

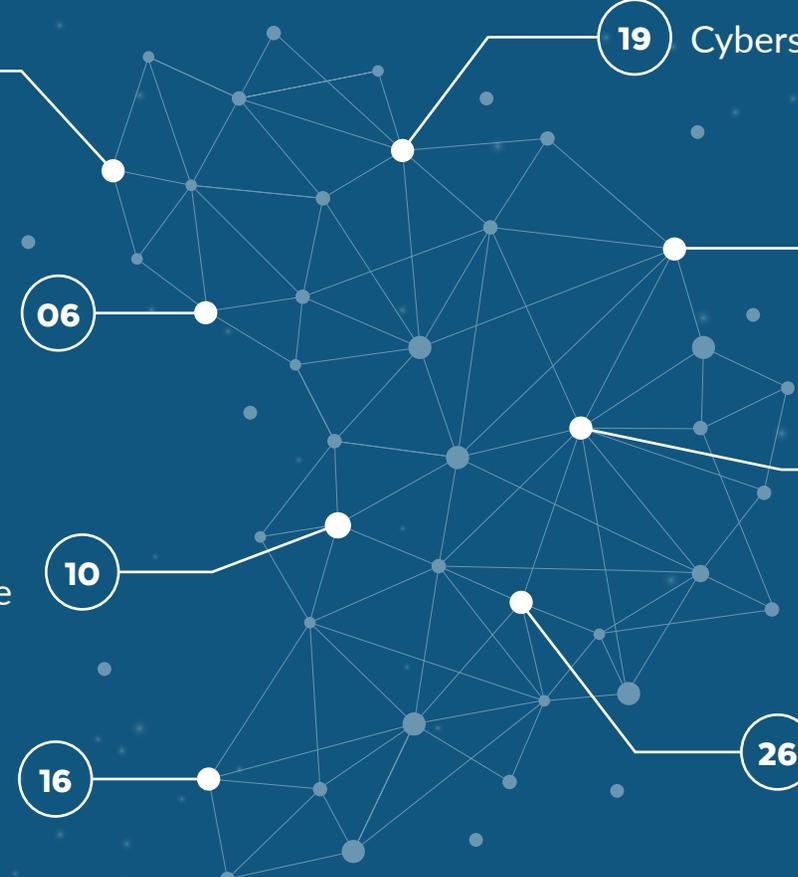


DATA + AI PREDICTIONS 2024

How generative AI is radically reshaping data science, cybersecurity, applications—and every organization's business strategy



TABLE OF CONTENTS



The AI Era Is Upon Us	03	19	Cybersecurity: The AI Challenge
Gen AI and LLMs Will Change Our Lives—Profoundly	06	22	Innovation Abounds
Gen AI and LLMs Will Transform the Enterprise	10	24	Contributors
Technical Roles in an AI World	16	26	About Snowflake

THE AI ERA IS UPON US

Not only was generative AI the big tech story of 2023, but with the constant firehose of articles, announcements and commentaries, gen AI sometimes seemed to be the only tech story of 2023. But don't let the Kardashian level of media dominance put you off—gen AI and the large language models (LLMs) that power it are worthy of the attention. And they are likely to remain the dominant tech trend for the foreseeable future.

In May, the market capitalization of NVIDIA, maker of AI-ready “superchips,” **soared past \$1 trillion**. And San Francisco is **filling up more** of its pandemic-vacated office space thanks to a **boom in AI** startups. We’ve been told that AI will **take our jobs, create** amazing **new jobs, identify diseases** and **discover new medicines, empower** and **detect** deepfakes, **undermine/transform education** and **destroy democracy, fuel** and/or **strangle** creativity, be **our friend** and **our therapist**, and not only **take our orders at the drive-thru**, but flip the burgers, too.

So naturally, when we gathered key leaders and experts at Snowflake to discuss the year ahead (and then some), generative AI dominated the conversation.

“Generative AI and LLMs are definitely the hottest topic right now—they’re sucking all the oxygen from the room,” says Mona Attariyan, Snowflake’s Director of Machine Learning. “It almost feels like there is nothing else in machine learning that anyone wants to talk about.”

And for good reason, says Christian Kleinerman, SVP of Product at Snowflake. “Lots of true disruption is coming. Mostly around end user experience and how people interact with technology,” he says. Based on what we’re seeing, it’s not just hype. “The dramatic change everyone’s talking about is real. Generative AI and related technologies will affect productivity, job roles and responsibilities. It will aid creative processes and create entirely different experiences.”





Amanda Kelly, Co-Founder of open source app framework Streamlit, which Snowflake acquired in 2022, notes that even as data has become more available for businesses, it has been locked in databases, with data scientists and other technical experts placed in the role of gatekeeper.

“But now generative AI is unlocking access to that insight directly for decision-makers,” she says. “We’re only starting to see how AI will assist us with cognitive work, the way robotic machinery that can lift tons helps with physical work.”

No surprise, then, that our 2024 data and technology predictions are weighted very heavily toward the impact of an AI revolution that’s centered around the power of LLMs and the transformative potential of natural language interfaces. Suddenly, our computers can figure out a lot, and we can essentially talk to the data like it’s a person. As we’ll see, this has implications for tech jobs, data strategy, cybersecurity and a lot more.

In the short term, the general impact of generative AI backed by LLMs will be felt as accelerated productivity and enhanced insight. These powerful tools won’t be doing the work so much as helping workers. The current crop of productivity tools for coding or text generation needs considerable human oversight. But that may change, because these productivity-enhancing tools are developing quickly. At some point, they may have a greater effect on staffing needs.

For now, think about enhancement, not replacement. Prasanna Krishnan, Senior Director of Product Management, Collaboration and Snowflake Marketplace, points out that the easy diagram of data → insight → action is really hard. “That looks like three steps, but the process can

really be 10 to get from data to insight, and 20 to get from insight to action,” she says. “Generative AI will help reduce the steps for the end user, for greater efficiency and velocity. Maybe the 30 steps between data and action become five. Or two.”

Snowflake SVP Sridhar Ramaswamy, who pioneered Google AdWords and more recently created Neeva, an AI-based search engine acquired by Snowflake in 2023, agrees that a lot of efficiency and convenience is on the horizon. Ramaswamy says that being able to talk to an application as though it’s a human being will bring innumerable little improvements.

“The other day, I had to fill out an online form that required me to enter a time, but it couldn’t understand ‘4pm’ or ‘1600’—I had to figure out it wanted 0400:00,” he says. “Eliminating that sort of friction to smooth our interaction with the myriad of online entities we deal with every day will be hugely positive. Getting things like better suggested sentences or an email template, or help understanding a complex piece of text, that’s all positive.”

This report will cover the near-term effects of generative AI and large language models, and will look a little further ahead, as well. It will touch on new cybersecurity threats and strategies, the evolution of the open source ecosystem and more.

But first, we explore the big prediction: Gen AI is changing everything, for better and worse.

GEN AI AND LLMS WILL CHANGE OUR LIVES— PROFOUNDLY



AI's impact on society will be huge and fast.

First, the big picture. After a solid year of generative-AI hype, we're seeing the inevitable backlash. While corporate leaders have voiced real concerns about cost and technical challenges slowing the rollout of generative AI and LLMs, this tech is not a bubble that's going to pop.

"It's comparable to the arrival of the smartphone," says Mona Attariyan, Snowflake's Director of Machine Learning. "Since the iPhone, the amount of time we spend accessing data and applications has gone through the roof, really changing how we move through our lives. The arrival of generative AI will be a similar step change, only much faster."

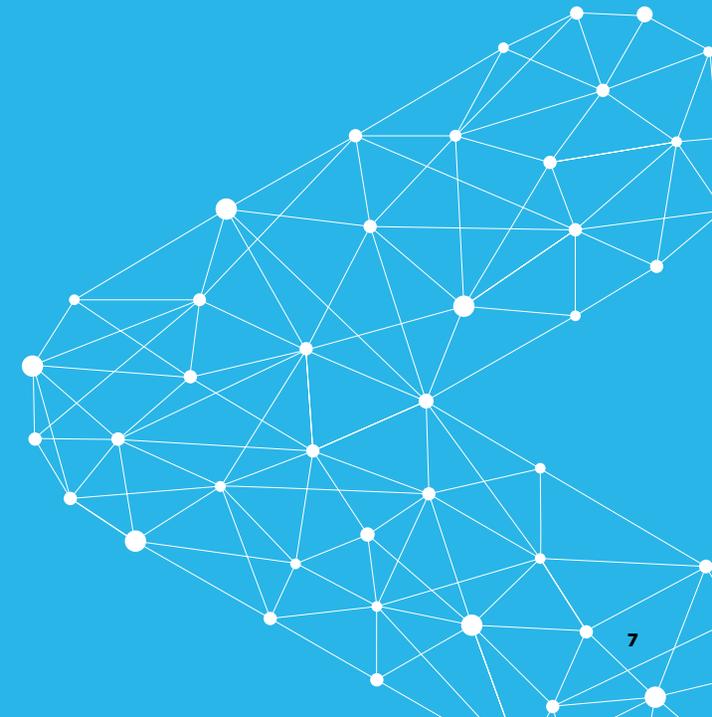
"The idea of AI disruption is very real," says Christian Kleinerman, Snowflake's SVP of Product. "There's a lot of opportunity to improve things in the business world, whether it's around making individuals more productive, or creating innovative end user experiences and interactions. It will change roles, responsibilities and skill sets."

"The tech industry has spent decades bringing data and digital technology to how we work, but for the most part, the core fundamentals of our day-to-day work haven't changed much," says Amanda Kelly, Co-Founder of Streamlit. "The industry is only now reaching the point where the tech industry is not just driving efficiency but empowering business people to truly change how they work."

In the short term, that means a true "democratization of data," as natural language interfaces allow business decision-makers to dive deep into data that previously required the help of necessary gatekeepers, such as data scientists, business analysts and other highly technical experts.

GENERATIVE AI WON'T BE A DIFFERENTIATOR FOR LONG.

“Increasingly, generative AI functionality will become table stakes,” says Neeva Co-Founder and Snowflake SVP Sridhar Ramaswamy. “There’s this massive rush that a lot of incumbents are embracing to offer smooth, fluid, frictionless access to functionality they currently have, and pretty soon, everyone will have to have it.”





Generative AI's negative effects, including job loss, deepfakes, and a deepening digital divide, will be hard to manage at first.

There's a tendency in the tech industry to gloss over the potential negative impacts of a new technology. While accentuating the positive is a virtue, a realistic appraisal cannot eliminate the negative. Snowflake's experts surfaced three immediate concerns that will be particularly challenging in the early years of widespread availability of generative AI and large language models.

First, there's the effect on jobs. Upbeat pundits emphasize that AI will assist humans, not take them out of the equation; however, when AI helps you do the work of two people, things might not look great for that other person.

"A lot of people involved in what we loosely call 'knowledge work,' quite a few of their jobs could vaporize," says Sridhar Ramaswamy. (CBS News reported that **4,000 jobs were felled** by AI in May 2023 alone.) He notes that the arrival of past technologies, including the PC, also eliminated a number of jobs and whole economic sectors that needed considerable time to rebound. Rapid change forced by widespread AI adoption, he says, would make it hard to quickly absorb displaced workers elsewhere in the workforce. "Both the private sector and governments will need to step up."

"It will be painful to adapt," says Mona Attariyan. "But I am very optimistic that we'll work it out as a society. There's quite a bit of awareness around potential downsides, and I see both business and government leaders starting to act and address these issues early on."

Deepfakes represent another pain point. While responsible creators of AI tools talk about adding digital watermarks to make it easy to identify fake content, we can expect bad actors to find workarounds—or toolmakers that don't feel so ethically constrained.

"In the next few years we can expect to see an assault on what we humans collectively think of as our reality," Ramaswamy says. "A world where no one can or should trust a video of you because maybe it was AI-generated. That's a very different reality from the one we're living in. That's a big issue."

"Disinformation will be painful," Attariyan agrees. "There are places where regulations are needed, and where companies and regulators need to act."

Finally, looking at the long term, Ramaswamy worries about a worsening digital divide.

"One of my big fears with generative AI is that these advances will exacerbate the divide between the haves and have-nots that has been happening over the past 20-30 years," he says. "I worry that advanced AI could exacerbate inequality across the globe. At the same time, I hope that by making information so much more accessible, this technology produces a new generation of young adults who better understand the issues and potential, and can counter that risk."

Generative AI will accelerate incremental innovation, but humans will still make the big swings.

Snowflake was founded on innovation. A couple of database experts at Oracle had an idea to break away and build a different kind of data platform. One of those experts, Benoit Dageville, believes algorithms will be unable to replace that kind of innovative thinking.

“What we’ll see with generative AI is more incremental, iterative innovation,” says Dageville, who is Snowflake’s President of Product. “AI can assist in research, prototyping, etc., so a company can pilot a new product quickly.”

Deeper innovation, he says, requires a leap that’s contrary to the data. “For instance, we created Snowflake at a time when all the data and opinions said that the future of big data was Hadoop,” he says. “We didn’t like Hadoop, so we went and built something contrary to the data and the prevailing beliefs.”

“Humans will still drive innovation and decision-making,” says Christian Kleinerman, SVP of Product. “The use cases are all human-plus-machine. Humans will continue to take responsibility, provide empathy; all the things it’s hard to program a machine to emulate. But for the heavy lift of computation or mechanical work, AI models will be—and already are—amazing.”

Dageville says he’s excited about what generative AI brings to the table, but absolutely does not see a threat to human creativity. Identify a new customer segment? AI will probably win at crunching the data. But give birth to a brand-new concept? AI’s just a tool in the creator’s hands.

“If you want to create a revolution, that revolution, by definition, does not exist,” he says. “So there’s no data on it.”

Ethical guardrails for AI will emerge, from both private and public sectors, faster than with other tech upheavals such as privacy.

To say that there are concerns about the potential of the latest AI advancements would dramatically understate the case. In March 2023, more than 350 tech experts, executives and others—led by Geoffrey Hinton, widely recognized as the father of AI—signed a one-sentence [open letter](#) stating, in its entirety, “Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.”

That’s one heck of a warning label.

Essential to responsible AI development and use will be action by industry players and governments to establish safe and ethical rules for the use of generative AI and other advancements. You might look at the less-than-satisfying treatment of privacy issues in the tech industry and predict the worst, but the failures and false starts around privacy, by both the public and private sectors, may lead to better AI guidelines, in a more timely manner.

“This AI boom came at a good time,” says Mona Attariyan. “We’re already bruised by the privacy issues, which are recent enough that we’re taking this next thing seriously. Eight years ago, society might’ve been a bit more naive if facing this AI moment. The players I see creating these technologies seem to be genuinely interested in helping to solve these issues.”

Sridhar Ramaswamy agrees with Attariyan that governments are also stepping up earlier in the cycle. In mid-September, the U.S. Senate hosted a private informational roundtable that included leaders from OpenAI, NVIDIA, Google, Meta and more. But, he cautions, quick regulatory intervention might not solve the problem(s).

“Narrow tech regulation is very hard,” Ramaswamy says. He points to Section 230, part of the 1996 U.S. Communications Decency Act, which provides immunity to websites for what third parties post. While it made the internet as we know it possible, the internet as we know it is rife with lies, hate speech and bullying. “We’ve seen that well-meaning regulation like Section 230 can play out in less than ideal ways.”

Anoosh Saboori, Snowflake’s Head of Product Security, thinks that industry leaders—particularly chief information security officers—will be highly motivated to step up and tackle many of the potential ethical issues and liability risks of generative AI. “Every time a new tech like a gen AI emerges, there’s a battle because CISOs own the risk. They’re going to demand a framework to evaluate these providers of generative AI. We’ll have to come up with ways to measure security, compliance and ethical use.”

In other words, the market will be the forcing factor and the government will follow.

“The signals I see from the government are that they’re already asking the industry to define what responsible AI means,” Saboori says. “I think they’re paying attention to the issue.”

GEN AI AND LLMS WILL TRANSFORM THE ENTERPRISE



Generative AI will supercharge the data strategy of tomorrow's leading businesses.

For years, companies have been urged, or admonished, to develop a comprehensive and forward-looking data strategy. Just as more and more businesses were ticking that box, AI advances threaten to render last year's plan moot.

Fortunately, our experts were unanimous that if you've already put in the work to create a solid data strategy, you're on the right track.

"The generative AI era does not call for a fundamental shift in data strategy," says Jennifer Belissent, Principal Data Strategist at Snowflake and a former Forrester analyst. "It calls for an acceleration of the trend toward breaking down silos and opening access to data sources wherever they might be in the organization."

Mona Attariyan, Snowflake's Director of Machine Learning, says that the change is to the level of importance of your data strategy—and the speed with which you need to execute and invest in it. "If you don't accelerate, you risk being left behind, virtually overnight."

That doesn't mean pursuing all of the seemingly limitless promised miracles of generative AI at once. Assistance with basic coding and copywriting is nice, but that's not what differentiates a business. Leaders should always prioritize the unique insights that can be drawn from their own data.

"There will be applications of generative AI to just about every aspect of your business," says Snowflake Co-Founder Benoit Dageville. "But continue to focus on your core needs, the things that keep you up at night, and invest in those."

Which means that how you manage and govern your data somehow becomes even more important. Snowflake CIO Sunny Bedi stresses that there is no room for shortcuts or half-hearted approaches.

"Governance is table stakes," he says. "Security, governance and compliance are the minimum price of entry to the world of generative AI and LLMs."

If there's one topic about which the new AI era challenges the IT status quo, both Bedi and Anoosh Saboori say it's where an enterprise puts its data.

"Generative AI will increase the trend toward centralized data on a managed service platform that provides the security and governance you need, and creates a single source of truth for LLMs and everything else," Bedi says. And he says he can recommend a great platform for that, if anybody's looking ...

Saboori, Snowflake's Head of Product Security, says that the many cloud-adoption laggards will have to step up quickly. "Gen AI will make it hard to keep your data on prem," he says. "That's still a significant discussion today, because the majority of workloads are still not in the cloud."

This push to the cloud, he notes, will complicate the security and compliance strategy: "Customers adopting a multi-cloud strategy often end up distributing data and models across many platforms with different security and governance capabilities. It becomes very challenging to manage your security and governance posture across such distributed environments."

AI WILL BE YOUR WORK BUDDY.

One of the most exciting ideas for workplace productivity, says Sunny Bedi, is the deployment of AI assistants that help new hires become—and continue to be—efficient and effective. He says that onboarding new workers is a complex process of educating the worker about systems, processes and culture, and providing prompt access to (only) the right systems and projects.

He says to expect AI assistants tuned to your specific department and role that will provide that orientation and accompany you throughout your tenure with the company.

"It will be tied to your persona," he says. "There's tremendous productivity we can provide through this. As your processes and needs become more mature, you're training the agent to do the same thing on your behalf next time. And when we hire a new person into your org, that person gets that full wealth of knowledge from the beginning."

The idea is already coming to fruition. In late August, Axios **reported that Walmart will launch** a generative AI app using an unnamed third-party LLM trained on corporate information. "My Assistant" will help its 50,000 non-store employees summarize long documents, create content and more. Expect many companies to follow suit.

LLMs/Gen AI will supercharge opportunities for data monetization.

Because enterprises have a lot of sensitive data to manage, expect most of their LLMs to be in-house, rather than turning to public tools like ChatGPT. Picture a shipping company buying a foundational model from the likes of OpenAI, then training it up with its own proprietary data. But they'll want more than that.

"Not only are companies going to feed their own data to their internal LLMs, but they're going to want to buy external data sets particular to their industry/market," says Jennifer Belissent, Snowflake's Principal Data Strategist.

She says business leaders will need a bigger picture than what their own data provides.

"I had a CDO tell me that with their own data, they can only see internally," she says. "They need regional trends, industry benchmarks, etc. They need to understand the context in which they are doing business to discover the opportunities that are out there."

This need means businesses can be more than just greater consumers of data, Belissent adds. She points to payroll software maker ADP, whose ADP Insights data product looks at pay growth across the United States by region, industry, gender and more.

"The rising demand for data brings the opportunity to sell your data to a very hungry market," she says.

LLMs will make usage-based consumption hot (okay, hotter).

Remember when installing software started with opening a box of two, six, a dozen CD-ROM discs? No one misses those days. Software pricing evolved from "buy software package once; buy again in five years for an upgraded version" to "software as a service subscription means pay per user (whether they use it or not) for access to up-to-date software." Already, the SaaS model is transitioning to consumption-based pricing: "pay for the data that you access or the compute that you consume," as we see with cloud providers.

Prasana Krishnan, Senior Director of Product Management, Collaboration and Snowflake Marketplace, says that the new AI era will significantly increase the adoption of consumption-based pricing models.

It's easy to charge for a software subscription, she says, even if that user-based model can be inefficient. "But when we talk about services or applications that leverage a model trained with a variety of data sources, those data owners should get paid as their data is used."

Thus the attractiveness of paying for data and applications on the basis of actual usage. Adopting these consumption models will be a challenge for organizations not used to pay-as-you-go.

"With true usage-based models, you need visibility into costs and to establish budgets," she says. Basically, when IT turns usage-based services over to its users, it's like parents giving their teenager a Mastercard. You definitely want to set a credit limit.

"We help customers do that on Snowflake—they'll be notified when they're close to hitting a budget threshold. It's all about understanding your spending, having controls to set limits and easily taking actions to optimize."

In addition to efficiency, she says that consumption-based pricing models will supercharge innovation. "We'll see more rapid prototyping with AI, because in that experimentation phase, usage-based pricing is really powerful," she says. "Otherwise you're stuck with negotiating a \$100,000 contract just to prototype an unproven idea."

That's not to say the subscription model is doomed. In some cases, usage-based pricing may be an entry-level tier; get a prototype into production where data or compute is being consistently used and you might decide to uplevel to a flat fee.

"Once a developer's idea moves from experiment to production and there's consistent, high usage," Krishnan says, "it could graduate to a subscription with a set, all-you-can-eat price."

It's not just gen AI—it's the apps that will revolutionize how we live and work.

The excitement around generative AI is often expressed as “being able to talk/interact directly to the data,” which is not exactly correct. The people who interact directly with the data are in a back room scripting out Python queries. The rest of us are going to interact with applications—like ChatGPT—that intermediate between us and the large language model. And of course the data itself is not sitting on whatever device we're using to access the app. Which means that the era of advanced AI is about more than the data and the algorithms.

“The revolution is not AI, or generative AI,” says Benoit Dageville. “It's AI plus cloud plus applications. What made the smartphone revolutionary wasn't the voice or text functionality—it was the apps.”

Smartphones begat app stores, filled with simple applications that perform discrete functions. That's how we're going to be interacting with AI.

“An explosion of applications has already been happening,” notes Christian Kleinerman. “But it will be significantly accelerated by generative AI and LLMs. There will be a major category of apps around those technologies, and many more applications will have AI-based search and natural language processing built in.”

“Everything is appifying,” Amanda Kelly agrees. “And I've increasingly seen very bespoke applications designed to handle a specific ETL job or pipeline or data type. LLMs will uplevel that, making it easier to create and use very useful, use case-specific applications.”

Thus, Dageville says, companies will need to look not only at AI adoption, but at how their users will interact with it. “As businesses try to figure out what to do with AI, they have to think about the app they need to build as the interface between that data and your users, whether they're employees or customers.”

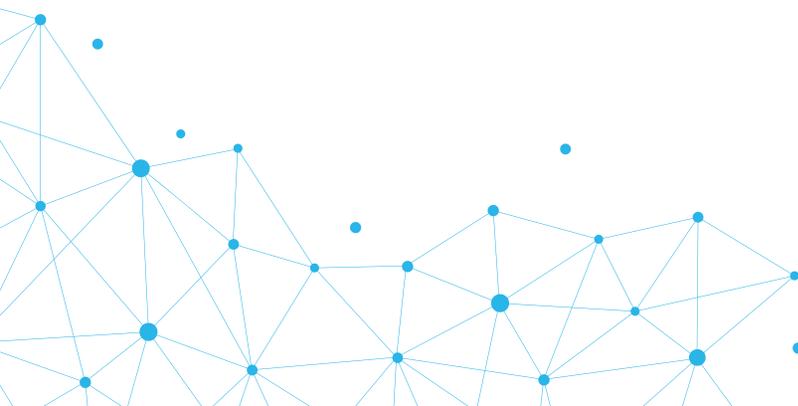
AI WILL NOT DESTROY THE INTERNET OR AD-BASED REVENUE MODELS.

One of the initial fears voiced in the wake of ChatGPT was that if people start asking AI to summarize internet content instead of looking it up themselves, the traffic that drives the ad-based internet will wither, and therefore, so will the internet as we know it.

Not so fast, says Snowflake SVP Sridhar Ramaswamy, who was instrumental in building and leading Google's AdWords program, and who tried to take on the ad-based internet by creating Neeva as an AI-driven, privacy-conscious, ad-free search engine. He says the net has held up just fine.

“The internet is a pretty robust and adaptive entity,” he says. “The ad model is incredibly powerful, and there's tremendous money riding on it. That model will continue to thrive in this AI era.”

Not only that, but expect ads to get more personal. “Your favorite celebrity will be talking to you directly in an ad,” he says. “Who knows, maybe a virtual Michael Jordan will be calling to say, ‘Hey Sridhar, buy the new Air Jordans.’”



The creation of AI models will be standardized, outsourced and specialized.

It will take time for wide adoption of LLMs in the corporate world. The “large” in “large language models” makes them expensive to train and run. Research is already well underway on how to make not-so-LLMs that run on fewer parameters, trained for specific functions. ChatGPT can summarize long pieces of text, suggest recipes based on what’s in your fridge, or plan your dream vacation to Buenos Aires. That’s a lot more range than you’ll need for an app that helps your CFO manage costs or that helps your sales team target new customer segments.

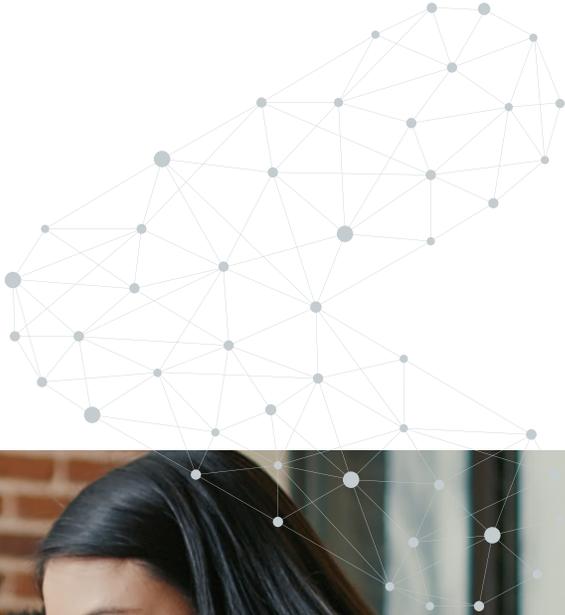
“There will be the few giants that will build the general ‘all things to all users’ model,” says James Malone, Snowflake’s Director of Product Management for data lakes and data storage. “But many more vendors will hyperfocus and hypertune on specific industries, use cases and data types.”

“There will be a robust and competitive market of foundational models, under a larger economy of custom model makers for various industries—healthcare, retail, finance, and manufacturing, for example,” says SVP of Product Christian Kleinerman. “Then enterprises will package those models together and add their own data.”

“Software vendors tend to acquire smaller companies to bring their specialties in-house,” Malone says. “But that won’t be the trend with AI. Smaller companies will exist to fine-tune models by specific need. There will be an AI supply chain that includes large general models and the smaller, industry-specialized models.”

So there will be hundreds of companies all looking to tweak and repurpose two or three big, expensive LLMs?

“The market for those foundational models will be robust, without significant consolidation, for at least five years,” Kleinerman says, adding that AI develops so quickly, and there are so many unknowns about adoption, that it’s hard to predict further out. “But at some point, that market presumably consolidates.”



TECHNICAL ROLES IN AN AI WORLD



Data scientists, data engineers and BI analysts are in for a fun ride.

LLMs and generative AI are going to have a big impact on the most technical of data users, and it's largely for the better. The traditional complaint in data science is that a lot of the actual work is basic data prep, boring stuff that's like having to mow the entire field and paint the white lines before you can play a little soccer. The promise of this new AI era is that a lot of the scut work will be automated by smarter AI tools.

1 Data engineering will evolve—and be highly valued—in an AI world.

You could forgive data engineers who worry that AI's going to eat their jobs. They needn't worry, says Mona Attariyan, Director of Machine Learning.

"Good data engineers will be highly valued," she says. "The primary job of a data engineer is to make sure the data comes in and is organized properly. You cannot leverage your data or the power of LLMs without the essential work of a data engineer."

The data pipelines built and managed by data engineers will be perhaps the first place to connect with LLMs to unlock value. Data engineers will be the ones who understand how to consume a model and plug it into a data pipeline to automate the extraction of value.

"Their expertise will still be needed, including in new and different ways," says James Malone, Director of Product Management. "Data engineers will need to solve unique use case challenges. They'll also be the experts who can oversee and understand the work that was formerly done by data engineers but will be routinely done by generative AI."

2 Data scientists will have more fun.

Some of the problems that data scientists tackle will go away. For example, building yet another model that does sentiment analysis may be irrelevant. Many of the problems will change and new ones have already emerged. For example, how to feed the right information into an LLM as part of a context, how to minimize hallucination, or how to prompt an LLM well are entirely new challenges that data scientists and engineers now face.

Streamlit Co-Founder Amanda Kelly says that generative AI will make data science jobs more attractive. "I've seen commentary from students, when you ask whether they'd like to become data scientists, that's critical of all the prep work, how you spend all your time doing tasks that aren't interesting—much of which will end up being handled by AI. That's the history of technology: It helps to automate away the things that are not the interesting work."

These fun-loving data scientists will need to adjust their skill sets a bit. Just as cloud infrastructure forced IT orgs to learn a new skill set by moving from builders of infrastructure and software to managers of third-party infrastructure and software vendors, data science leaders will have to learn to select and integrate the right external vendors of AI models.

"You can think of a data scientist or business analyst as an intermediary between raw data and the consumer," says Adrien Treuille, who co-founded Streamlit with Kelly. "The work has to be accurate, and if their report or dashboard comes with hundreds of caveats, the data scientist needs to do more work."

When this becomes the work of the LLM, Treuille says, the data scientist will go from creating the dashboard to helping the LLM answer questions correctly. "The data scientist will have to tell the model, 'Here's a bunch of context about our data and how it works, and therefore when interacting with the CEO, know all these things,'" he says. "Data science becomes the art of formalizing and embodying the company's metadata in such a way that the LLM can answer accurately on behalf of a broad range of users."

3 BI analysts will have to uplevel.

Today's analysts generally create and present a canned report, following up with new queries when an executive has a specific supplemental question. In the future, execs will expect to interact directly with data summarized in that overview report. This self-service will free analysts to work on deeper questions, bringing their own expertise to what the org really should be analyzing.

That upleveling, says Snowflake CIO Sunny Bedi, is inevitable, and that makes it incumbent on analysts to up their game. "If I offer you either eight dashboards or an app that lets you query the data in natural language, which would you prefer?" Bedi asks. "The BI person writing those dashboards will want to uplevel their role to solving problems the AI interface can't."

4 Developers expect to be 30% more efficient using generative AI assistants.

CIO Sunny Bedi asked his dev team to estimate how much of the code that they're writing could've been handled by a gen AI tool, and they consistently estimated 30%. While this has yet to be proven, that level of efficiency would be a real game changer, he says.

"Beyond that initial efficiency," he adds, "there's reusability and sharing. A piece of AI-generated code that makes the dev 30% more efficient today can be used again on other projects, increasing the overall efficiency."

Plus, he observes that testing and quality assurance could eventually be assisted by an AI agent, leading to faster, higher-quality deployments. But the coders will still need to code. For now.

"You can't debug code or edit writing unless you have some skills as a coder/writer yourself," Attariyan says. "In the near term, these tools are about more quickly doing things that you could definitely just do yourself."

Like others, she hedges when asked to look out more than five years. "It's hard to predict," she admits. "Things are moving so fast, systems improving so much, it's actually not clear how much supervision various AI outputs will need in the long run."

WHAT TO WORRY ABOUT WHEN YOU'RE DONE WORRYING ABOUT AI ADOPTION

The big generative AI questions today are about how to embrace this powerful technology. We asked our experts what else executives and consumers should worry about.

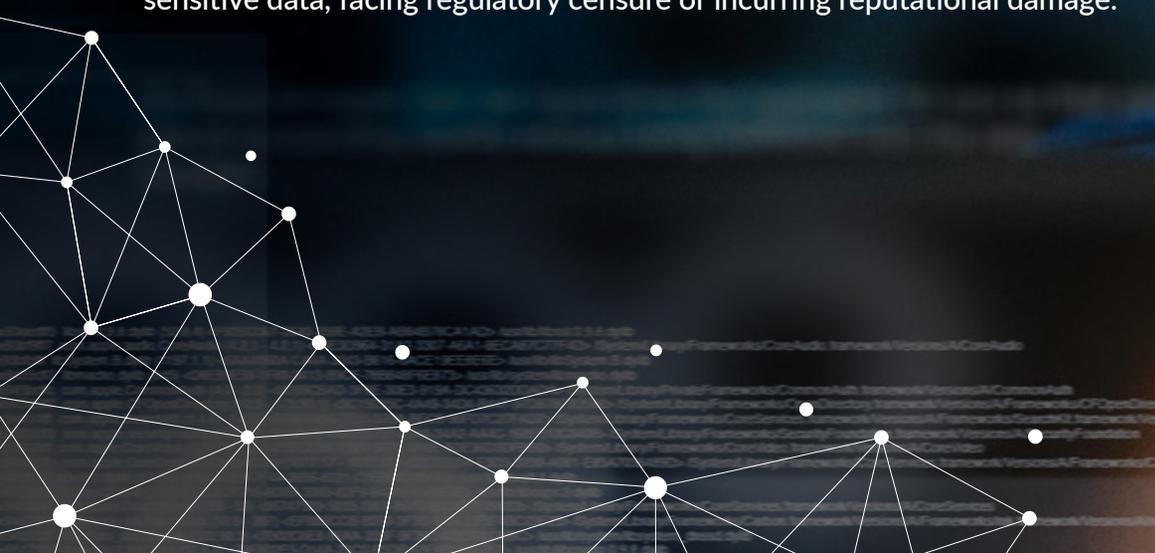
- **CEOs need to think about regulation.** "Today, business leaders are thinking, 'How do I leverage gen AI? Am I moving fast enough?' But they need to balance fast with fastidious. As regulations like the EU's AI Act proliferate, leaders will want to ensure they are in compliance. They'll be asking, 'How do I make sure my current and future uses of AI are safe from risks of regulatory fines and reputational damage?'" — Jennifer Belissent, Principal Data Strategist at Snowflake and a former Forrester analyst
- **And data governance.** "More data is being generated, more parties are interested in the data, and there are more tools for working with the data. And the challenge of governing data is only going to get harder. If you fine-tune an LLM with your proprietary data, sensitive data might emerge. AI compounds the familiar problem because you have less control of and visibility into what the machine is doing. Data lineage, provenance governance just become even more important." — Christian Kleinerman, SVP of Product
- **Consumers will (and should) demand transparency.** "It's possible that the average consumer isn't giving a lot of thought to how AI will affect them. Many people are often comfortable with or inattentive to how businesses use their data. But as LLM-powered AI begins making more decisions that affect us, like whether you can have a loan, a job interview, a medical procedure, people will demand more transparency into how data models affect them." — Jennifer Belissent

CYBERSECURITY: THE AI CHALLENGE

We've already alluded to some of the risks inherent in deploying new AI technology. It's not that the advent of generative AI and LLMs is especially perilous—all new technologies demand that we consider their risks. But AI developments are fast-moving and startling in their breadth and capability, and thus will be extra-challenging to security teams.

CISOs have to guide responsible adoption of these highly valuable tools. Some companies [including [Apple](#), [Amazon](#) and [JPMorgan](#)] have reportedly restricted the use of free AI chatbots and coding assistants to generate text or code because inevitably proprietary data would end up being put into those models. That's an immediate concern. But in the longer term, generative AI is too powerful, too appealing. If security leaders don't provide responsible alternatives, frustrated staff will create workarounds and shadow IT.

CISOs must help strike a balance in which innovation is made accessible, for the good of the enterprise, in ways that limit the risk of compromising sensitive data, facing regulatory censure or incurring reputational damage.





LLMs will be secured in-house.

Earlier, Christian Kleinerman and James Malone discussed the AI supply chain that will allow businesses to construct large and not-quite-large language models within their secured environments. That's essential to good security, says Mario Duarte, Snowflake's VP of Security.

"Of course, businesses will maintain generative AI tools and LLMs within their security perimeters, but this creates its own security challenges," he says.

"First, can you trust the external data sources and open source models? Second, a lot of the risk with any technology has to do with misconfiguration," he says. "Just as user error can leave your data vulnerable on a CSP, human error can make LLM-based tools vulnerable. And no one has experience maintaining and securing those models yet."

Another frequently discussed concern has to do with bad data poisoning a model, particularly if deliberately introduced by an adversary. "While 'inaccurate outputs from our data tools' might seem less like a direct cybersecurity concern," Duarte says, "it's actually a form of social engineering, and thus right in the security wheelhouse."

The AI data supply chain will be a target of attack. Eventually.

Digging into the potential vulnerability of the data itself, it's important to realistically assess the risk. We're talking about an adversary playing a relatively long game by injecting false or biased data into foundational LLMs. Picture a propaganda operation in which a political actor plants content that clouds the truth about a nation-state conflict, election integrity, or a political candidate.

"You basically see that happening already with bots being used on social media to influence the zeitgeist," Duarte says. "So it's not far-fetched to imagine an operation specifically designed to influence a foundational model trained on the open internet."

And from political shenanigans to business attacks is not an enormous leap. "I know of a case where Company A was damaged when a rival fed a journalist a false rumor that A was about to be acquired," Duarte says. "That rumor created a hesitancy among potential customers that hurt Company A and therefore helped the rival."

Plant some stories, misinform a foundational model, and down the line, an LLM could offer inaccurate or secretly biased advice about a certain company or business strategy. But Duarte is quick to note that this should not top any security leader's list of worries.

"The vast majority of cyberattacks are about money, and this sort of attack doesn't have an immediate financial payoff," he says. "So it's not a very practical threat. It's something to be aware of, but I don't see that as being a significant worry in the next 12 months, at the least."

And furthermore, a lot of the solution to this problem is stuff that any good security org is already doing. People can talk about finding new security approaches for new AI tools, but Duarte says that probably the best defense is a longstanding best practice: Make sure your partners and suppliers have earned your trust.

"Most of our existing security controls and practices apply to generative AI as well," he says. "So vetting your vendors' practices and controls remains a sound and effective way to improve your security posture, regardless of these new attack vectors."

Gen AI will improve intruder detection.

There's more good news on the AI + cybersecurity front. One of the biggest problems in IT security is how much time passes between a bad actor gaining access to your systems and you figuring it out (possibly because they successfully complete the attack) and shutting the barn door. Various industry reports put median dwell time at about two weeks, which is a long time when data moves at the speed of light.

Anoosh Saboori, Head of Product Security, says AI will significantly improve automated detection of intruder activities.

"AI can be used to hugely improve different aspects of security, from natural experience of interacting with security products—lowering the bar for non-security experts to ask meaningful questions about the security of the system—to shortening the time it takes to detect anomalies and intrusions, all the way to automated response and forensic analysis. For example, we already have a lot of behavioral data that can help us understand malicious or inconsistent behavior," he says. "Variations compared the baseline for where and when Employee X signs in, and which systems are accessed, can be flagged as someone using stolen credentials or malicious insider activity. Generative AI will be much, much better at judging and flagging that activity."

AI will be a huge boon to cybercriminals before it becomes a help to security teams.

Okay, now the bad news. Cybercrooks will benefit from widespread deployment of advanced AI tools before their targets can set up AI in their own defense.

"Legitimate businesses are careful about adopting and using new technologies—there's cost, regulatory requirements and reputational risk if it's done poorly," says Mario Duarte. "Bad guys won't wait. So at first, they'll have the full firepower of LLMs and gen AI, and defenders will be playing catch-up."

Eventually, he says, the playing field will even out. But expect a lot of pain in the meantime.

Generative AI will make lowbrow cyberattacks smarter.

The idea of cybercriminals deploying one of the most advanced technologies on the planet might conjure images of sci-fi-level malevolence, of elite, sophisticated attackers launching brilliant attacks we don't even have names for yet. Sure, probably, but up first will be a lot of the basic attacks that already work so well.

"Phishing is still a big, big deal," Duarte says, noting that most phishing emails are pretty clumsy and dumb. "Generative AI will make this already effective attack vector more successful. I don't think people are ready for that."

Cyberattackers will continue to shift left.

Our final thought is less an AI issue than a DevOps/DevSecOps matter. Those practices shift testing and remediation to the left of the software development lifecycle diagram (toward developers) and automate much of the right-side production work. A happy benefit of the automation is that it eliminates opportunities for human error, which often gives cybercriminals their entry point.

"The whole 'shift left' aspect of DevOps/DevSecOps is enabled by automation. And automating functionality in the production environment means less human error for attackers to exploit," Duarte says. "So attackers are now looking for ways in via developer environments, because that's where human mistakes can still be discovered and exploited."

It's harder for security teams to defend against such attacks, Duarte adds, because it's much harder to create baselines for acceptable dev activity than for an automated, well-managed production environment.

"Development is by nature chaotic and experimental, so understanding what's normal and abnormal in a dev environment is very difficult," he says. "But we have to figure it out. This is where you throw everything—humans, machine learning, AI—at understanding what suspicious behavior looks like."

And despite the difficulties, he says he's confident that in the long term, security teams will be able to keep up with shift-left attacks.

INNOVATION ABOUNDS

The issues most likely to affect businesses in the next one to five years center on the adoption of large language models and generative AI, and the eternal priority of security and risk management. But that's not to say those are the only things going on with data and technology.

"There are many other technologies that will be impactful or transformative that we're just paying less attention to during this AI craze," says Snowflake SVP Sridhar Ramaswamy.

"That's the nice thing about the world we live in. There's no dearth of impressive things going on, even if in Silicon Valley it looks like gen AI is all that matters. There are lots of things going on that will have enormous impact."



Ramaswamy mentions, off the top of his head:

- **Advancements in self-driving cars.** Developments around autonomous driving are changing the fundamentals of transportation.
- **The ongoing revolution in battery design.** Longer-lasting batteries will not only give electric vehicles more range, but affect stationary power storage, the batteries in whatever you're reading this report on, and the impact of the electrical grid on climate change.
- **Biotech breakthroughs.** He points to the rapid development of the COVID-19 vaccine, which compressed perhaps 10 years of work into something less than 18 months.

"I speak to doctors and bio professors, and they're most excited about the project to map human RNA," he adds. "They think it will be the key to solving many diseases, to detect and cure cancers. The Dean of Medicine at Brown thinks the RNA project will have a bigger impact than the human genome project."

We do indeed live in exciting times. At Snowflake, we're looking ahead to a future in which more data is better governed, more effectively analyzed, and continues to directly contribute to successful business outcomes and positive societal change.

The future is bright. See you there.

CONTRIBUTORS



MONA ATTARIYAN

Mona Attariyan is Snowflake's Director of Machine Learning, and holds a Ph.D. in computer science from the University of Michigan. Previously, she served as Head of Analytics and Machine Learning at Redfin and held various technical roles at Google.



BENOIT DAGEVILLE

Benoit Dageville is the Co-Founder of Snowflake, and today serves as the company's President of Product. Before founding Snowflake in 2012, Benoit spent more than 16 years as a data architect and developer at Oracle.



SUNNY BEDI

Sunny Bedi is the Chief Information and Data Officer at Snowflake. Before that, he spent over a decade at NVIDIA, where he was VP of Corporate IT and Head of Global IT, among other roles. Sunny also serves on the board of DevSecOps company GitLab.



MARIO DUARTE

Snowflake's VP of Security Mario Duarte has more than 20 years of experience as a security professional, having built and managed security teams and developed and implemented security programs for a large number of private and public organizations.



JENNIFER BELISSENT

Jennifer Belissent, Ph.D., is Principal Data Strategist at Snowflake. She spent over a decade at Forrester Research as an internationally recognized expert in data sharing, the data economy, and data leadership, including best practices for building world-class data organizations.



AMANDA KELLY

Amanda Kelly is Co-Founder of Streamlit, which was acquired by Snowflake in 2022, and she now serves as a Product Director focused on bringing Streamlit and other great developer experiences to Snowflake. Prior to that, she was Director of Operations at Zoox and led a team at Google X. Amanda holds an MBA from Stanford.

CONTRIBUTORS



CHRISTIAN KLEINERMAN

Snowflake's SVP of Product Christian Kleinerman has more than two decades of experience working for the world's largest enterprise tech companies. He worked at Google as YouTube's Senior Director of Product Management and held a number of management roles at Microsoft.



SRIDHAR RAMASWAMY

Sridhar Ramaswamy was Co-Founder of Neeva, acquired in 2023 by Snowflake, where he is now a senior vice president. He spent more than 15 years at Google, where he started as a software engineer and rose to SVP of Ads & Commerce. Sridhar earned a Ph.D. in computer science from Brown University.



PRASANNA KRISHNAN

Prasanna Krishnan is Snowflake's Senior Director of Product Management, Collaboration and Snowflake Marketplace. She is the founder of SmartyPal and has held product management positions at Microsoft and an executive position at Jetsetter, acquired by Tripadvisor. Prasanna holds an MBA from Wharton.



ANOOSH SABOORI

Snowflake's Head of Product Security, Anoosh Sabori, holds a Ph.D. in computer and information systems security/information assurance from the University of Illinois. Prior to Snowflake, he held product management roles at Google Cloud and Microsoft.



JAMES MALONE

James Malone is Director of Product Management at Snowflake. Prior to Snowflake, he held a number of technical management roles at large companies such as Google, Disney and Amazon. James holds an MBA from the University of Washington.



ADRIEN TREUILLE

Adrien Treuille is Co-Founder of Streamlit, and today focuses on next-gen LLM opportunities at Snowflake. Previously, he was VP of Simulation for autonomous vehicles at Zoox and an AI project leader at Google X. Adrien also taught computer science at Carnegie Mellon for eight years.



ABOUT SNOWFLAKE

Snowflake enables every organization to mobilize their data with Snowflake's Data Cloud. Customers use the Data Cloud to unite siloed data, discover and securely share data, and execute diverse artificial intelligence (AI) / machine learning (ML) and analytic workloads. Wherever data or users live, Snowflake delivers a single data experience that spans multiple clouds and geographies. Thousands of customers across many industries, including 639 of the 2023 Forbes Global 2000 (G2K) as of July 31, 2023, use the Snowflake Data Cloud to power their businesses.

[Learn more at snowflake.com](https://www.snowflake.com)



© 2023 Snowflake Inc. All rights reserved. Snowflake, the Snowflake logo, and all other Snowflake product, feature and service names mentioned herein are registered trademarks or trademarks of Snowflake Inc. in the United States and other countries. All other brand names or logos mentioned or used herein are for identification purposes only and may be the trademarks of their respective holder(s). Snowflake may not be associated with, or be sponsored or endorsed by, any such holder(s).