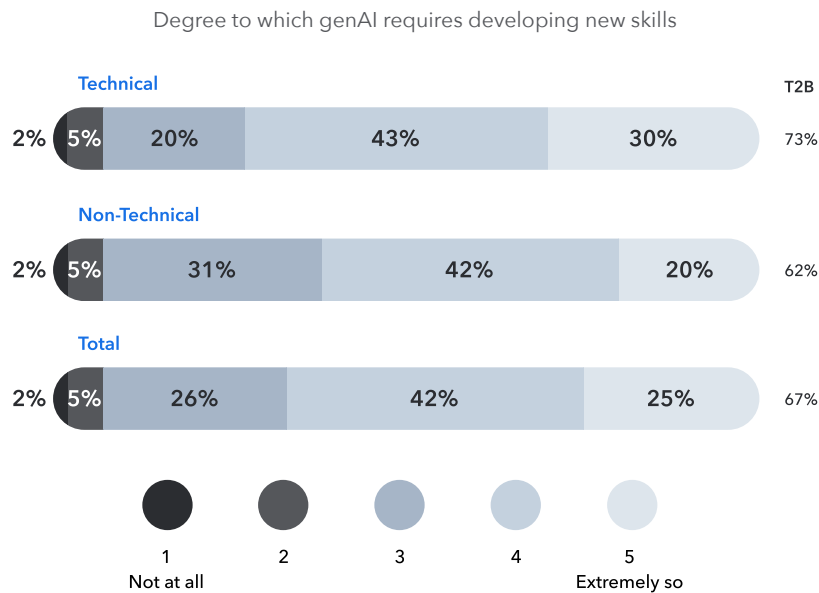


High AI knowledge n=148, 4 n=332, 3 n=244, Little AI Knowledge n=96

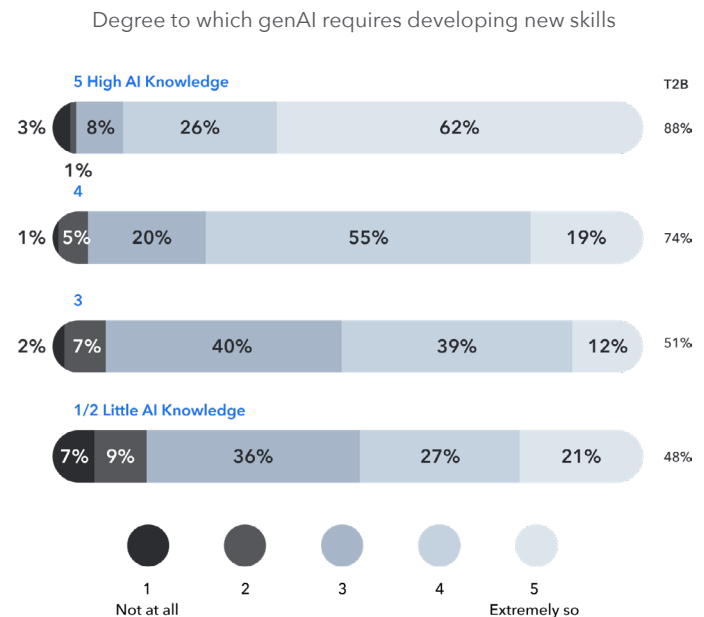
Most anticipate that genAI will require them to develop new skills

Over two-thirds of all respondents, 67%, rated the degree to which they'll have to develop new skills as a 4 or 5. A scant 2% said not at all. Technical respondents are more firmly convinced than non-technical respondents that they will need to learn new skills as a result of genAI.

Here too, the level of genAI knowledge determines the degree to which respondents anticipate the need to develop new skills. Those who consider themselves the most knowledgeable are overwhelmingly likely to think that genAI will require learning a significant amount of new skills. Even respondents who didn't claim to know much about genAI seem to have a good inkling that they'll need to learn new skills, but they perhaps don't yet know to what extent.



Total n=820, Tech n=390, Non-tech n=430





Tools, use, and guidance

On average, those using genAI tools, whether professionally or personally, are using more than one – regardless of role type or level of genAI knowledge

Overall, respondents who use genAI tools professionally use slightly more tools than those who use them personally. As with other areas of our analysis, the average number of tools increases with levels of genAI knowledge.

This “more than one tool” statistic – and among all but the least knowledgeable, more than two – supports the idea that people across all manner of job roles are experimenting and trying out various options to find what genAI capabilities work best for them.

Which genAI tools are you using?

	5 High AI knowledge	4	3	1/2 Low AI knowledge	Total
Professionally	3.73	2.93	2.52	1.56	2.91
Personally	3.28	2.80	2.48	1.59	2.74
Total	3.51	2.86	2.50	1.58	2.82

Total n=1409, High AI Knowledge n=291, 4 n=629, 3 n=413, Little AI Knowledge n=76

By far, the most commonly used genAI tool is ChatGPT, for both professional and personal use


No surprise here – ChatGPT is the most well known and the free version is readily accessible to all. Bing makes a respectable showing as the second most frequently used tool, with over a third of respondents indicating professional or personal use.

More interesting is how many of the other paid, often more specialist, genAI tools have a significant percentage of users, notably GitHub Copilot, Adobe Firefly, Jasper, and Writer. There's a significant subsegment of users for whom these tools are becoming mainstream.

The fact that close to a seventh of respondents indicate using Falcon and Meta's LLaMA, in one capacity or another, hints at the experimentation being done in building custom genAI tools.

Among the fairly small set of "other" tools mentioned, we see on one side genAI capabilities that are integrated into other commonly used tools like Canva and Wix and on the other highly technical investments like proprietary genAI tools. This suggests that we're moving quickly toward a "something for everyone" approach, spanning experts and novices across a range of functional uses.

GenAI tools used



Professionally n=699, Personally n=710

GenAI isn't just for writing blog posts (but yes, that too)


People anticipate using genAI for a wide range of use cases, not just creating content. And their organizations already are. Does that mean that the “tsunami of crap” that many expect genAI to produce might not be so overwhelming after all? Time will tell, but these results give hope that useful output will result regardless.

We asked survey respondents two questions to understand more about what challenges they anticipated genAI would solve for them and where in a professional capacity these tools are currently being used. First, we asked what types of challenges and needs they anticipated generative or other forms of AI to solve for them. Second, we followed up by asking respondents to tell us where they or, to the best of their knowledge, others in their organization were already using genAI.

Among the top challenges or use cases respondents see genAI solving are indeed some that are content-related: researching a content or technical topic, creating a draft, or creating an outline. But there are several important areas not directly related to content, per se, like testing applications, writing code, or cleaning up data.

This validated a hunch we had: for all the focus in the media and elsewhere on genAI “producing content,” it’s actually being used in many different ways, often as part of a broader process, and not necessarily to produce “final product” content.

Use cases for genAI



Total n=820



80%

Technical documentation



80%

Product descriptions



77%

Graphics and charts



72%

Marketing banners



72%

Code



70%

Photo-realistic images



69%

Blog posts



68%

SEO content



68%

PR or other headlines



68%

User interfaces



68%

Video



65%

SEO metadata



65%

Audio

Respondents indicated that they and their colleagues are already using genAI in a wide variety of specific areas, either directly by respondents or, to their knowledge, by others in their organizations.

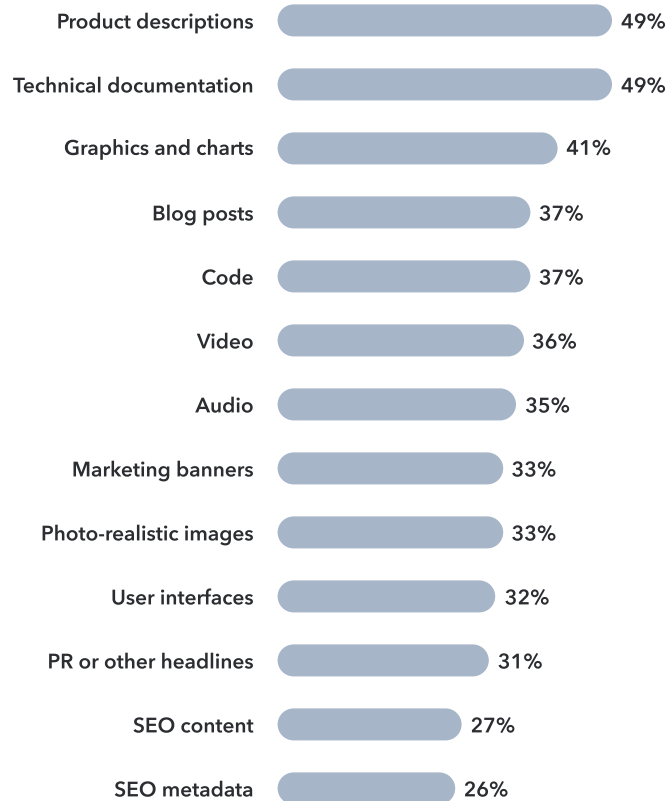
Producing technical documentation and product descriptions rank highest based on respondents themselves or others in their organizations using genAI in the process. Graphics and charts, marketing banners, and code round out the top five current uses.

Current usage areas for genAI

Total n=820

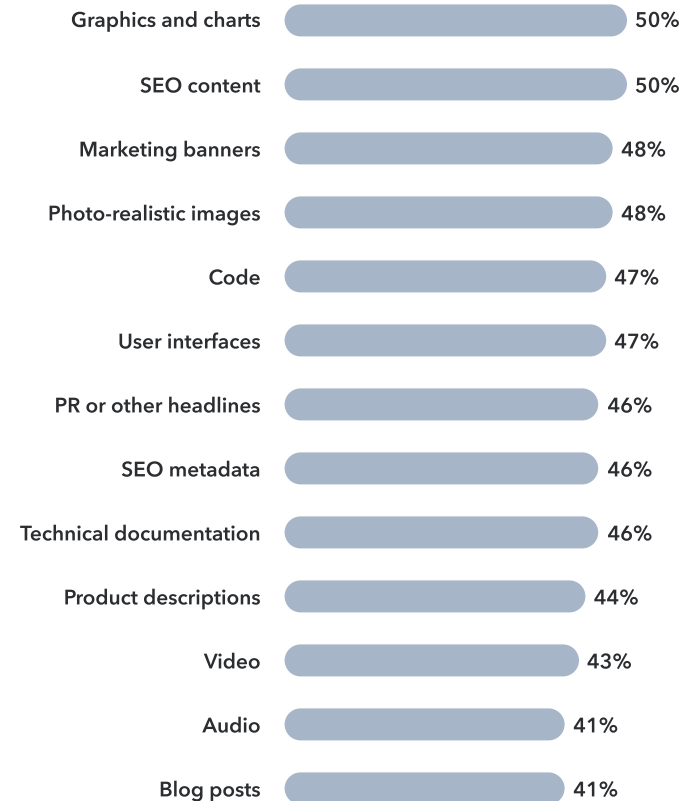
In comparing where respondents say they're using genAI tools and where others in their organization are, we see clear indications that most people think that others are using these tools more widely than they are. Whether hype or that universally felt fear-of-missing-out, most people seem to think their own usage is behind the curve.

Current AI usage areas (Me)



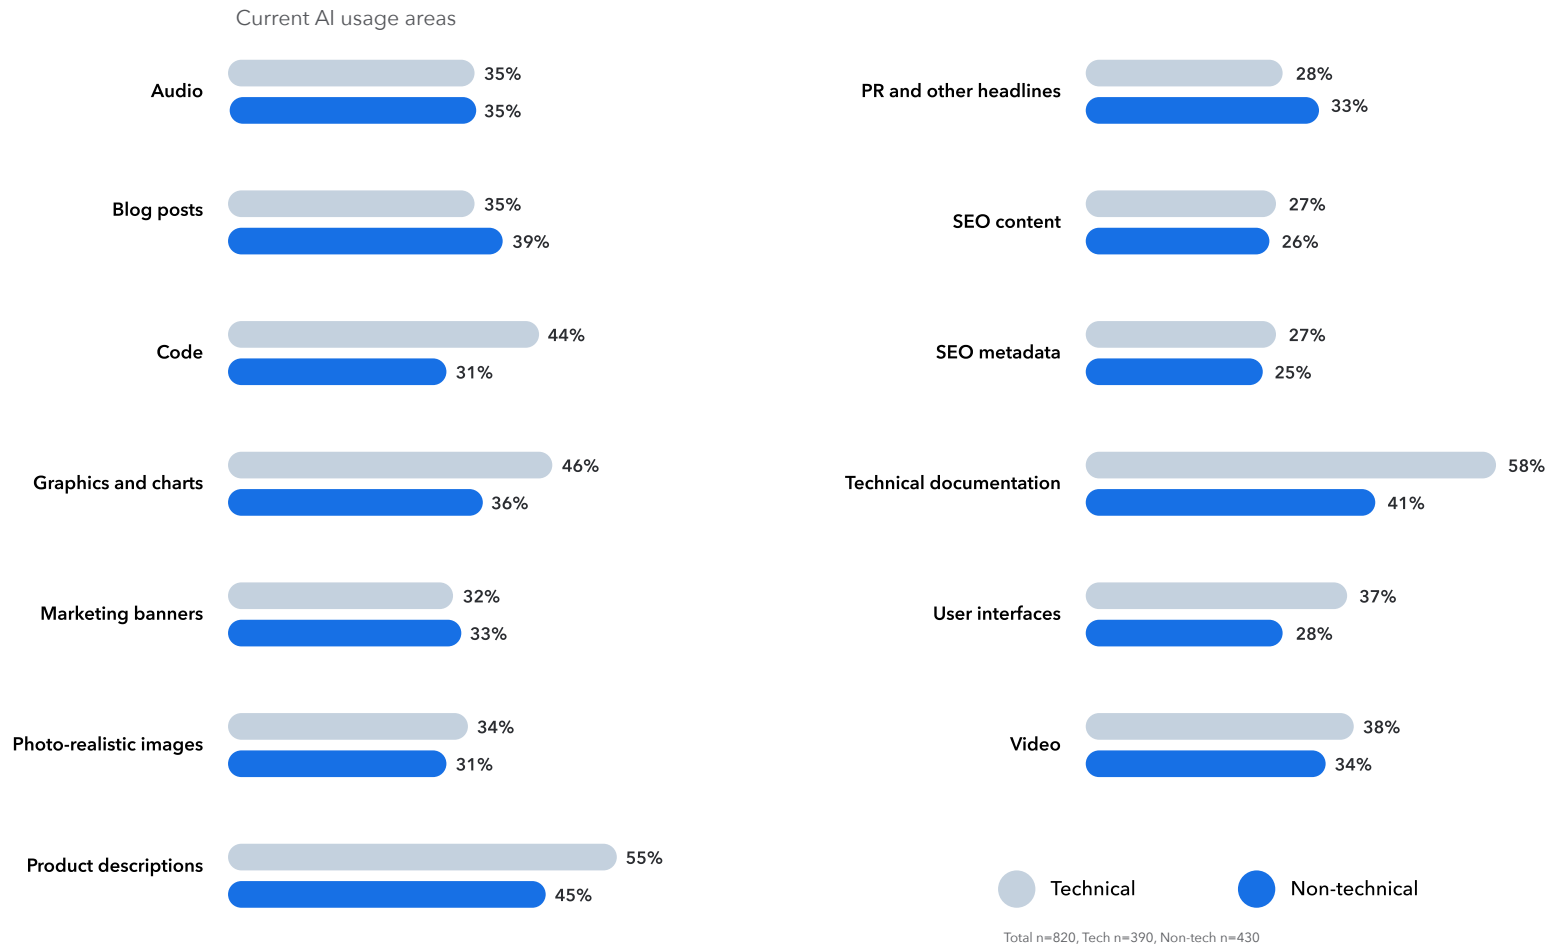
Total n=820

Current AI usage areas (Others in my organization)



Total n=820

While there's a fair amount of overlap, the areas in which technical and non-technical users indicate that they themselves are using genAI reflect different priorities and objectives, but still have significant overlaps across both more technical tasks (i.e., coding) and less technical tasks (i.e., blog posts).



Perspective from the Prof: Sam Maglio's take



Sam Maglio
Professor of Marketing and
Psychology, University of
Toronto Scarborough

Wait, wait ... do tell me

Uncertainty. Let's take a minute to talk about how much people hate it. Consider a lottery that costs \$1 to play. There's a 60% chance that it pays \$2, leaving a 40% chance of winning nothing. That's a pretty good bet! With a 60% chance of \$2, you can expect to walk away with \$1.20 for every \$1 you play. But lots of people don't think like that. There's still a chance of losing the dollar, so they refuse to play.

Throughout the results we're seeing, respondents are telling us how much uncertainty they feel about genAI. Based on how far things have come in the short year since the release of ChatGPT, using these tools seems, similar to the gamble above, like a good bet. But it's also a risky one. That explains why people are dipping their toes in the genAI water and not doing a sprinting cannonball.

Ask people what they think of genAI and it's mostly positive, plus a hefty dose of "I don't know," "I'm not sure," and "no comment."

Those unknowns take some of the shine off the otherwise glistening, positive promise of genAI. Appreciating the benefits it might bring, our respondents want to know how to dial down the uncertainty.

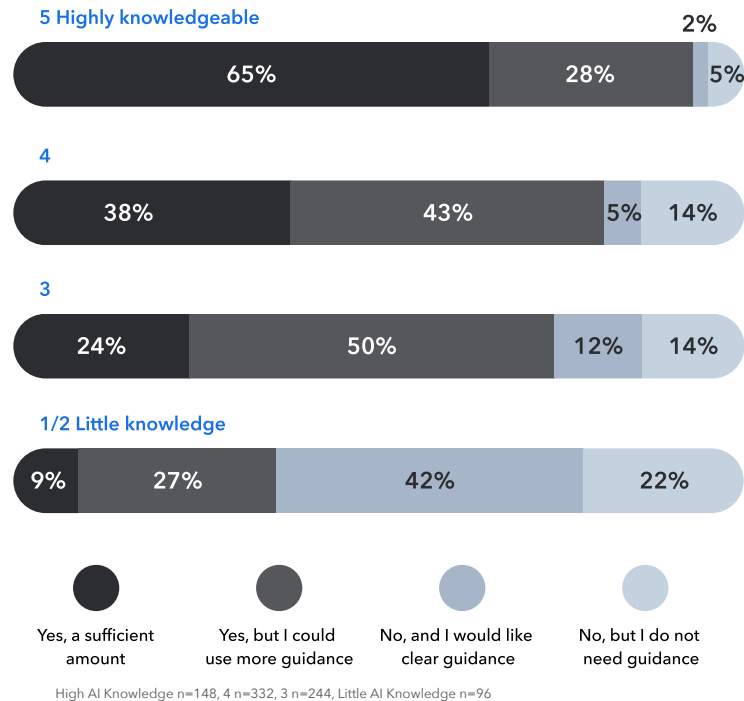
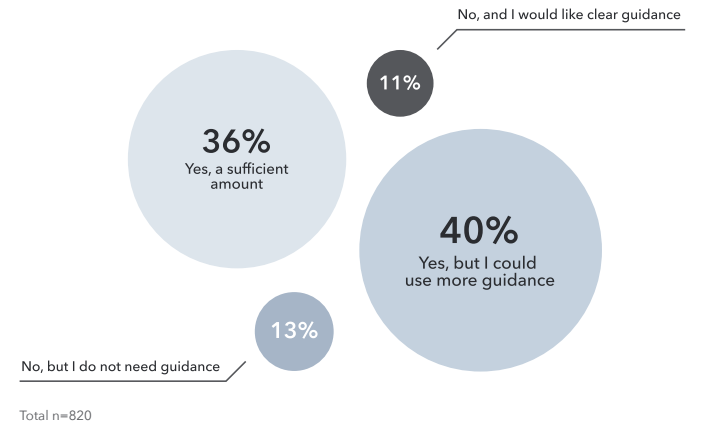
Nowhere is this on clearer display than when you ask people about how they use generative AI and compare it to how they think others use the same tools. For nearly every type of task, and for the technical and the non-technical alike, professionals think that they're using generative AI some but that others are using it more. What are those other people doing? Yes, they're presumed to be using it more to make marketing banners and improve SEO. But what they're really doing, according to our respondents, is "other." That is, people are thinking that there's this awesome tool in genAI and that everyone else is finding ways to use it that they themselves can't fully appreciate.

Despite their enthusiasm, people overwhelmingly want more guidance

Although 36% say they have been given a sufficient amount of guidance from their organization on how to use genAI responsibly, more than half of respondents, 51%, would like more. Forty percent say they already have some guidance but want more and an additional 11% have none but would like some.

Only those who rated their knowledge a 5 are likely to say they have a sufficient amount of guidance. The self-professed experts out there may feel like they know what they're doing with genAI, but just about everyone else would like more reassurance. Importantly, the people who consider themselves the least knowledgeable are most likely, on one hand, to say that they don't have but would like clear guidance and, on the other, to say they don't need it – presumably because they aren't using genAI anyway.

Amount of guidance provided on use of genAI




The vast majority support requirements to disclose the use of genAI

The overwhelming majority of those in our survey, 76%, believe that the use of genAI should be disclosed, whether internally or to customers. Only 12% say that it should not. This is one area in which we expected to find some noteworthy geographic differences. To our surprise, these views are highly consistent across regions, with no significant differences.


Once again, the most meaningful contrasts emerged among levels of knowledge. If anything, those who are more knowledgeable about genAI are more likely to be in favor of disclosure.

We asked respondents to explain in an open-ended response why they felt the way they did about disclosure requirements. Reasonings varied and were often insightful. Consistent themes include, most commonly, the need for transparency and honesty (23% of respondents). The second most common response (11%) expressed uncertainty – people weren’t sure or couldn’t explain why they felt the way they did. (Six percent of respondents provided no answer or an incoherent one, but we’ll chalk that up to survey fatigue.)

Need for genAI disclosure



Total n=820, USA n=203, Canada n=104, Australia n=102, Europe n=360



High AI Knowledge n=148, 4 n=332, 3 n=244, Little AI Knowledge n=96

Although 5% of respondents pointed in some way to fear of what genAI might do, most responses articulated a more nuanced view. They touched on themes around the potential of genAI to improve work or operations (10%), the need to improve awareness or understanding (8%), the evolving capabilities of genAI (6%), the ethical obligations or requirements of disclosure (6%), and the need for regulation or policy guidelines (6%).

The majority of respondents who believed use of genAI should be disclosed clearly indicated that, for one reason or another, it's better to make known where these tools are used. Of those who don't believe there should be disclosure requirements, many indicated a concern that use of genAI may provide some material advantage that businesses shouldn't have to make public. Others felt that since there's already some level of human quality control or involvement, disclosure isn't necessary. A small but distinct group doesn't think disclosure is necessary because they don't think genAI will amount to much or make a meaningful impact.

And then there are responses like this, encapsulating a tech-elitist view that, though far from predominant among our survey population, certainly exists: "I feel like it isn't a hard concept to grasp, if you need to be explained and/or taught about everything regarding AI, you shouldn't be using it." (Considering this view came from a male in the 18-24 age group, we may well classify this as the ignorance – and certainty – of youth!)

Most intriguing to us, respondents who had differing views on the need to disclose use of genAI often had very similar reasonings. For example, the idea that "it depends" came from respondents who variously said they thought use of genAI should be disclosed, shouldn't be disclosed, or weren't sure.

One respondent who felt the use of genAI should be disclosed said: "For our industry, I do not think disclosure of using AI is necessary but in other industries, senior leaders need to be transparent with customers if they use AI in their systems, e.g., health sector." One who did not believe there should be requirements to disclose the use of genAI had a similar view: "If it's just a part of the process and it's not fully controlled by the AI, I don't think it's necessary." Most of the "it depends" camp weren't sure about the need for disclosure, with rationales like this: "It depends on the type of content and how it is used. For example, code generated by AI would not need to be disclosed, but images or other art that is published might."

Across the range of comments, there is a palpable sense of cautious optimism. Most of our survey respondents see the potential and possibilities of genAI and want to encourage positive development while avoiding potential harm or misuse.

Perspective from the Prof: Sam Maglio's take



Sam Maglio
Professor of Marketing and
Psychology, University of
Toronto Scarborough

And action!

No wonder everyone wants more transparency and disclosures around usage. On some level, they're worried that genAI, if not regulated properly, will put others at an advantage over them. As a result, people strive to stay current (paying for genAI out of pocket, playing with more than just ChatGPT, prioritizing employers who use these tools). If there's an advantage to be had, they call dibs.

On the risky chance that others (say, someone competing for the same promotion) might get an unfair leg up, they want guardrails to level the playing field. People may well have an existential fear that, in 20 years, genAI will take their job or wipe out humanity. In the shorter term, we're seeing another kind of AI-anxiety: concerns that genAI will help others get ahead of them.

At the institutional level, this means that businesses need clear internal policies around how genAI can – and cannot – be used in the workplace. You could imagine a near future where companies not only have HR professionals (to govern matters among employees) but also AI professionals (to govern matters between employees and the technology

they use). Even with an internal issue like this settled, businesses also need to consider outward-facing matters where their AI usage meets their customers.

We see from our respondents what people have been saying for a long time: The more we trust genAI, the more we'll embrace it. They've also, for almost as long a time, been telling scientists like me how to make AI more trustworthy: by making it less uncertain. People fear the unknown, but a little bit of help goes a long way to demystify AI. They might see genAI as a black box – but, if you let users put [their own tweak](#) on the algorithm, they [trust it more](#). The same thing happens, according to [my research](#), when people watch an algorithm make a mistake and then learn from it. People might not totally get how genAI works, but people do get people. An easy route to making algorithms less uncertain and scary is to make them more human-like.

People are begging for ways to make the future of generative AI more predictable, regulated, and equitable. As in so many other areas of life, it's best to give them what they want.




Looking forward

Most people prefer genAI capabilities to be integrated into other tools they (already) use, but want the ability to turn it off

More than half, 56%, want genAI to be integrated. A sizable minority, 29%, prefer them to be standalone. Most people, 72%, want the ability to turn integrated genAI capabilities off and on.


Though the majority of respondents (56%) favor integrating AI capabilities into existing tools, those with greater genAI knowledge are more likely to prefer these capabilities be integrated (70% of 5s and 62% of 4s). Those with the least knowledge are fairly evenly split between integrated, standalone, and no preference (see data on following page).

Preference for integrated vs. standalone AI tools



Total n=820

Ability to control use of genAI capabilities

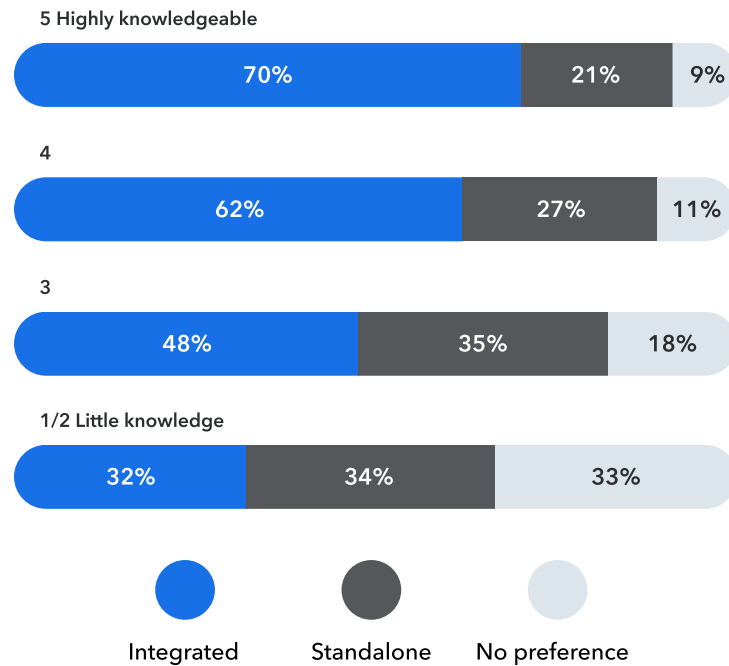


Total n=820

Looking forward

The majority of all respondents (72%), regardless of knowledge level, favor the ability to turn genAI capabilities off and on if they're integrated. The least knowledgeable are more likely to have no preference.

Preference for integrated vs. standalone AI tools



High AI Knowledge n=148, 4 n=332, 3 n=244, Little AI Knowledge n=96

Perspective from the Prof: Sam Maglio's take



Sam Maglio
Professor of Marketing and
Psychology, University of
Toronto Scarborough

Where experts and non-experts agree

It makes sense that experts and non-experts would diverge in their thinking about generative AI. Similarities in how they think – the psychology behind their judgments about generative AI – lead to differences in what they think. So it really jumps out at you when these two groups agree. We found this in two important areas.

First, both self-described experts and non-experts want the ability to flip off the genAI switch. This makes it sound like algorithm aversion is alive and well. People can be hesitant when it comes to taking advice from an algorithm, especially in certain domains. They'll let an algorithm **tell them** which tax prep software to buy but not what kind of clothes to buy with their refund. They'll **trust an algorithm** on how best to drive to a movie theater but not what to see once they get there. Of course, despite all the enthusiasm for this new technology, users still have reasonable reservations about it. The fact that everyone still wants the option to pump the brakes on genAI tells me that people will always want a blend of human and machine.

Second, experts and non-experts want to make sure everyone knows who's doing the work. Three-fourths of both groups insist that users of genAI disclose that they've used it. Some respondents answered the survey from the perspective of being a diligent employee, like the person who said, "Employers should know who is doing the work – employees or AI." But these respondents also make decisions as consumers, like the person who said, "Consumers need to know if recommended action is from a human."


This sense of wearing both hats was put well by the respondent who said, "One should work transparently both within the company and with customers." Disclosure around generative AI is the new frontier in business ethics and corporate social responsibility. Consumers have always prioritized and will always prioritize these values. Companies that meet the moment with openness stand to benefit the most.

More than two-thirds of organizations are considering plans either to apply an existing LLM to their own content or to train a proprietary LLM


Just 31% of our survey respondents said they were unaware of any such plans in their organizations, 18% already have plans and a small, but forward-thinking 6% have projects underway. Of those organizations with projects or plans, 49% are applying an existing LLM and 42% are training their own. We seem to be on the cusp of a major wave of tailored genAI use.

These results validate our view that, for all the philosophical or academic discussions of “artificial general intelligence” emerging at some point in the future, the significant near-term value of genAI for most organizations is in a tailored approach. GenAI tools that reference vetted, validated, and approved inputs are more likely to produce useful outputs, whatever the objective or context.

Plans to train LLM on proprietary content



Plans to use LLM



One of the most pointed questions these data raise is which approach to building a tailored LLM will be the fastest and most effective. Training a proprietary LLM using, say, open-source tools can deliver substantial value, but is not for the faint of heart. This approach requires clearly understood use cases, specialist engineering capabilities, and significant resources. The ability to fine-tune or apply retrieval-augmented generation (RAG) to existing LLMs may present faster and less expensive (or less demanding) alternatives. In all of these cases, the volumes of proprietary content businesses already have become valuable assets. We look forward to examining this topic more thoroughly in future research.



Recommendations

The knowledge gap in genAI has major implications for businesses and individuals alike. Those most knowledgeable about genAI are far more enthusiastic about this new technology than others. They are more efficient because they save more time using it and they anticipate a much larger need to learn new skills.

Exposure and access do not necessarily mean unfettered use. There are many well-founded reasons to be selective about how and where genAI tools are used. However, finding opportunities to use and experiment with them, for all employees regardless of job role, is the most likely avenue to identifying useful applications as well as potentially problematic or harmful outcomes. In this case, knowledge really is power.

Given the rapid rise of genAI and the demonstrable enthusiasm among those who are most knowledgeable about it, businesses have much to gain by making sure all of their employees have access to these tools and the guidance they need to work with them appropriately. People who don't have the opportunity to work with genAI on the job may seek out other opportunities to do so, either outside work or in other roles elsewhere.

Business leaders must take action

Encourage everyone throughout the organization to experiment with genAI – and give them access to tools

Even in businesses that anticipate little benefit from genAI, ignorance could prove exceedingly risky. Knowledge should not be compartmentalized within only a few teams. The potential impact is far-reaching and firsthand experience should be as well. As Prof. Maglio points out, [current research](#) suggests that non-technical users may be the most likely to gain from its use. Even negligible benefits may yield valuable lessons (see our comments on [survey methodology](#)).

Provide a clear set of guidelines and encourage experimentation

Our survey results clearly indicate that employees across a range of roles are eager to see what genAI can do for them, but they're concerned about doing something harmful or unethical. Business leaders must provide the guardrails that allow employees to experiment without fear of inadvertent missteps. The narrow or broad guidelines should be determined by the type of business and even the work of particular teams.

Build a plan for custom genAI tools

Tailored LLMs, whether trained from scratch or leveraging existing models, are poised to be the next step-change in genAI. The ability to leverage existing, proprietary content to drive reliable, relevant outputs has broad potential benefits. So does the use of a range of additional, specialized tools that employ genAI. We expect businesses to see the biggest benefits from these investments. Above all, recognize that genAI tools, for all their potential, are a means to an end.

Individual professionals also have agency

Seek out opportunities to learn, whether at work or elsewhere

A growing number of businesses are actively encouraging and facilitating experimentation with genAI tools among their employees, but this isn't universally the case. For many reasons (well justified or not), some businesses may restrict or prohibit use of genAI tools. Gaining experience with and building understanding of genAI capabilities can come from any sphere, including personal use and outside courses.

Be prepared for change

GenAI has been widely available for just over a year. We are only at the start of using it and understanding what it can do. At the same time, capabilities are evolving even more rapidly than expertise. We can only anticipate how and where it will provide the biggest benefits and make the most impact. Those most knowledgeable about genAI are also the most likely to say there is a lot to learn. Growth mindsets and lifelong learners rejoice.

Appendix: Methodology and demographics

A few words about our survey sample

Any survey will have some bias in its responses, whether because of the subjects it covers, the respondents sought, who responds, or the way questions are posed. Ours is no exception. With a subject like genAI in particular, it's fair to say that people who feel they know more are more likely to respond. That's certainly reflected in our survey population.

It also means that although we met our survey objectives for the number of responses in 10 countries and four regions, an approximate balance between technical and non-technical

roles, a reasonable mix of industry representation and company size, and a mix of age ranges, the gender balance among our respondents skewed somewhat more male than female. The full demographics of the survey are available in the [demographics section](#).

Overall, we feel confident that our survey population presents a good approximation of the diversity among our customer base, though knowledge of or experience with a content management system (CMS) was in no way a qualification to participate in this survey.

Methodology

The survey for this research was developed by Contentful, then validated and fielded by PureSpectrum on behalf of Contentful during December 2023. Respondents were part of voluntary research panels and contacted via email to complete an online survey. We set response quotas by country and had soft targets for a roughly even number of technical and non-technical respondents in each, job levels that covered mid-level responsibilities and seniority, a reasonable distribution across industry sectors, and a range of company sizes.

PureSpectrum and Contentful jointly analyzed the data using Decipher. Additional translation and classification of open-ended responses was conducted by Contentful using ChatGPT 3.5 (also known as the free version).

We found ChatGPT's translations to be consistently high quality, something we validated by requesting a one-for-one output of the original response and the translation. We did a significant amount of spot-checking among these translations and were satisfied with the results.

We also asked ChatGPT, based on the original questions posed, to identify important themes and then categorize the responses within those themes, counting how many fell into each category. Here is where ChatGPT 3.5 produced some less impressive results. The tool was fairly good at identifying themes among the open-ended responses. It was far less good at consistently categorizing individual responses accurately. In some cases, ChatGPT only provided counts

of the responses within each category, without identifying which theme it had classified individual responses within.

When asked to provide each response and the category within which it fell in a table, ChatGPT 3.5 had difficulty with long lists of responses. Any more than about 100 responses at a time caused the system to break down and either stop providing responses or provide nonsense responses. Even with providing 100 responses at a time, the tool couldn't get through the entire list of responses for some questions. (Perhaps we should have paid for access to GPT 4!)

The other major difficulty the tool had was accurately classifying responses, especially when nuance or informal language was involved. Considering how challenging this can be even for experienced humans, we weren't particularly surprised. As a result, we took significant time to manually review or add classifications to all responses, and in some cases, to the list of themes we used to classify them. Despite this considerable manual effort, we estimate that ChatGPT 3.5 saved at least five hours of work over several partial days of effort (perhaps 12 hours of manual work). We probably learned at least as much as ChatGPT did in the process.

Prof. Sam Maglio at the University of Toronto's Rotman School of Business joined the Contentful team in the substantive analysis of the survey data. His unique perspectives as a psychologist examining perceptions and attitudes toward machine learning and artificial intelligence have provided invaluable contributions to the findings we share in this report.

Regional breakdown

Country breakdown

Country	Percentage
USA	25%
UK	12%
Canada	13%
Australia	12%
Germany	13%
France	7%
Denmark & Norway	6%
Netherlands	6%
Mexico	6%

Total n=820

Respondent job level and role

What best describes the level of your position at your organization?

Position	Percentage
Entry-level	0%
Mid-level	4%
Managerial level	59%
Director level	21%
Head of Department	16%
Vice-President level	0%
Executive level (C-level)	0%

Total n=820

Position	Percentage
Technical	48%
Non-technical	52%

What best describes the job function that you work in?

Technical position	Percentage	Non-technical position	Percentage
Web developer	1%	Content creator	5%
Front end designer/developer	1%	Content designer	3%
Back end designer/developer	1%	Content editor	2%
Full stack developer	1%	Content strategist	2%
Software developer	4%	Communications	10%
Data analyst	8%	Digital marketing manager	5%
Cloud architect	1%	Information architect	3%
DevOps manager	4%	Marketing leader	6%
Technical account manager	7%	Public relations	6%
Mobile app developer	1%	Social media marketing	3%
IT	18%	UX designer	1%
		UI designer	1%
		Other	6%

Total n=820

Survey demographics

Industry

Which of the following best reflects the industry you work in?

Industry	Percentage
Financial services and Banking	9%
Public sector	10%
Healthcare	8%
Retail	11%
Manufacturing	9%
Technology	21%
Media and Communications	5%
Professional services	6%
Real estate	2%
Construction, Engineering and Architecture	4%
Transportation and Warehousing	5%
Entertainment	3%
Energy, Mining, Oil and gas	2%
Other	4%

Total n=820

Company size

How big is your company?

Employees	Percentage
I'm self employed	0%
Under 50	7%
50-499	37%
500-999	19%
1,000-1,249	9%
1,250-4,999	12%
5,000-9,999	6%
10,000+	10%

Total n=820

Gender

Gender	Percentage
Female	44%
Male	55%
Other	1%

Age

Age	Percentage
18-24 years	10%
25-34 years	31%
35-44 years	26%
45-54 years	19%
55-64 years	10%
65+ years	4%