
Capgemini
Data-powered
Innovation
Review
Wave 10

Foreword

Niraj Parihar

Chief Executive Officer,
Insights and Data, Capgemini

As you flip through this edition of the Data-Powered Innovation Review, you'll notice the intricate patterns of desert sand etched across the pages. These aren't just beautiful visuals – they speak to the very nature of data and AI in 2025: structured yet shifting, shaped by invisible forces, and constantly redrawn by time, pressure, and perspective.

Just like the desert surface, our business landscapes are being reshaped – not in centuries, but in quarters, months, weeks even. Agentic AI, small-scale models, embedded intelligence, autonomous teammates, and future-first architectures are emerging not as futuristic ambitions, but as active forces in the now. The articles in this issue don't simply describe what's changing – they show how the very foundation of business, strategy, and decision-making is being reimagined.

There is power in structure: in building models, processes, compliance frameworks, and digital twins. But the stories in this edition remind us that adaptability is now just as critical. From rebooting your business from the year 2030, to embedding AI as a co-pilot in manufacturing and policymaking, to challenging the "bigger is better" myth of large-scale language models – this issue champions speed, frugality, resilience, and the courage to break patterns, not just refine them.

It's not about plugging AI into legacy workflows and hoping for the best. It's about embracing a shift toward AI-native thinking, where value creation, risk management, supply chain intelligence, and customer engagement are all redesigned with data-powered systems that learn, adapt, and act – often before we ask them to.

Yet amid this transformation, one thing remains constant: our responsibility. As much as these technologies offer scale and speed, they also raise serious questions about sustainability, bias, trust, and control. It is our job not only to innovate but to curate the future we are building, and to shape the dunes, not just admire their patterns.

On behalf of everyone contributing to this tenth edition, I invite you to explore these shifting landscapes and to see in them both a challenge and an opportunity. The sand may move, but our direction need not. Let's build with purpose, not just precision.

Here's to thinking from the future – and acting in the present.

Editor's note

Every edition of the Data-Powered Innovation Review starts with a spark – an idea, a conversation, a glimpse into the future we're all building (or dodging, for that matter). And while for some of us this may be the beginning of a well-deserved summer break, others are deep in delivery mode or strategizing for what's next. Wherever this edition finds you though, we hope it offers the right balance: light enough to browse, rich enough to spark action.

You'll find plenty of materials in this magazine to inspire that, from (but of course) agentic AI to frugal LLM models, humanoid teammates, AI-native supply chains, and compliance engines with a bite. This mid-year edition maps a landscape that's moving fast amid highly unpredictable times. And for a change, that is far from the cliché it often turns out to be. But there's a direction too: we're moving from simple AI plugins to built-in, native intelligence. From systems that support to systems that co-create. From workflows to networks. From pilots and proofs-of-concept to full-on enterprise reboots. The pendulum swing keeps going on.

If you'd like to keep the inspiration coming, our Data-Powered Innovation Jam podcast continues to explore new ideas with a mix of insight, a slight touch of irreverence, and quite a few unexpected musical references. Hosted by our AI rockstars Robert Engels and Weiwei (plus me), it's your companion to the magazine – or simply a soundtrack to thinking out loud. You can find it wherever you get your podcasts, or on our page.

Ron Tolido

Chief Technology Officer,
Insights and Data, Capgemini

Thank you to all the contributors, collaborators, and creators of this edition – with a special mention of our partner blackshark.ai. Whether you're scanning this between meetings, on the move, in a moment of pause, or with your feet firmly on the sand of the beach, we hope it stirs something worth pursuing.

And if the dunes shift tomorrow? All the more reason to read today!

Table of content

The serendipity engine 06
Agentic AI is what personalization was trying to be
Neerav Vyas, Capgemini

Factory settings: Human plus humanoid 10
How robots that look like us are reshaping the workplace
Alexandre Embry, Capgemini Engineering

Smart and green 14
Supply chain resilience – the AI way
Sudarshan Sahu, Capgemini

Nature, engineered 18
AI and quantum computing redesigns enzymes – and the R&D playbook
Franziska Wolff, Capgemini Engineering
Fabian Gwinner, Capgemini Engineering

Insights, upgraded 22
Bringing agentic AI to KPIs
Dnyanesh Joshi, Capgemini

Pixels with purpose 26
AI meets geospatial data to make maps think and satellites work smarter
Brett Clark, blackshark.ai
Weiwei Feng, Capgemini

Your business called. It wants a reboot from the future 30
Plugging AI into an org chart isn't transformation
Dinand Tinholt, Capgemini

Gen Garage 34
Where tomorrow's talent builds today's AI for good
Aishwarya Kulkrni, Capgemini

Ø ETL 38
Less extract, more impact
Sjoukje Zaal, Capgemini

Born today, leading tomorrow 42
How AI teammates will transform the way we collaborate, create, and lead
Hemavathi Thiruppathi, Capgemini
Venkadesh Thirupathi, Capgemini

The fast and the frugal 46
Outrun the big models – without draining the grid
Joakim Nilsson, Capgemini
Johan Müllern-Aspegren, Capgemini

The compliance upgrade 50
From policy overload to AI-powered safety and alignment
Dr Abi Giles-Haigh, Capgemini
Liz Henderson, Capgemini

Mind the gateway 54
Between the platform and the promise: Gateways keep Gen AI on track
Kumar Chinnakali, Capgemini

Agents of confidence 58
Designing trust, autonomy, and alignment in an AI-native world
Robert Engels, Capgemini
Mark Roberts, Capgemini Engineering

© Capgemini 2025. All rights reserved

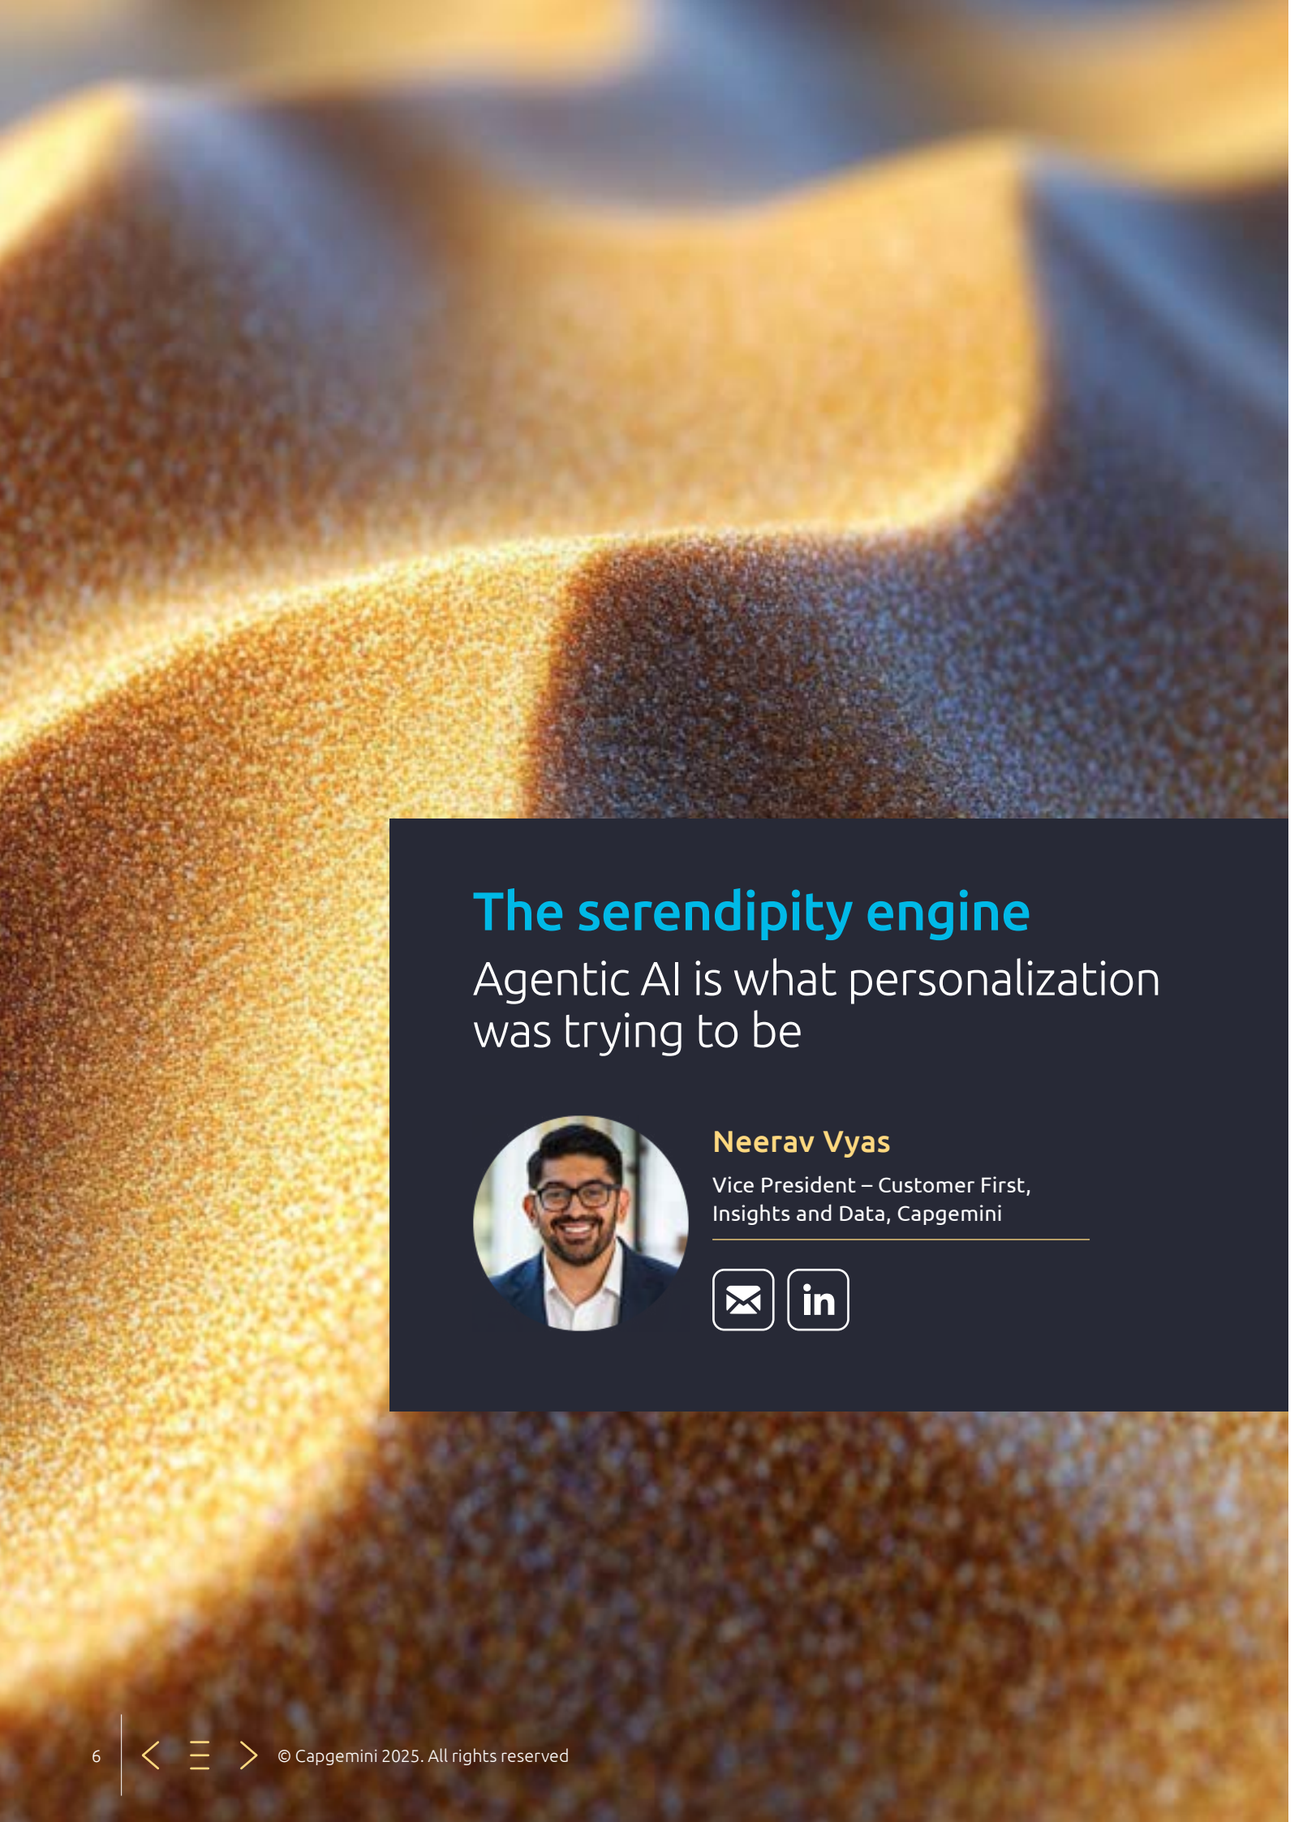

The serendipity engine

Agentic AI is what personalization was trying to be

Neerav Vyas

Vice President – Customer First, Insights and Data, Capgemini

It's not about personalization. Personalization is trite. We need instead to think about delivering serendipity. Serendipity at scale or obsolescence with haste are the challenges of modern brands. In a world of AI agents, how we create and scale experimentation cultures to support and grow personalization programs looks vastly different than how the pioneers at Amazon and Netflix sustained their success. Those who embrace an agentic-forward approach will leapfrog the competition and hockey-stick the impact on business outcomes by vastly increasing the pace of successful serendipitous personalization.

The promise of tailoring experiences has been a game of pattern recognition for years. We built elaborate taxonomies of user behavior and crafted recommendations based on digital breadcrumbs left behind by customers. This was akin to understanding a city solely by its traffic patterns. Agentic systems promise something far more profound: the ability to know the individual resident, their unspoken needs, their fleeting desires, before they themselves are fully aware. They'll also do so while dramatically lowering the effort of humans to deliver and support these experiences. Agentic systems will deliver serendipitous experiences at unprecedented scale.

Imagine experiences where we're not reacting to clicks or linear events in your journey but anticipating your intellectual wanderings. Picture agents assigned to you, digital familiars constantly observing and assisting, learning your explicit and implicit choices and the subtle shifts in your attention, focus, and intent. This isn't about serving you more of the same; it's about understanding the trajectory of your interest, the nascent idea struggling to form in your mind, and proactively offering the missing piece. This constantly builds serendipitous recommendations to elevate your experiences.

Imagine an interface that doesn't just present information, it curates it. It knows when to offer a broad overview and when to delve into granular detail. This isn't personalization as we know it; it's individualization, custom digital experiences crafted in real time that understands you in a way that feels almost intuitive. This is hardly far-fetched; Amazon is already a hyper tailored personalized experience where the UI for the site is customized to each customer.

We're moving to a world where the arc of the customer journey is designed and orchestrated by AI agents across different platforms and touch points. Linear sales are archaic remnants of the past. As we move forward, users will experience dynamically generated paths optimized for their individual needs and motivations with agents orchestrating AI interventions at a pace and level of complexity that far exceeds what linear and rule-based journeys can achieve.

How will we know if these hyper serendipitous experiences are truly working? At Amazon and Netflix, the key metric of success for personalization teams was the number of experiments launched. They knew trying to predict winners and losers was folly. They found that launching more experiments meant uncovering more winners, which meant more innovations that improved business outcomes.

Our current methods, A/B tests and multivariate analyses, feel almost crude in the face of agentic potential. We meticulously craft variations, divide our audience, and wait for the statistical tea leaves to settle. But agentic AI allows for a far more fluid, almost organic approach to experimentation.

Imagine thousands of micro-experiments blooming simultaneously, each tailored to the unique contours of an individual user. Instead of broad strokes, we can paint with the finest of brushes, understanding not just if a change works, but for whom, and under what specific circumstances. Picture AI not just testing pre-defined hypotheses, but generating them, sifting through the intricate tapestry of user interactions to uncover patterns and possibilities that would remain invisible to the human eye.

The very notion of a static control group begins to dissolve. Agentic systems can dynamically allocate users to different experimental branches, not through rigid segmentation but through an ongoing assessment of individual fit. Successful variations aren't rolled out in phases; they are subtly woven into the experience of those for whom they are most likely to resonate, a real-time optimization that feels less like a deployment and more like a natural evolution.

> *Imagine an interface that doesn't just present information, it curates it. It knows when to offer a broad overview and when to delve into granular detail. This isn't personalization as we know it. It's individualization.*

Of course, this brave new world is not without its shadows. The sheer volume of data required to fuel these autonomous agents raises profound questions about privacy and trust. The potential for algorithmic bias to be amplified, creating digital echo chambers of unprecedented intensity, is a genuine concern. And the very autonomy of these systems demands a new level of vigilance, a constant grappling with the ethical implications of algorithms that can act, decide, and shape our experiences with minimal human oversight.

The transition to an agentic AI-driven world will not happen overnight. It will require significant advancements in AI research, infrastructure, and ethical considerations. However, the potential to create truly individualized and adaptive experiences, coupled with a more dynamic and insightful approach to experimentation, makes the agentic horizon an incredibly exciting one. Businesses that embrace this paradigm and proactively address the associated challenges will be well-positioned to unlock unprecedented levels of customer engagement, satisfaction, and innovation. The future of personalization and experimentation is no longer just about understanding users; it's about having AI agents that can truly understand each user and act in their best interest.

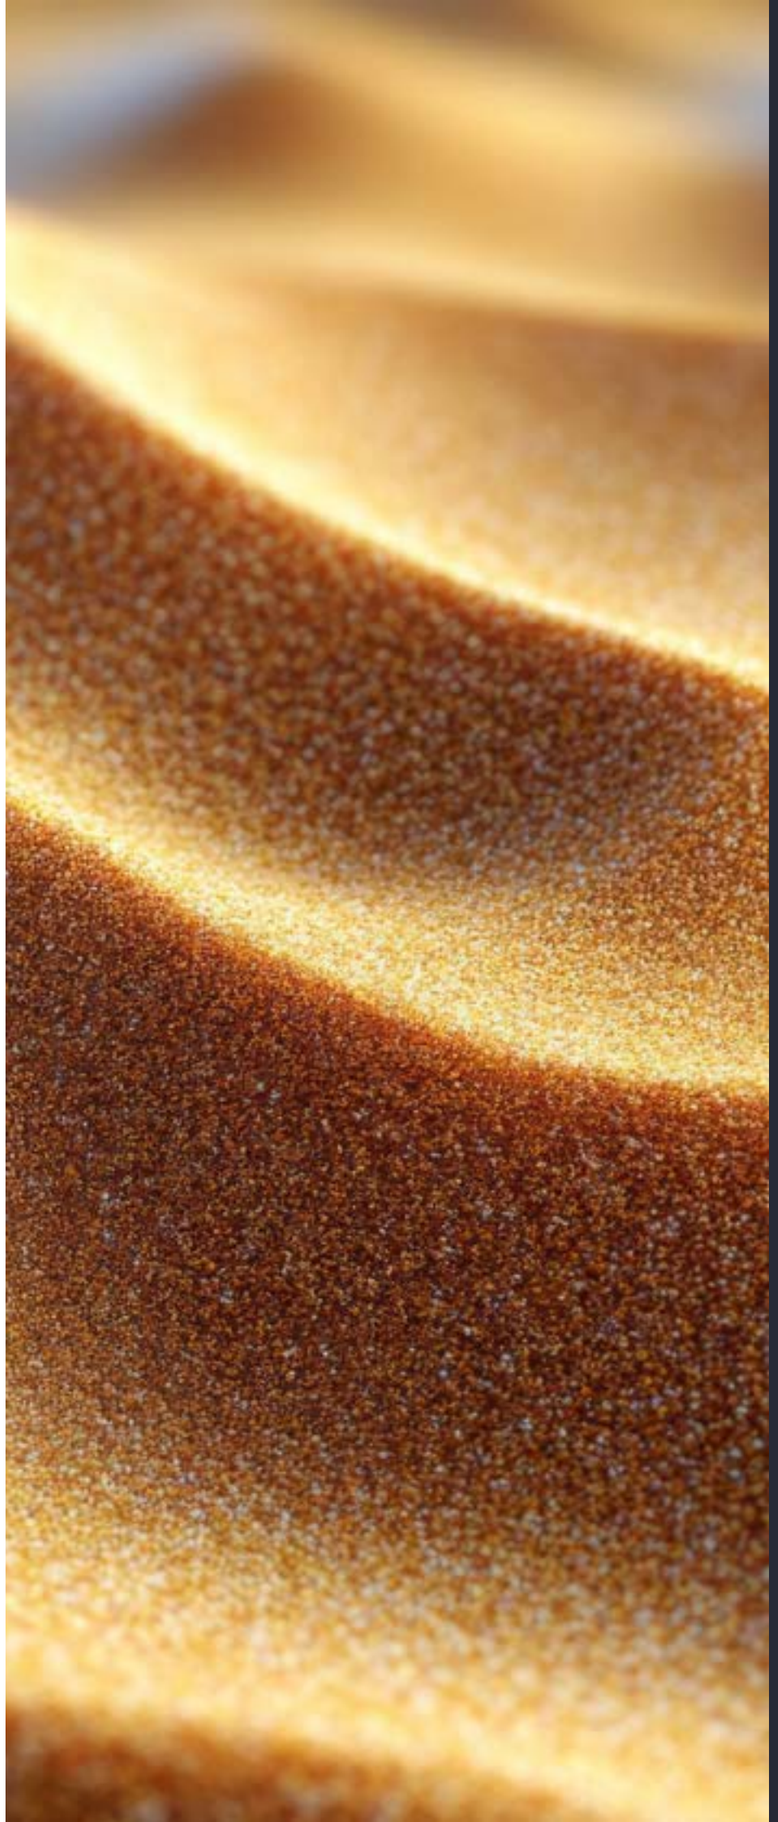

Start innovating now

Hyper-personalized experiments

Instead of testing a single change on a broad audience, agentic systems can conduct thousands of micro-experiments simultaneously, tailoring variations to individual users based on their agent's understanding. This allows for the identification of highly granular and context-dependent optimizations. Begin with leveraging agentic systems to assist with more micro-segments, starting in the dozens and gradually increasing over time.

Dynamic experiment allocation

Instead of fixed traffic splits, agentic systems can dynamically allocate users to different experiment variations based on their individual characteristics and the predicted effectiveness of each variation. This can lead to faster learning and reduced opportunity cost. Start with leveraging agentic systems to create your test versus control plans, and leverage multi-armed bandits to assist with split allocation and optimization.

Autonomous hypothesis generation

Agentic AI can go beyond testing pre-defined hypotheses. By continuously analyzing user interactions and outcomes, it can identify patterns and autonomously generate new hypotheses for experimentation, uncovering insights that humans might miss. Begin by leveraging agentic systems to summarize insights from past experiments and gradually evolve to leveraging agents to suggest new hypotheses to be tested and validate.

#SerendipityAtScale #Serendipity #AgenticAI #AgenticSerendipity #AgenticCX #AgenticExperiences #Personalization

Factory settings: Human plus humanoid

How robots that look like us are reshaping the workplace

Alexandre Embry

Vice President, CTIO, Head of the Capgemini AI Robotics and Experiences Lab, Capgemini

Once confined to science fiction, humanoid robots are stepping onto factory floors – not to replace workers, but to work alongside them. With the convergence of AI, robotics, spatial computing, and digital twins, enterprises now face a profound shift: automation with arms, legs, and reasoning skills. These human-shaped machines can adapt to existing environments, learn new tasks, and scale operations without disruption. But the real breakthrough isn't just technical, it's collaborative. Humans, humanoids, and agentic AI systems are about to become one team.

More flexibility, less disruption, better scaling

Until recently, human-shaped robots – complete with a head, arms, and legs – were only speculative. But the first production models are now a reality from companies such as California-based Figure AI, which is building a factory expected to manufacture 12,000 humanoid robots per year. These will be game-changers for enterprises across multiple sectors – including manufacturing, life sciences, automotive, aerospace, defense, energy, utilities, and consumer products.

Robots and automation are nothing new on the factory floor, but legacy deployments involved purpose-built machines and dedicated assembly lines. When manufacturing changes were required, a company would have to specify new robot designs and rebuild factories – which translated into significant financial investment and production disruptions.

Humanoid robots address these drawbacks by making the robot as adaptable as a human worker – capable of mimicking human gestures thanks to the rapid development of sensors and other hardware. What's more, they are no longer just machines: humanoids are autonomous and adaptable physical reasoning agents equipped with cognitive capabilities. Artificial intelligence merges with robotics to take a physical form, making this the next big thing in AI and leading to the ultimate stage of automation on the shop floor.

Humanoids can be deployed to automate brown field operations without rebuilding, and humanoids can be easily integrated into existing industrial operations to perform undesirable, dangerous, labor-intensive, or repetitive tasks. Organizations can start small with a few humanoids working on focused activities, and then simply add more robots to scale up over time. And when production needs change, companies will no longer have to retool. Instead, they can retrain both human employees and humanoid robots to undertake new operations in flexible workspaces.

Investigating the pathways to true convergence

That said, there's work to be done to unlock the full potential of humanoids. These robots consist of an incredible package of sensors, controllers, motors, processors, and other hardware. At Capgemini's dedicated AI Robotics and Experiences Lab, we're exploring how to apply our expertise in agentic AI and LLM, computer vision, digital twins, data analytics, robotics, and sector-specific industrial processes to humanoid robots. Our goal is to help shape the technical convergence of human-centered, digital-physical interactions between humans, systems of robots, AI-agents.

The hub of this convergence is a new virtual space in which most of the digital-human interactions will happen. In this space, unified and pre-trained data models, AI agents, edge AI, and digital twins merge with an interaction layer that leverages technologies such as real-time 3D and spatial computing. Data comes from an aggregation layer linking various data sources – including IS/IT/OT systems, sensors, and machines. Agentic AI and physical AI then enable autonomous digital-physical interactions, with virtual AI-agents being able to act in the real world by interacting with AI agents housed in robots.

What we do at AI Robotics & Experiences Lab

Technical convergence of **human-centered natural digital-physical interactions** with systems of **robots, agents, machines.**

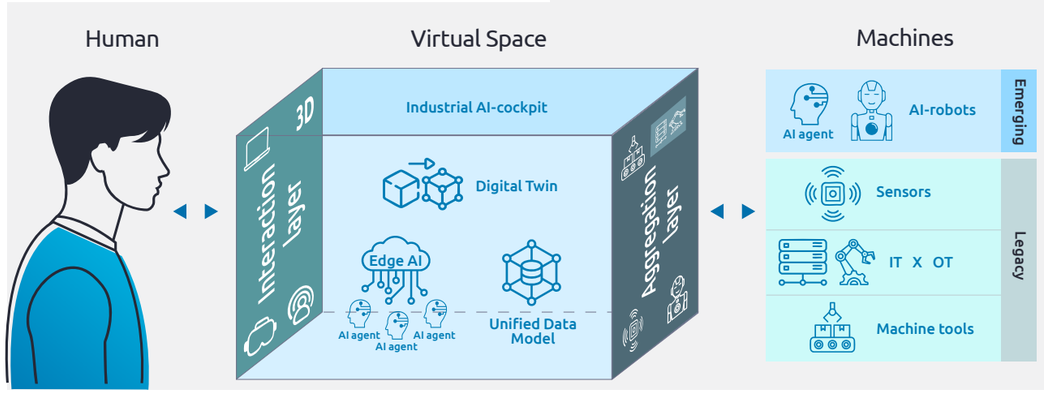

The challenge is how to best leverage technology enablers to unlock the productivity gains of humanoid robots with minimal operational disruption. To address this, Capgemini is adding agentic AI-enabled decision-making capabilities to a network of interconnected digital twins that replicate a physical industrial environment. These enhancements apply several key concepts in robotics and AI research:

- Vision language action models, which enable agents to understand and execute natural language instructions in a visual environment
- Reinforcement learning and simulation to transfer learning, which enable the sharing of knowledge from various data sources, such as videos or motion capture, to train in virtual environments and then translate into real-world scenarios
- Teletraining, which enables a human to remotely control the agent to demonstrate desired actions
- System 1 and system 2 reasoning models, which are dual-process architectures that mimic a human brain's ability to both generate fast, reactive responses and engage in slower, deliberate planning.

Collaboration in action

On the factory floor of the very near future, Capgemini envisions humans, AI-robots, and multi-agent systems working as a team to enhance operational efficiency, precision, and safety. Each participant will make important contributions to this, leveraging their unique abilities.

Robot humanoid

- **Precision handling:** Humanoid robots ensure precise handling and placement of parts, reducing the risk of damage and ensuring consistency.
- **Efficiency:** Robots can perform repetitive tasks quickly and accurately, increasing throughput and reducing cycle times.
- **Safety:** By taking over potentially hazardous tasks, robots enhance workplace safety and reduce the risk of injury to human operators.

Human

- **Skill and adaptability:** Human operators bring their expertise and adaptability to the assembly process, handling complex tasks that require fine motor skills and decision-making.
- **Guidance and oversight:** With AR projections, humans receive real-time instructions and visual cues, ensuring accuracy and reducing the likelihood of errors.
- **Problem-solving:** Humans can quickly identify and address unexpected issues, providing a level of problem-solving that complements automated systems.

Multi-agent System

- **Real-time monitoring:** The agentic systems continuously monitor the assembly process and the workstation environment, ensuring each step is executed correctly and identifying any deviations.
- **Anomaly detection:** The system detects potential anomalies and provides alerts and corrective actions, maintaining workflow integrity and preventing errors.
- **Quality control:** The agentic system makes real-time decisions on whether parts require inspection, ensuring only high-quality products proceed to the final assembly.

Human operators will bring their expertise and adaptability to bear upon the production process. Using spatial computing, they'll supply guidance and oversight. They'll also handle complex tasks that require fine motor skills and decision-making. And because humans can quickly identify and address unexpected issues, they'll provide a level of problem solving expertise that enhances automated systems.

Humanoid robots will ensure precise handling and placement of parts, reducing the risk of damage and ensuring consistency. They can also perform repetitive tasks quickly and accurately, increasing throughput and reducing cycle times. And by taking over potentially hazardous tasks, robots will enhance workplace safety by reducing the risk of injury to human operators.

Agentic systems will continuously monitor the operational processes and environment to ensure each step is executed correctly. They will identify deviations, alert the team, and suggest corrective actions to help maintain workflow integrity and prevent errors. They can also make real-time decisions about whether operations require additional action or reinforced quality inspection process. Eventually, robots driven by virtual agents will be able to perform some of these actions.

Built to Collaborate

We're resetting the factory settings — not to replace the human workforce, but to reconfigure it. With humanoid robots now stepping out of fiction and into physical workflows, a new kind of collaboration is taking shape. These machines may walk like us and learn like us, but their greatest value lies in how well they work with us. The future of industry won't be defined by humans or humanoids alone but by how seamlessly they operate as one team. The convergence is here. The collaboration starts now.

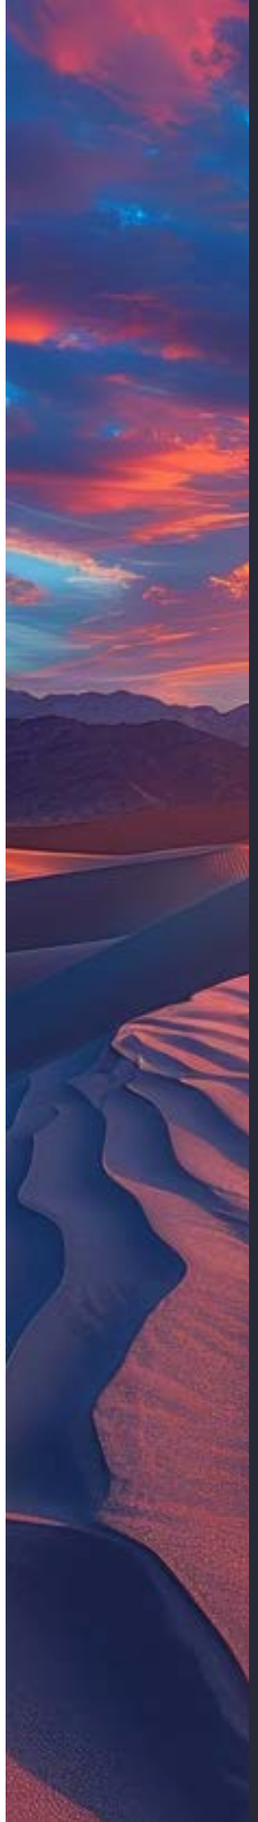

Automation now comes with arms, legs, and reasoning skills.

Start innovating now

Establish your foundation

Leading companies are already undertaking complete digital transformations of their operations. This is essential for deploying humanoid robots – so make completing these transformations a priority.

Don't delay

Humanoid robots are here today and will soon be deployed in real-world environments – across multiple industrial sectors – to provide early adopters with significant competitive advantages.

Support, don't replace, humans

The best outcomes will be achieved if humanoids are deployed as part of teams that also include human workers and multi-agent systems. Now is the time to determine how each team member's strengths can best be leveraged.

#AIRobotics #PhysicalAI #EmbodiedAI #DigitalTwin

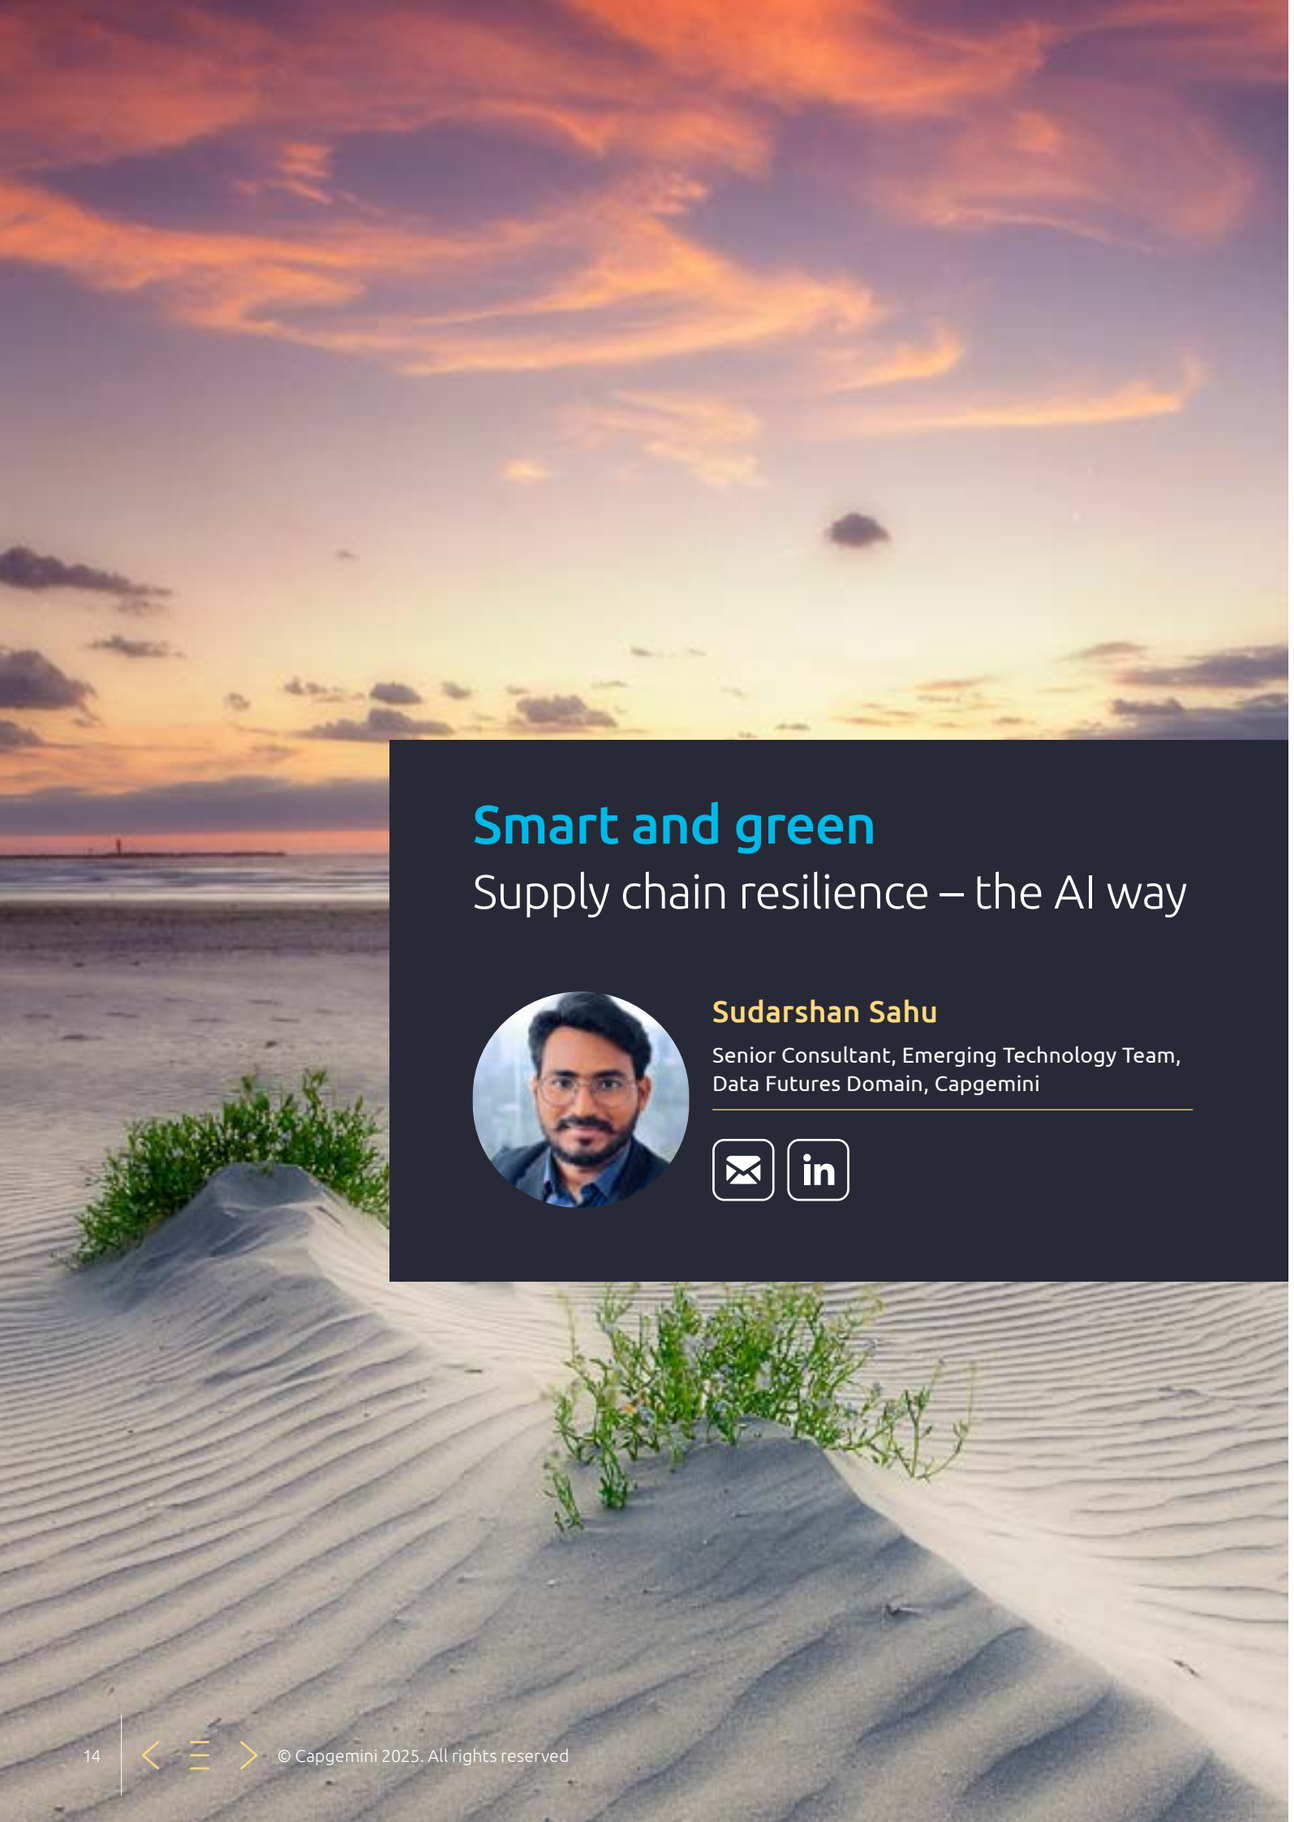

Smart and green

Supply chain resilience – the AI way

Sudarshan Sahu

Senior Consultant, Emerging Technology Team, Data Futures Domain, Capgemini

Climate change isn't a distant threat. It's a reality to deal with now. Businesses need to rethink how they operate, especially when it comes to supply chains, which are crucial for global trade. Just like in the movie Interstellar, where survival depended on data, AI, and adaptability, today's supply chains need to be flexible and smart to handle disruptions and climate challenges. AI-powered insights and actions are like the movie's robot TARS: helping predict risks, optimize logistics, and reduce waste. Data ensures that every decision is as precise as a gravity equation. AI enhances precision in supply chains by analyzing vast amounts of data in real time, predicting risks, and optimizing logistics. It's the key to transforming supply chains into smarter, greener, and more resilient systems that balance profitability with ecological responsibility.

Supply chains aren't just stretched – they're under siege. Disruption is no longer the exception; it's the norm. That's why resilience – the ability to anticipate, adapt, and recover fast – has shifted from nice-to-have to non-negotiable. A recent report from The Business Continuity Institute delivers the reality check: **80 percent of organizations** faced supply chain disruptions last year, most more than once. That's an uptick and, despite better planning, proof that we're still reacting more than we're preparing. Meanwhile, sustainability pressures are mounting. With supply chains responsible for over **60 percent of global carbon emissions**, according to the World Economic Forum, they're no longer just operational engines – they're climate liabilities, too.

Let's face it: what we're doing right now isn't cutting it. The cracks in our supply chains are showing, and incremental fixes won't be enough. It's time for bold moves. If we want supply chains that can truly withstand shocks and stay ahead of the curve, we need to lean into smarter, faster, more adaptive solutions. That's where AI steps in, not just as a tool but as a game-changer. With its ability to forecast disruptions, optimize operations, and accelerate response times, AI is shaping the supply chains of the future. To stay ahead, companies must embrace green supply chain management (GSCM), where sustainability is built into every step. AI supercharges this shift, turning GSCM into a smart, data-driven engine. From cutting carbon to driving circular economies, AI enables supply chains that are not just efficient, but truly green.

Resilience is not yet autonomous

Supply chains are navigating a perfect storm: geopolitical instability, extreme weather, shifting consumer expectations, and growing uncertainty in global trade. Disruptions are no longer outliers; they're part of the operating environment. While many organizations are embedding risk management into supply chain strategy, execution is still stuck in manual mode. Too much effort goes into collecting, cleaning, and stitching together data, which leaves little room for insight, foresight, or speed. AI and machine learning are still underused, and critical response actions often rely on human intervention alone. The result? Slow reactions, mounting workloads, and talent focused on firefighting instead of forward-thinking.

What's missing? It's technology that doesn't just capture and store data, but actively turns it into prescriptive insights and clear, actionable recommendations. Unfortunately, most tools in the market today still fall short of that promise. Instead, businesses are left stitching together manual processes and siloed teams to make sense of a rapidly changing environment. To build truly resilient supply chains, we need to shift from reactive, human-heavy models to intelligent, tech-augmented systems. The future isn't about replacing people; it's about empowering them with tools that amplify their decision-making, speed up response times, and free them to focus on what matters most.

How AI and data are changing the game

Data and AI are at the core of this transformation, delivering unmatched insights, predictive accuracy, and optimization potential. By leveraging real-time data and predictive analytics, AI can identify potential risks – such as supplier delays, extreme weather, or geopolitical issues – before they impact operations. This early warning capability allows businesses to proactively mitigate threats through alternative sourcing, dynamic rerouting, or inventory adjustments. AI also enables scenario modeling, helping organizations test various disruption scenarios and build contingency plans with data-backed confidence. As a result, companies can maintain continuity, reduce downtime, and ensure customer satisfaction, even in the face of unexpected challenges. In today's volatile global environment, AI is no longer a luxury but a critical enabler of resilient and future-ready supply chains.

AI-enhanced supply chain resilience framework

The AI-enhanced supply chain resilience framework strengthens supply chain agility and robustness by harnessing advanced AI technologies. It integrates real-time data from IoT devices into a centralized system for comprehensive analysis. Through predictive analytics and machine learning, the framework forecasts demand and detects potential risks – like supplier disruptions or market shifts – enabling proactive risk mitigation and smarter decisions in areas like inventory and logistics.

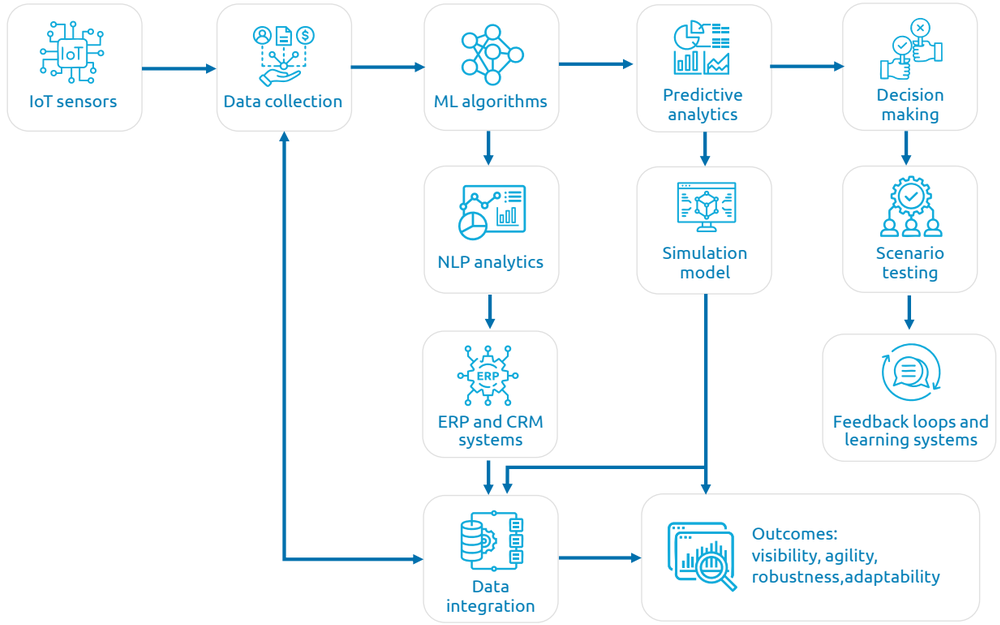

AI-driven communication tools improve collaboration with suppliers and stakeholders, ensuring seamless, transparent information flow. Continuous monitoring and adaptive feedback loops allow the supply chain to respond swiftly to changing conditions, driving ongoing improvement and innovation. By adopting this framework, businesses gain end-to-end visibility, reduce vulnerabilities, and ensure operational continuity, ultimately building a more resilient and high-performing supply chain.

Leveraging AI enables businesses to streamline operations, improve efficiency, cut costs, and elevate customer experiences. One powerful application is demand forecasting, in which AI analyzes historical data to accurately predict customer needs. This leads to smarter inventory management, minimizing overstock and stockouts while optimizing capital use. Another key use case is route optimization. AI-driven tools evaluate factors like weather, traffic, and transport costs to determine the most efficient delivery paths. This reduces time and expenses while ensuring faster, more reliable service that meets growing customer expectations.

How organizations can harness it effectively:

Areas	Real-time supply chain visibility	Predictive risk and disruption management	AI-optimized transportation and logistics	Smart demand forecasting	Sustainable supplier selection and monitoring	Emissions monitoring (Scope 1, 2, & 3)	Waste and packaging reduction
Trends, subtopics, and solutions	Organizations are moving from siloed systems to real-time, unified visibility across multi-tier supply chains	Supply chains are shifting from reactive to predictive by leveraging AI for scenario-based risk modeling	Carbon is being treated as a cost variable; logistics is now optimized for sustainability and speed	ML models now incorporate real-time data (weather, promotions, social sentiment) to minimize overproduction	AI is being used to evaluate and onboard suppliers based on carbon impact and ethical sourcing	With ESG and CSRD mandates, companies are automating emissions tracking beyond direct operations	Packaging is being redesigned using AI to reduce material use and emissions while maintaining performance
	IoT integration, cloud platforms, digital twins	Risk intelligence, climate impact modeling, external data ingestion	Smart routing, last-mile optimization, load consolidation	Machine learning forecasts, external signals, demand sensing	ESG risk scoring, tier-2 visibility, supplier collaboration	Carbon accounting, Scope 3 modeling, integrated reporting	Smart packaging design, waste analytics, lean production
	Implement AI-powered control towers and IoT sensors for live tracking of inventory, emissions, and disruptions	Use machine learning to predict climate, labor, or geo-political risks and trigger automated contingency plans	Deploy AI to reduce empty miles, optimize delivery routes, and improve fleet utilization	Implement demand sensing platforms to right-size production and prevent unnecessary emissions and waste	Use AI-based platforms to rate suppliers on ESG, emissions, and compliance performance	Implement carbon analytics tools that integrate supplier and logistics emissions into enterprise reporting	Use analytics to pinpoint where waste occurs and feed insights into product and packaging redesign
			Value proposition				
Use cases	Unilever uses real-time dashboards powered by AI to monitor over 150,000 suppliers globally – enabling quick, sustainable decisions	Cargill leverages climate forecasting and AI to adjust transport routes proactively, avoiding disruptions and reducing GHG emissions from detours	DHL cut CO_2 emissions across Europe using AI-powered routing that boosted fuel efficiency and reduced idle time	Walmart uses store- and SKU-level ML forecasts to reduce food spoilage and avoid unnecessary restocks – saving costs and emission	Apple evaluates its global supplier base using renewable energy and carbon benchmarks to ensure responsible, low-emission sourcing	Microsoft uses automated AI-based platforms to track Scope 1–3 emissions across its supply chain and tie them to internal carbon pricing mechanisms	Nestlé uses AI to analyze product line data and reduce overpacking, saving thousands of tons of plastic and improving recyclability

According to IDC, 55 percent of Forbes Global 2000 OEMs are projected to have revamped their service supply chains with AI and, by 2026, **60 percent of the top 2,000 companies** based in Asia will use generative artificial intelligence (Gen AI) tools to support core supply chain processes as well as dynamic supply chain design, and will leverage AI to reduce operating costs by five percent. This signifies a widespread adoption of AI to improve efficiency and gain a competitive advantage in supply chain management. Further, generative AI can be harnessed to monitor global events and proactively identify emerging risks. It can automatically generate risk assessments, simulate scenarios, and suggest strategic mitigation plans, empowering supply chain teams to manage risks more effectively. Its conversational interface enhances user experience and accelerates response times. Over time, this evolves into a system-guided, data-driven approach, drawing from a rich library of scenarios and mitigation strategies to deliver contextual, timely responses to risk events.

Considering all of the facts

The fusion of data and AI isn't just a tech upgrade – it's a strategic shift for building supply chains that can bend without breaking. Organizations that embed intelligence into their operations now won't just survive the next disruption – they'll lead the transition to greener, faster, more adaptive ecosystems. By the end of 2025, global supply chains will be reengineered out of necessity and powered by innovation. AI will help companies but it will also help nations stay resilient, competitive, and climate-conscious. It will redefine how we make, move, and manage everything. And like TARS in Interstellar, the most effective systems won't just follow instructions – they'll anticipate, adapt, and act as true copilots. What supply chains need now isn't just visibility, it's vision.

> *In a world shaped by disruption and defined by climate urgency, data is gravity and AI is the force that holds tomorrow's supply chains together. Smart, green, and resilient isn't a goal: it's the new baseline for survival and success.*

Start innovating now

Give your supply chain an AI-enabled sixth sense

Plug your supply chain into real-time feeds – from IoT sensors to storm trackers – and let AI act like your all-seeing oracle. Spot trouble (like delayed shipments or political curveballs) before it hits the fan.

Make generative AI your strategic partner

Leverage generative AI to generate real-time risk assessments, simulate disruption scenarios, and recommend mitigation strategies, all in a conversational interface.

Build a digital twin – your virtual supply chain lab

Think of it as a flight simulator for your supply chain. A digital twin lets you mirror operations in a virtual space to test "what-if" scenarios – from port delays to carbon constraints – without breaking a sweat in real life.

#SmartSupplyChain #ResilientSupplyChain #SupplyChainInnovation #SupplyChainTech #AIForGood #TechForSustainability

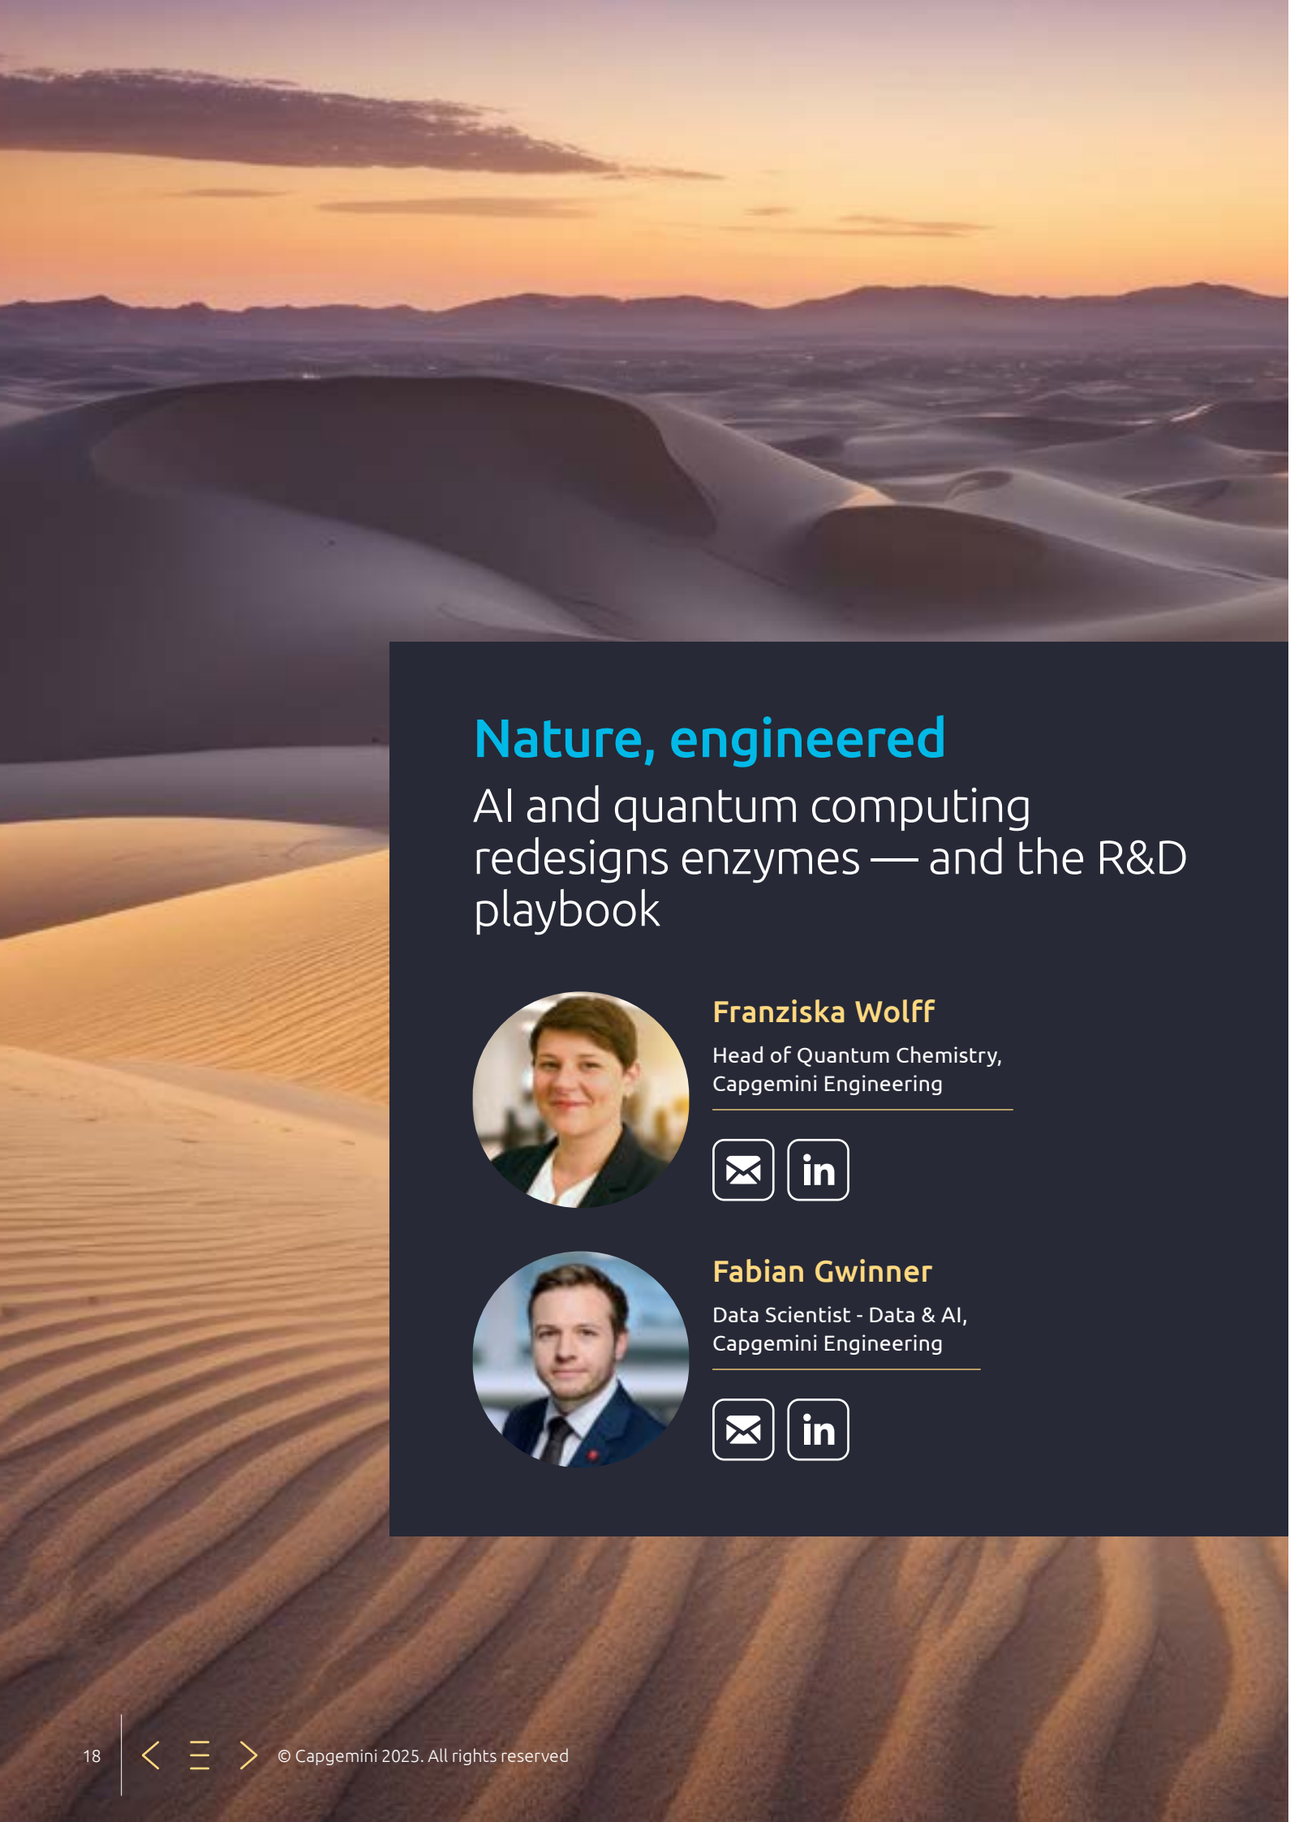

Nature, engineered

AI and quantum computing redesigns enzymes — and the R&D playbook

Franziska Wolff

Head of Quantum Chemistry, Capgemini Engineering

Fabian Gwinner

Data Scientist - Data & AI, Capgemini Engineering

Metalloenzymes are fascinating proteins with wide-ranging applications – from pharmaceuticals and skincare to agriculture and sustainable materials. We're pioneering a new approach to accelerate and transform how these proteins are designed by combining AI models for drug discovery with quantum computing. This hybrid method brings unmatched precision to enzyme engineering and opens new doors for developing high-performance, eco-friendly solutions across industries – from smarter medicines to greener consumer products.

Nature has perfected biochemical processes over billions of years. Today, we can deliberately engineer these processes to create better, healthier, and more sustainable products. Traditionally, understanding and designing these processes has been inefficient and expensive due to limited knowledge of molecular interactions within enzymes, leading to trial-and-error approaches. AI and quantum computing tools are invaluable here, significantly accelerating these processes by enabling accurate digital modeling of compounds, enzymes, and proteins, and testing various modifications to protein structures or binding ligands. A groundbreaking approach is now combining classical AI with quantum computing to design metalloenzymes.

Metalloenzymes, proteins incorporating metal ions as essential catalytic centers, drive critical reactions like oxygen transport and pH regulation. They are crucial for drug development and offer potential for industry and consumer innovations in skincare, cleaning, bio-degradable packaging, and increasing stress tolerance in agriculture, among others. Understanding metalloenzymes at a fundamental level would allow us to modify their structure for new functionalities or identify optimal targets. This capability could revolutionize multiple industries, from pharmaceuticals to sustainable chemical manufacturing, enabling novel biomaterials and consumer products with enhanced performance and reduced environmental impact. Traditional experimental approaches to studying and engineering these complex systems are prohibitively time-consuming and expensive, often requiring years of laboratory work and millions in research funding. The question is: how can we achieve this level of molecular engineering more efficiently?

Enhanced natural engineering for industrial applications

Traditional approaches to harnessing nature's capabilities are limited by computational constraints and molecular complexity. By integrating the predictive power of AI with the high accuracy of quantum computing, we are developing a hybrid classical-quantum computing approach to overcome these barriers and engineer nature's processes with unprecedented precision. In our AI-Driven Protein Engineering and Quantum Simulation (AIPEQS) project, we combine classical AI with quantum computing to optimize metalloenzyme design, accurately predicting and optimizing enzyme activity using natural principles and physics to validate structures. This integration addresses the fundamental challenge of accurately predicting and optimizing enzyme activity while maintaining computational efficiency. Our work builds upon research showcased in the **quantum computing for drug discovery challenge** and our own work on **quantum-centric data-driven approaches for reactivity predictions for drug discovery**, which demonstrated quantum computing's practical application in molecular property estimation.

The main components of our project consisted of two workflows:

1. Classical-compute based AI pipeline for protein structure prediction
2. Quantum-centric data-driven property prediction of the catalytic reaction within the protein.

1. Classical-compute based AI pipeline for protein structure prediction

Our AI pipeline for drug discovery includes three key steps: protein structure prediction, metal-ion location prediction, and ion redocking, all performed on GPU-heavy computing resources. We utilized **Boltz-1** to predict protein structures with high accuracy, ensuring proper co-factor incorporation. With **Metal3D**, we predicted metal ion binding sites, which was crucial for correct metal coordination. Lastly, with sophisticated redocking methodologies, we simulated and validated dynamic interactions between metal ions and proteins, ensuring stable complexes.

2. The quantum computing enhancement

The quantum computing enhancement aims to precisely describe the active core of proteins, aiding in understanding their mechanisms and functionality. Analyzing, understanding, and predicting the unique catalytic properties of metalloenzymes allows for their targeted use in developing effective and innovative products. However, modeling metalloenzymes is challenging due to the quantum mechanical nature of metal-ligand interactions. Traditional methods struggle with these quantum interactions and protein dynamics. Our approach combines classical and quantum computing to simulate these systems effectively. We developed software for embedding, Hamiltonian simulation, and data-driven modeling, further enhanced with QM/MM methods to incorporate into artificial enzyme design. This integration enables accurate predictions and optimizations for efficient protein engineering.

How quantum and AI model oxidative stress

To illustrate our hybrid AI-quantum approach, we examined superoxide dismutase (SOD), a crucial metalloenzyme in cellular defense. SOD mitigates oxidative stress by converting superoxide radicals into oxygen and hydrogen peroxide, protecting proteins, lipids, and DNA from damage. This protective function is vital because unchecked free radicals can damage proteins, lipids, and DNA, ultimately compromising cellular integrity and function. The significance of SOD extends beyond basic cellular biology to applications, for example in dermatological formulations. Skin tissue is particularly vulnerable to oxidative stress due to constant exposure to environmental factors such as UV radiation and pollutants. These stressors generate reactive oxygen species that accelerate aging processes and contribute to various skin pathologies.

As mentioned above, the modeling of metalloproteins like SOD is challenging due to its complex metal-centered active site and quantum mechanical catalytic mechanism, and that makes it a perfect candidate for this study.

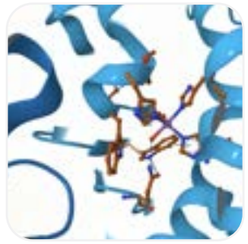

Benefits of a hybrid approach

Our hybrid AI-quantum approach offers significant advantages over traditional methods. By modeling molecular interactions at the quantum level, we achieve unprecedented accuracy in understanding complex electronic interactions during metalloenzyme catalysis. This will reduce the number of experiments needed, shorten research timelines, and open new design spaces. Unlike purely data-driven methods, our approach integrates fundamental physical principles, leading to more effective therapeutic enzymes and industrial catalysts in areas where extensive data sets are lacking. Economically, it reduces R&D costs by replacing trial-and-error with precise computational modeling, democratizing advanced enzyme engineering.

Best practices

Our AIPEQS project delivers an efficient model on quantum devices for sophisticated QM/QM/MM embedding schemes; this maintains high fidelity at catalytic centers while still creating meaningful results. On the protein design approach, we matched AI models to specific protein families and achieved an improved prediction accuracy and efficiency. We recommend rigorous validation protocols that verify computational predictions through both structural and functional analyses, creating a feedback mechanism that continuously refines prediction accuracy. Additionally, strategic optimization of computational resources ensures maximum efficiency while maintaining necessary accuracy.

Future trends

We currently see the R&D landscape shifting from observation- to knowledge-driven methodologies, enhancing research success and reducing costs. The convergence of quantum computing, AI, and biotechnology is creating new innovation opportunities, necessitating a rethink of traditional R&D frameworks. This paradigm shift enables more focused research with higher success rates and lower costs. The integration of quantum computing, AI, and advanced biotechnology is driving unprecedented innovation and creating new methodological possibilities.

Protein engineering is at a critical inflection point, moving beyond trial-and-error toward a knowledge-driven paradigm enabled by AI and quantum convergence. Engineered proteins now sit at the heart of innovation across pharmaceuticals, sustainable materials, consumer products, and even energy storage. Our hybrid approach makes this transformation more systematic, efficient, and accessible. As computational methods advance, we expect faster discoveries, lower R&D costs, and a new wave of bio-inspired breakthroughs. The future of bioengineering won't just borrow from nature — it will help us engineer it with purpose. And that future is already in motion.

Moving with new technology, we also move into atomic-scale modelling

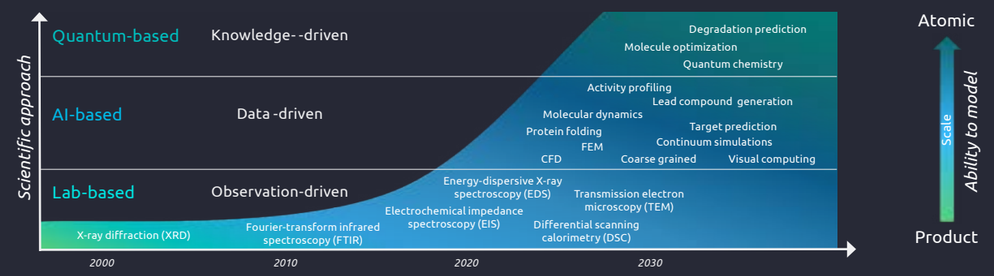

By modeling molecular interactions at the quantum level, we achieve unprecedented accuracy in understanding complex electronic interactions during metalloenzyme catalysis. This will reduce the number of experiments needed, shorten research timelines, and open new design spaces.

Start innovating now

Rethink R&D processes: From data-driven to knowledge-driven

Transition from traditional data-driven methodologies to a knowledge-driven approach. This shift focuses on understanding problems at their roots and leveraging advanced technologies to transform research and development processes. By adopting a knowledge-driven paradigm, you can enhance research success and reduce costs, moving beyond trial-and-error to precise, informed decision-making.

Leverage advanced technologies

Utilize cutting-edge technologies like AI and quantum computing to accelerate research and development. These tools enable accurate digital modeling and simulation, providing deeper insights into molecular interactions and optimizing research approaches. Implementing these technologies will democratize advanced enzyme engineering, allowing smaller groups to contribute alongside industry leaders.

Innovate with hybrid approaches

Combine the potential of current technologies with machine learning and classical computing algorithms in a hybrid way with new technologies like quantum computing and generative AI. This method allows development of adaptable approaches that can evolve as new technologies mature and improve, ensuring continuous innovation and efficiency.

#KnowledgeDrivenR&D, #Innovation, #AI, #QuantumComputing, #QuantumChemistry, #ComputationalChemistry, #MultiscaleModeling, #GenerativeAI, #HybridApproach, #AdvancedTechnologies, #EnzymeEngineering, # ProteinDesign, #MetalloProteins, #SustainableProducts, #ResearchAndDevelopment, #TechInnovation, #FutureOfR&D

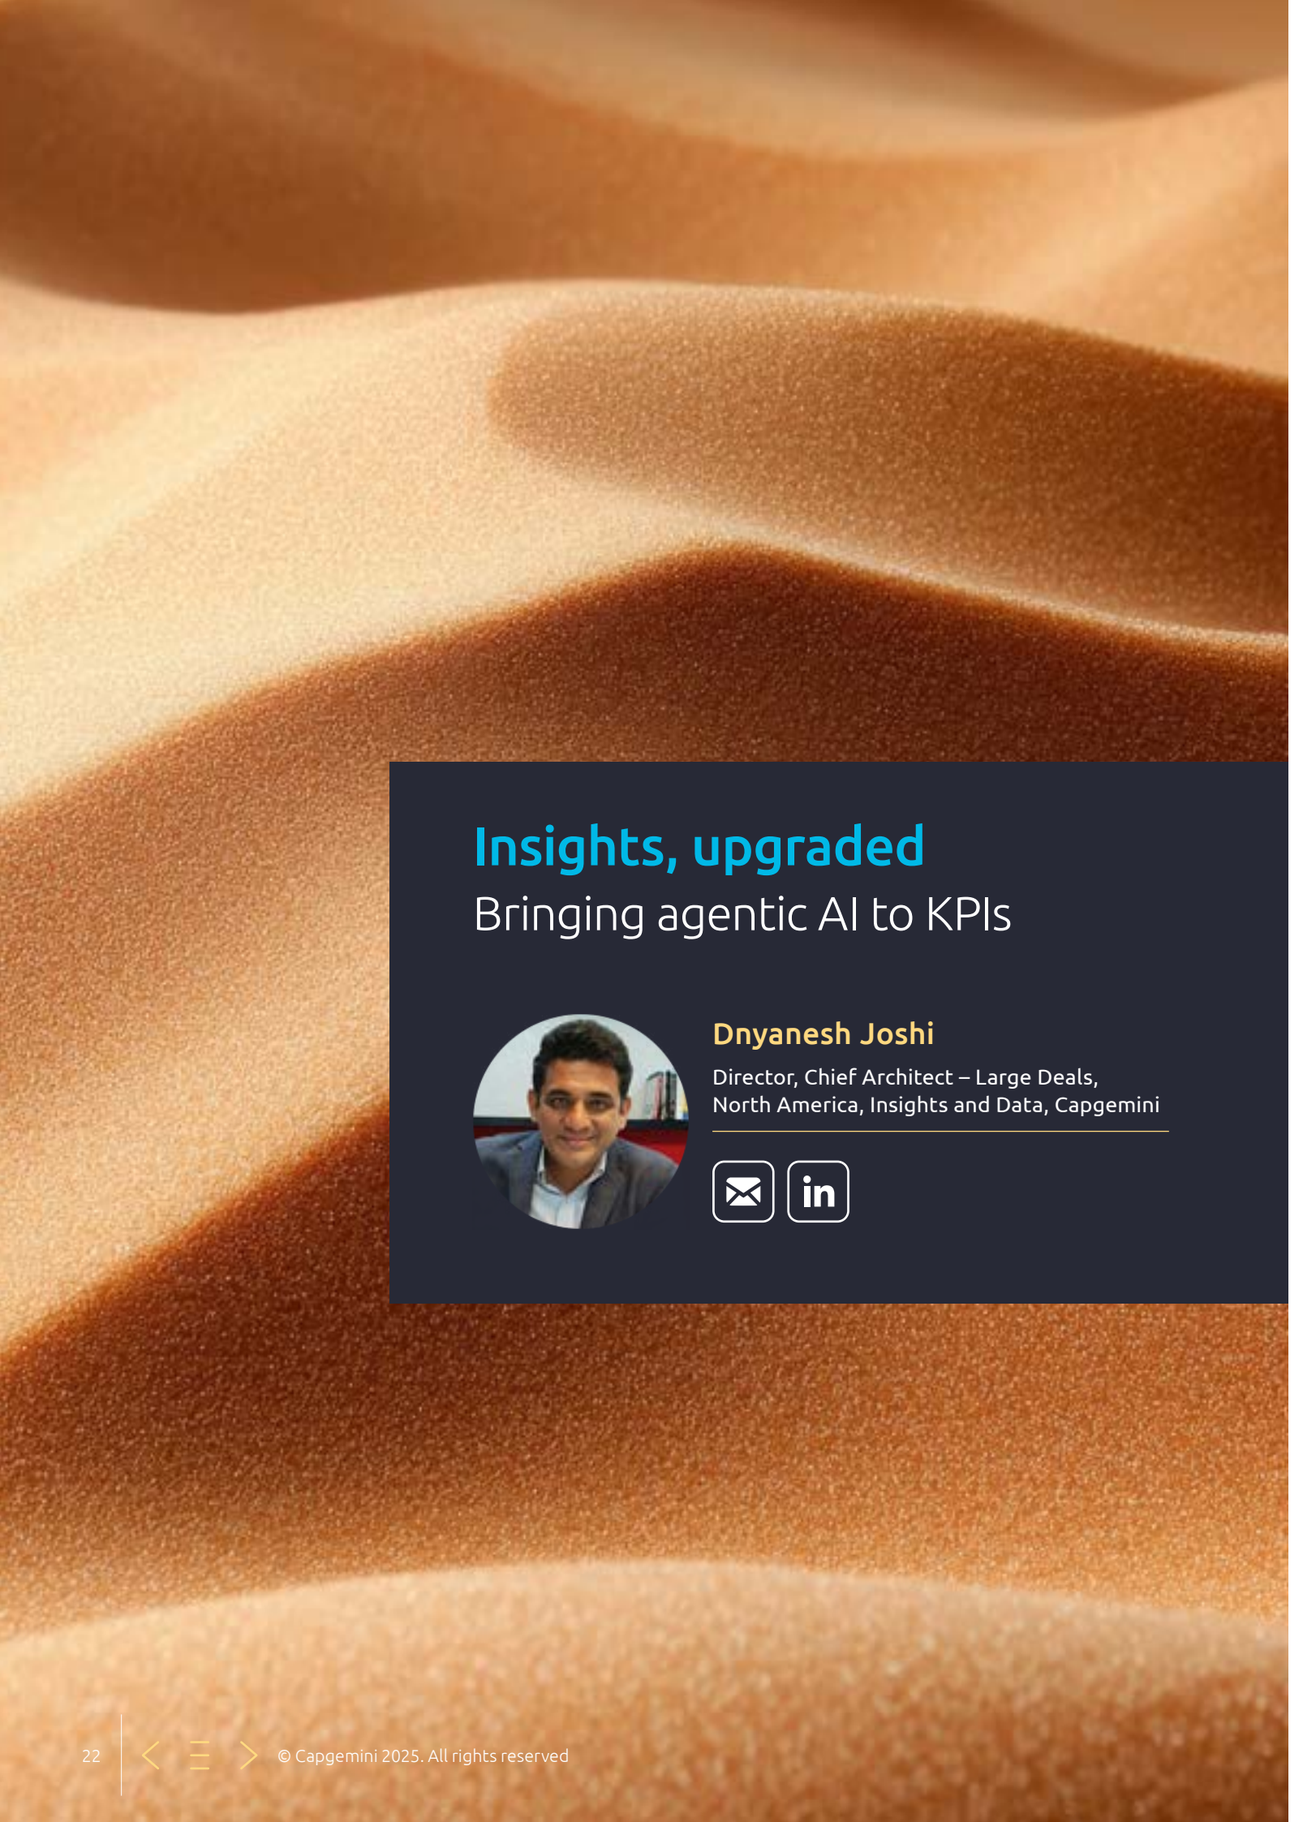

Insights, upgraded

Bringing agentic AI to KPIs

Dnyanesh Joshi

Director, Chief Architect – Large Deals, North America, Insights and Data, Capgemini

Dashboards don't make decisions. People do. But what if your business insights were already aligned with your KPIs, personalized to your role, and ready to act on today, with no perfect data platform required? Our Gen AI Strategic Intelligence System by Capgemini (GSIS) brings agentic AI into the boardroom, enabling organizations to boost performance using the data they already have. From finance and supply chain to IT and revenue growth, GSIS delivers explainable, persona-aware insights that connect strategy to action – and help enterprises move from legacy reporting to next-generation intelligence.

It's clear executives understand Gen AI's potential to deliver tangible benefits. In **Harnessing the value of generative AI: 2nd edition**, the Capgemini Research Institute noted 80 percent of organizations increased their investment in Gen AI in 2024, compared to the previous year. What's more, led by large enterprises, 24 percent of organizations surveyed have already integrated Gen AI into some or most of their locations or functions. Researchers also learned 82 percent of organizations plan to integrate AI agents into their operations within the next one to three years, as these evolve from supportive tools into independent entities that can operate autonomously.

The disconnect between data and insights

But organizations trying to apply AI to extract business insights using legacy BI systems face many major challenges.

- Analytics systems often fail to support strategic foresight and transformative innovation – instead providing business users with yet another dashboard.
- The results are often, at best, a topic for discussion at the next team meeting – not sufficient for a decision-maker to act upon immediately and with confidence.
- Systems typically fail to personalize their output to provide insights contextualized for the person viewing them – instead offering a one-size-fits-nobody result.
- Systems often aggregate data within silos, which means their output still requires additional interpretation to be valuable.

In short, many legacy systems miss the big picture, miss actionable meaning, miss the persona – and miss the point.

An advisor that works with today's data

Those shortcomings drove the development of the Gen AI Strategic Intelligence System by Capgemini (GSIS). This agentic AI advisor leverages several new technologies to maximize the power of Gen AI to access and analyze an organization's data. These include Anthropic PBC's Model Context Protocol (MCP) to effortlessly connect AI agents to any data source, and Google's Agent2Agent protocol (A2A) to enable multiple AI agents to seamlessly collaborate.

The result is the equivalent of engaging hundreds of domain experts to tackle each business unit's most pressing challenges.

Capgemini's solution uses virtualization and governance to capture the organization's information in a data fabric. This alleviates the need for the enterprise to make the significant investment required to build a unified data platform. Agentic AI agents are then able to analyze this data, collaborate with other agents throughout the organization, and make specific recommendations to boost the domain's core KPIs.

The solution's recommendations are based on a complete view of the enterprise – something impossible for human experts to deliver. What's more, the solution employs explainable AI and a chain of thought reasoning. This enables human users to understand how the AI arrived at its conclusions, so they trust the solution's results.

Applicable to multiple domains

The Gen AI Strategic Intelligence System by Capgemini (GSIS) offers hyper-personalization at scale across various domains. As an example, an enterprise's finance domain can deploy it to fundamentally change the DNA of performance reporting by:

- Delivering a unified view of the truth by embracing both ERP and non-ERP data sources with no data re-engineering required
- Providing independent performance tracking and recommendations for next best actions via natural language processing of queries from the CFO, Chief Risk Officer, or financial controller
- Enabling persona-based dynamic KPI reporting using Gen AI and conventional AI analytics to optimize relevance, depending on whether the solution is working for a CFO, FC, or RO
- Leveraging six key AI services for improved, data-driven decisions – covering revenue assurance, financial performance, financial close, cost reduction, working capital, and risk and compliance.

Key benefits for the finance domain include improvements to cash flow, working capital, asset utilization, operational efficiency, and cycle times. The solution can also help reduce material and operational costs as a percentage of revenue, inventory costs, reporting times, and budget versus actual variances. And it can decrease the manual effort involved in financial processes by 20 to 30 percent.

Gen AI Strategic Intelligence System by Capgemini

Fundamentally change the DNA of performance reporting

- Independent performance tracking and next-best actions/improvements via NLP-based querying by CFO/FC/CRO
- Finance persona-based dynamic KPI reporting via pre-canned models using Gen-AI/ conventional AI analytics
- No data re-engineering approach based on a unified view of truth of ERP and non-ERP data sources in a data fabric and in data products
- Finance knowledge-based delivery through six key AI services: revenue assurance AI, fin performance AI, fin close AI, cost reduction AI, working capital AI , risk and compliance AI

Outcomes at scale

As CFO I want to...

Reduce CapEx, OpEx while ensuring **revenue leakage prevention**

As FC I want to...

Ensure effectiveness of internal controls and accuracy of **financial statements**

As RO I want to...

Ensure 100% compliance with data protection regulations, improve risk identification and mitigation rates, and eliminate fraud incidents

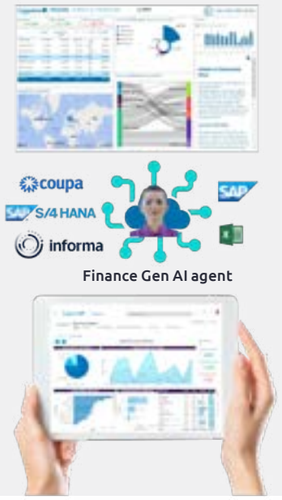

Finance Gen AI agent

The Gen AI Strategic Intelligence System by Capgemini (GSIS) delivers comparable benefits to other domains, too.

Supply chain: Insights on inventory, supply planning, logistics, spend, and orders. Potential benefits include more accurate inventory, plus improved procurement, logistics, transportation, and order management.

Revenue growth: Insights on interactions, intelligent customer profiling/segmentation, sentiment analysis, and marketing. Expected benefits include sales growth, higher profit margins, more sales opportunities, lower acquisition cost, higher conversion rate, and improved ROI. It can also help formulate targeted revenue growth management campaigns.

IT SDLC/operations: Insights on automation and digital employees. Potential improvements include higher productivity within SDLC and IT Ops.

Innovation enables long-term investment

To unlock the full value of generative AI at scale, organizations must still pursue best-in-class data strategies. These include having the right platform in place to build secure, reliable, and scalable solutions, adopting an enterprise-wide governance framework, and implementing all appropriate guardrails.

That said, the Gen AI Strategic Intelligence System by Capgemini (GSIS) delivers goal-oriented insights aligned with business objectives. It's ready to empower the organization, today, through actionable roadmaps for sustainable growth and competitive advantage. This helps executives realize business value almost immediately, including additional revenue that can then be invested in larger, longer-term Gen AI initiatives.

> *Many legacy systems miss the big picture, miss actionable meaning, miss the persona – and miss the point.*

Start innovating now

Don't wait

It takes time and resources to transform an organization's data landscape. Embrace analytics solutions that work with the existing landscape, even as the enterprise plans for tomorrow.

Focus agentic AI on KPIs

Even a small improvement to a core KPI can provide enormous benefits, so this is a great starting point for every domain within an organization.

Reinvest in data

Gen AI solutions can generate new revenue. Smart organizations will reinvest this in best-in-class data strategies to support larger, longer-term Gen AI plans.

#DataPowered #StrategicPerformanceManagement #DigitalAgentsInStrategicDecision-making #GenAI #DigitalEmployees #SCMDigitalAgents #FinanceDigitalAgents #RevenueDigitalAgents

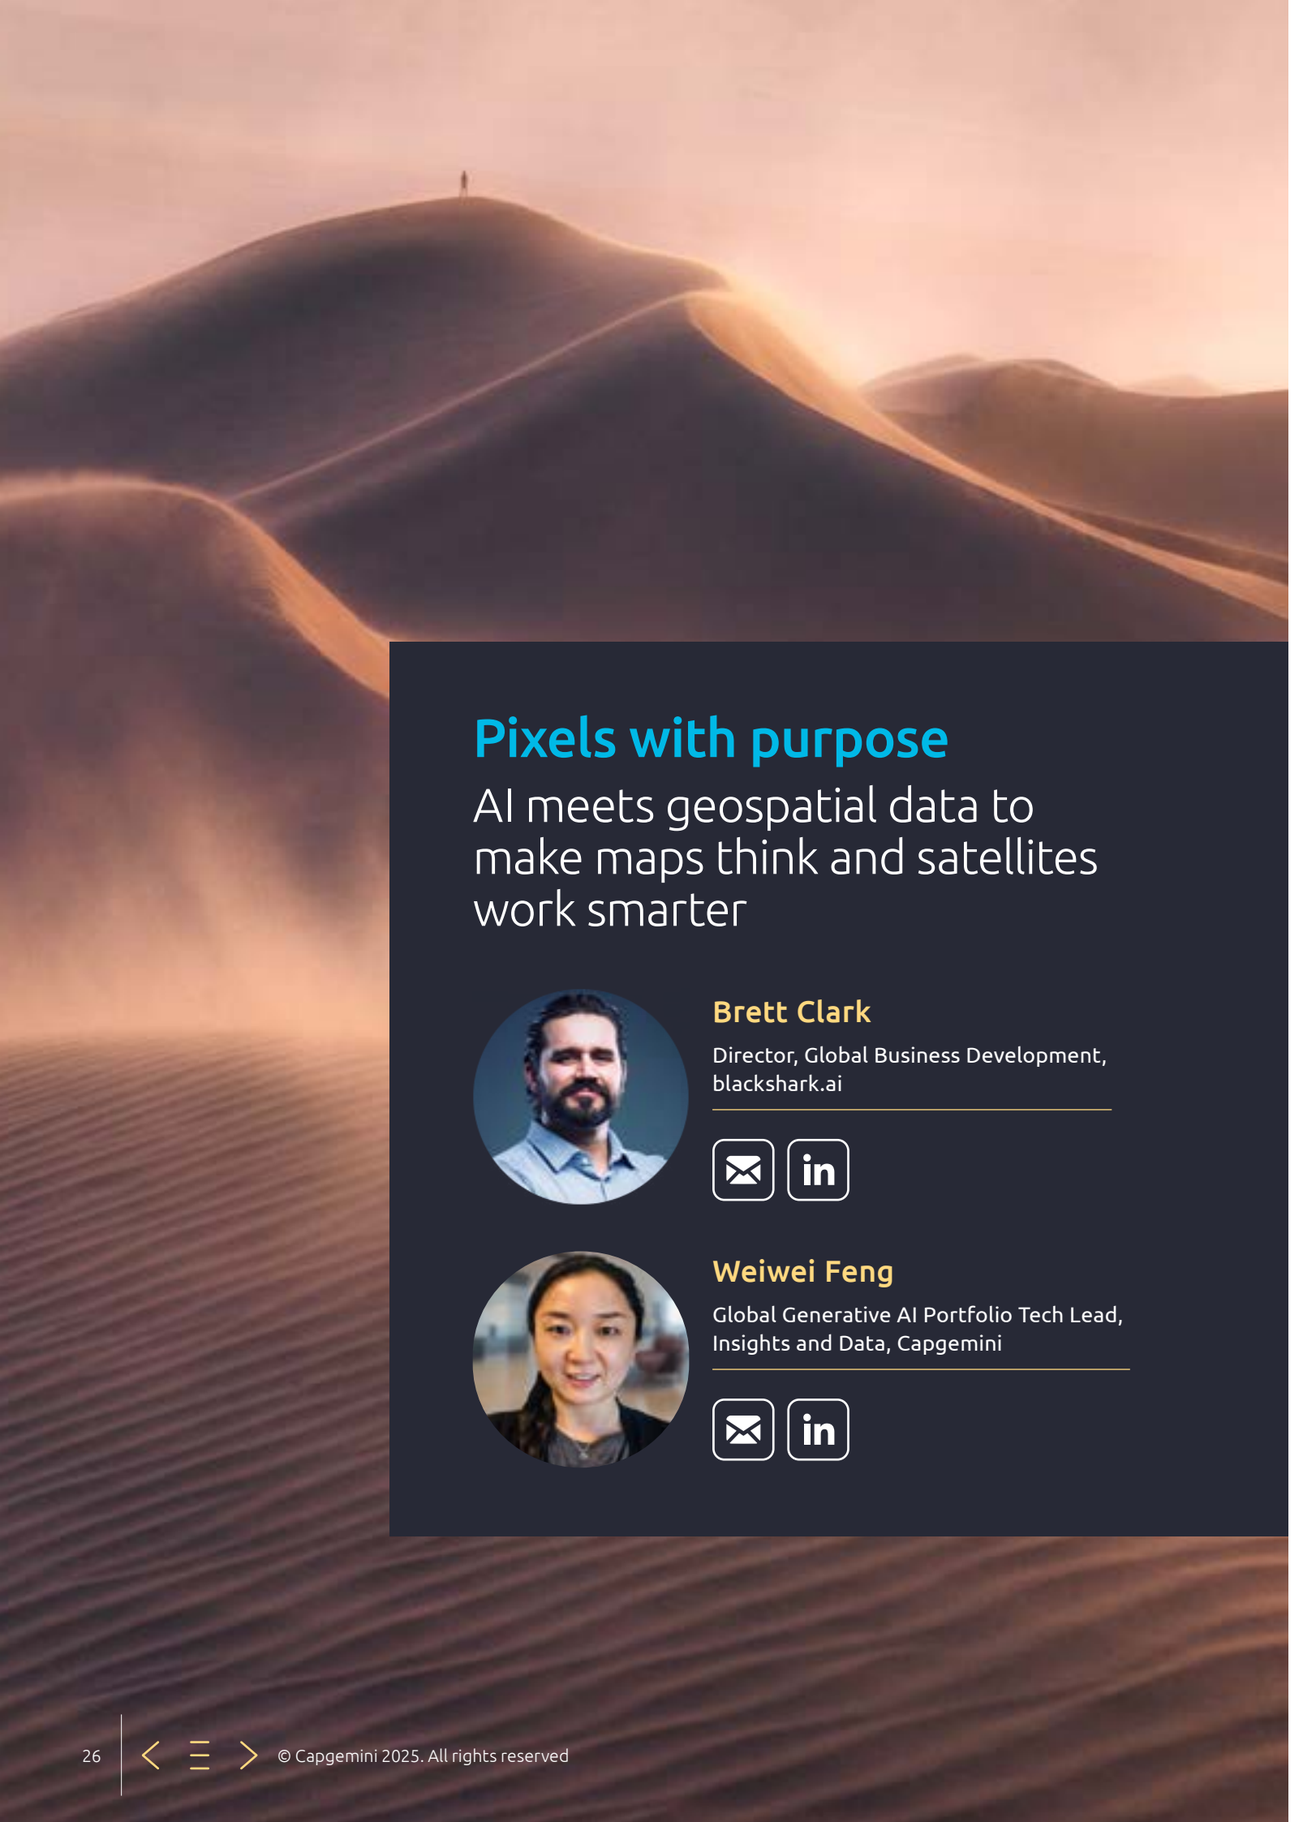

Pixels with purpose

AI meets geospatial data to make maps think and satellites work smarter

Brett Clark

Director, Global Business Development, blackshark.ai

Weiwei Feng

Global Generative AI Portfolio Tech Lead, Insights and Data, Capgemini

Geospatial data used to be passive – maps, images, layers. Today, it thinks. Thanks to AI-powered platforms like Blackshark.ai, geospatial intelligence is shifting from static pixels to smart, purpose-driven systems. From creating a digital twin of the entire Earth to enabling real-time detection of threats, trends, and opportunities, blackshark.ai is redefining what's possible with GIS. AI is breathing new life into satellite imagery, sensor data, and any other bitmaps – and every industry should start paying closer attention to what their pixels are really saying.

Organizations in most economic sectors can benefit from geospatial analytics, but there are a few standout examples.

Transportation and logistics companies use GIS to optimize route planning, improve fleet management, and better oversee supply chains. Increasingly, companies are looking to integrate GIS with other data sources to make more timely decisions and respond faster to unforeseen events.

GIS improves infrastructure management by helping public- and private-sector organizations design and construct more resilient buildings and other physical assets. It also monitors existing assets to identify issues so they can be addressed before major failures occur. GIS can also enhance how remote inspections are performed, improving worker safety and efficiency.

Other sectors that benefit from GIS include defense and intelligence, automotive (especially as autonomous vehicles are developed and introduced), retail, agriculture, government, insurance, real estate/property management, and environmental stewardship.

New realities present new challenges

There are, however, several barriers to unlocking the value of geospatial data.

Widespread deployment of satellites plus the rising popularity of drones and other technologies that capture images of the world around us mean organizations have plenty of visual data. But they don't really know how to make sense of it. The ability to extract useful intelligence from pixels has traditionally been a challenge.

While satellite imagery is public, many organizations have their own sources of visual data, such as aerial photos or drone images. Realizing the full potential of geospatial analysis requires a solution that can integrate open and commercial third-party datasets plus internal business data. Doing so effectively requires a solution that can accommodate a variety of data formats, diverse standards, and inconsistent models.

As the volume of data grows, many traditional, desktop-based tools used for geospatial workflows are no longer sufficient. But the new solutions that replace these must be efficient in terms of computational requirements, storage, and processing capabilities.

Focused solutions are better

It's important to avoid the trap of trying to create a single solution to perform every GIS-related task. Using only relevant data speeds up deployment time and makes it easier to adjust the solution to address evolving requirements. As an example, at blackshark.ai, our AI-powered solutions can be trained to ignore certain types of visual information, so they don't waste time and computing resources processing irrelevant data within an image.

Organizations must also be aware that capturing and using geospatial data must be done in compliance with all applicable personal information and data privacy laws. GIS solutions must remain on top of current and evolving regulations across multiple jurisdictions.

GIS on a global scale

Fortunately, new solutions that enhance traditional GIS applications with AI and other advanced analytics tools are addressing the issues. An example is **blackshark.ai**'s solution, which combines an AI-enabled geospatial platform with an automated 3D engine generator. blackshark.ai's platform extracts buildings, roads, vegetation, and other features from GIS data. It then applies AI and other analytics tools to create a digital twin of the Earth.

This is a realistic, 3D replica of the entire planet that includes detailed terrain and 1.4 billion structures. Significant buildings are rendered accurately, while others are what we call geo-typical: accurately sized and shaped but synthetically clad in regional-specific architecture, created from a library of thousands of 3D objects that form a procedural, generative building grammar. This generates realistic images while reducing data processing and storage requirements. The digital twin can be completely refreshed in just over two days – an exceptionally fast turnaround given the scale of the task – and supports large-scale simulation, training, and visualization applications.

We originally developed digital twins to support Microsoft Flight Simulator 2020, but its planet-blanketing scale immediately attracted interest from the defense and intelligence community. This provided blackshark.ai with a solid foundation for growth. For example, we've now enhanced the platform so it can ingest visual data from a wide variety of sources. These include camera phones, drones, Google Street View, video stills, thermal and infrared imaging, and even historical records such as maps and vintage photographs.

Applying intelligence to GIS

To extract actionable intelligence from this data, blackshark.ai has developed ORCA HUNTR. This no-code solution enables subject matter experts to train, test, and exploit the platform's object detection and classification abilities directly, with no need for a data scientist between the business user and the solution.

ORCA HUNTR works on any rasterized pixel information. This can include information that isn't normally considered "visual." For example, we have converted radio frequencies into waterfall plots, then used the resulting visual data to train ORCA HUNTR to identify active attempts at RF jamming.

A compelling competitive advantage

GIS delivers significant value for organizations. It supports data-powered decision-making to help identify opportunities for innovation, optimize operations, and improve risk management. Enhancing GIS with AI and other advanced analytics technologies makes it even more essential for enterprises to embrace. When pixels become intelligent, maps become mission-critical – capable of predicting disruptions, revealing hidden patterns, and simulating entire environments in real time. That's not just evolution; it's transformation. And it's what turns geospatial data from background noise into a competitive signal.

When pixels start thinking, maps stop being maps – and start becoming intelligence.

Start innovating now

Plan for GIS

Every industrial sector can benefit from geospatial data. If you're not using it yet, it's time to figure out what it can do for your organization.

Data isn't the problem

Most organizations generate plenty of geospatial data. What's needed is useful ways to interpret it.

Ease of use is critical

Embrace a no-code solution that democratizes GIS data so subject matter experts can leverage it directly.

#GeospatialIntelligence #DigitalTwins #AIForGeospatial #RemoteSensing #BlacksharkAI

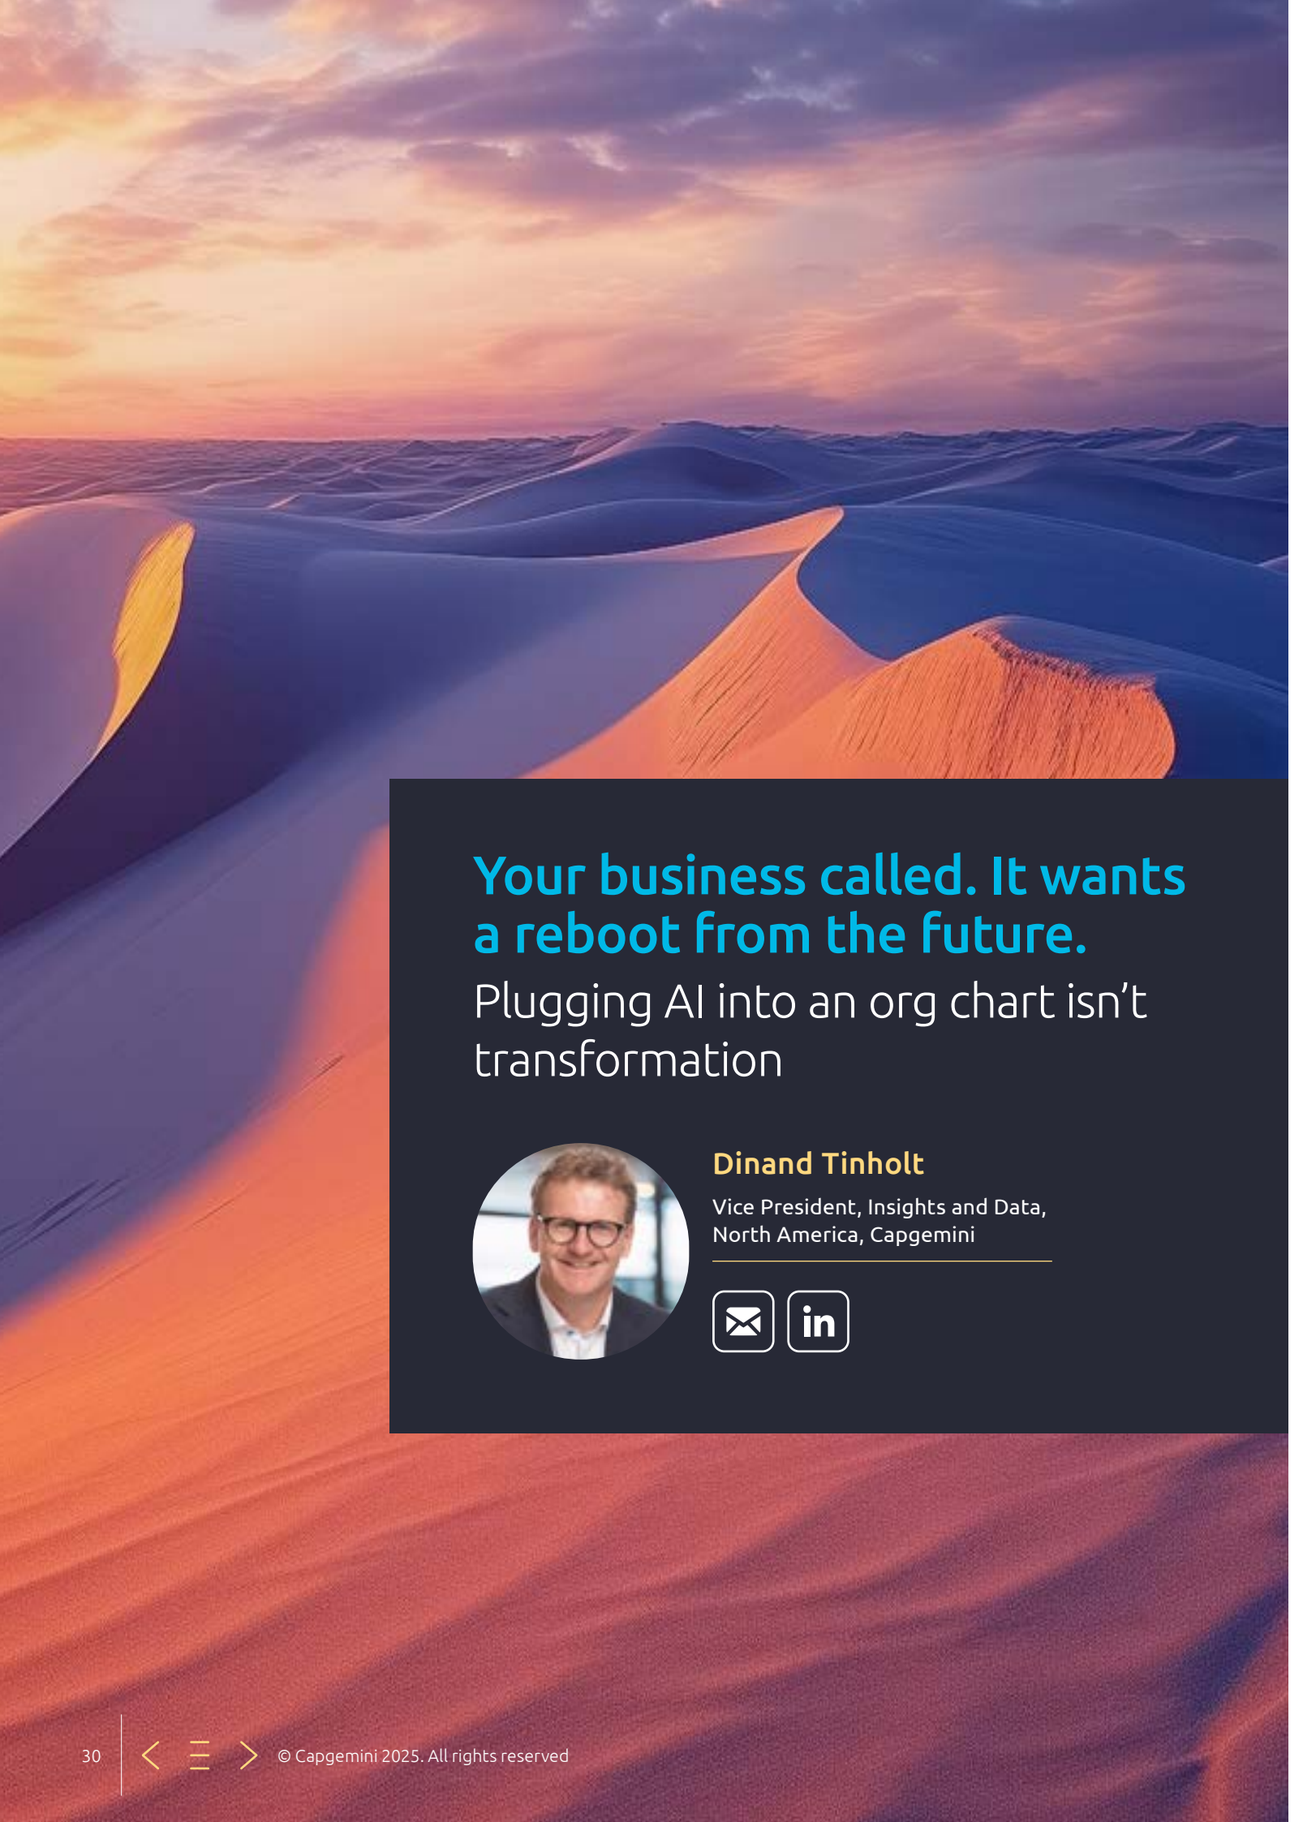

Your business called. It wants a reboot from the future.

Plugging AI into an org chart isn't transformation

Dinand Tinholt

Vice President, Insights and Data, North America, Capgemini

Most companies tweak. The bold ones transform. But a new mindset is emerging – one that goes full hacker-mode on the enterprise itself. No legacy constraints, no silos, no sacred cows. What if the business were rebuilt from scratch, as if it were being launched five or 10 years from now, with agentic AI not as a plug-in but as a strategic co-founder? This isn't about optimization – it's about reimagining the purpose, structure, and intelligence of the organization from the ground up. Agentic AI can help redesign the operating model, shift decision-making, and unlock entirely new ways of creating value – all by thinking from the future, not just about it.

Somewhere between your last strategic offsite and your current AI working group, the business world changed. Again. Only this time, it wasn't a new framework or another agile manifesto. This time, it's a new brain.

Agentic AI doesn't just automate what you already do – it questions why you do it in the first place. It proposes. It critiques. It simulates. It rewires how decisions are made, who makes them, and when they happen. And unlike your last re-org, it doesn't wait six months to show any impact.

The problem is that most organizations are still trying to install AI into 20th-century workflows like it's a plugin. "Let's just optimize procurement," becomes a board-approved project that takes 18 months, involves many large teams, and results in slightly faster procurement.

But what if you stopped optimizing? What if you hacked your own business?

Design from 2030, not 2020

Let's imagine that you're not burdened by legacy systems, organizational charts, compliance rituals, or "we've always done it this way" syndrome. Imagine you're starting from scratch. The year is 2030. Agentic AI is not just embedded in your processes – it's embedded in your people. Or rather, your people are embedded in a new kind of system: one where intelligence is ambient, decisions are real-time, and the line between human and machine is not blurred, it's collaborative.

Would you really recreate the same silos? Would you copy-paste last year's operating model into the future and call it transformation?

In this future-forward model, decisions don't just happen faster – they happen better. Picture an AI agent that pulls real-time data from your supply chain, simulates three economic scenarios, aligns them with your top KPIs, and taps your planning team with a message like: "If we shift production to Mexico for the next three weeks, we protect margins and avoid inventory bottlenecks. Proceed?" It's not fantasy. It's already emerging in pilot programs. Retail, logistics, finance – they're not waiting for permission. They're experimenting. Quietly, sometimes clumsily, but they're moving.

What's missing is boldness at scale. The audacity to say "If AI can hack my business, so can I."

Rebuilding your business from a blank slate sounds like the stuff of retreats and vision decks. But it's more than that. It's a challenge to strip things down to first principles. What decisions truly matter? Who needs to make them? What should be instantaneous, and what still needs reflection? Where do humans shine, and where do they hold things up?

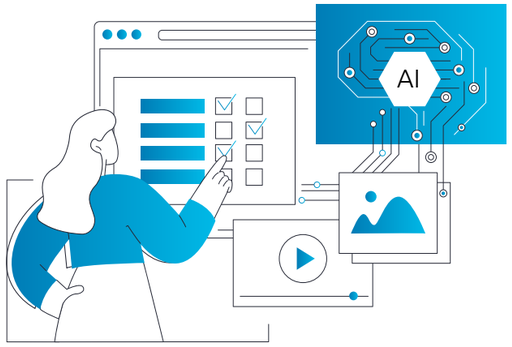

From workflows to intelligence networks

The old mindset builds workflows. The new mindset builds networks. You don't need another dashboard. You need a decision system. One that's alive, adaptive, and, yes, sometimes smarter than you.

When you start thinking this way, the org chart begins to look suspiciously like a museum exhibit, a relic from a time when communication flowed one way and information took the scenic route. In an agentic enterprise, power is less about title and more about your ability to navigate systems, interpret signals, and co-create with machines. Managers stop managing tasks and start designing interactions. Strategy is no longer a quarterly slide deck, it's a living process that evolves in real time, shaped by continuous inputs and autonomous agents that never sleep.

And leadership? Leadership becomes less about having the answers, and more about asking the right questions. Less about command and control, more about trust and iteration.

Yes, it sounds radical. So did the cloud. So did putting your ERP in someone else's data center. So did letting your intern post on LinkedIn. But here we are. Still talking about "pilots" and "use cases," while the real opportunity is to reimagine the business entirely.

> *The old mindset builds workflows. The new mindset builds networks.*

AI-native government

An example? In government, an AI-native model wouldn't just digitize existing services – it would reimagine how policy is shaped and delivered. Agentic AI can simulate the real-time impact of legislation across demographics, recommend adjustments to improve equity or efficiency, and even co-design citizen services based on behavioral signals, not just historical data. It's not bureaucracy with better bandwidth. It's governance with built-in intelligence.

A less embarrassing future

So, if yet another AI pilot is on the table – same org, same roles, same handoffs – it might be time to pause. Because layering new tech on old thinking rarely leads to transformation. Instead, ask the only question that truly matters: What would we build, if we started today, from five years in the future? Not to optimize what's already there, but to rethink what should be there. In that future-first mindset, agentic AI becomes more than a tool – it's a co-architect of the business itself. And the goal isn't to make the past slightly more efficient. It's to make the future slightly less embarrassing.

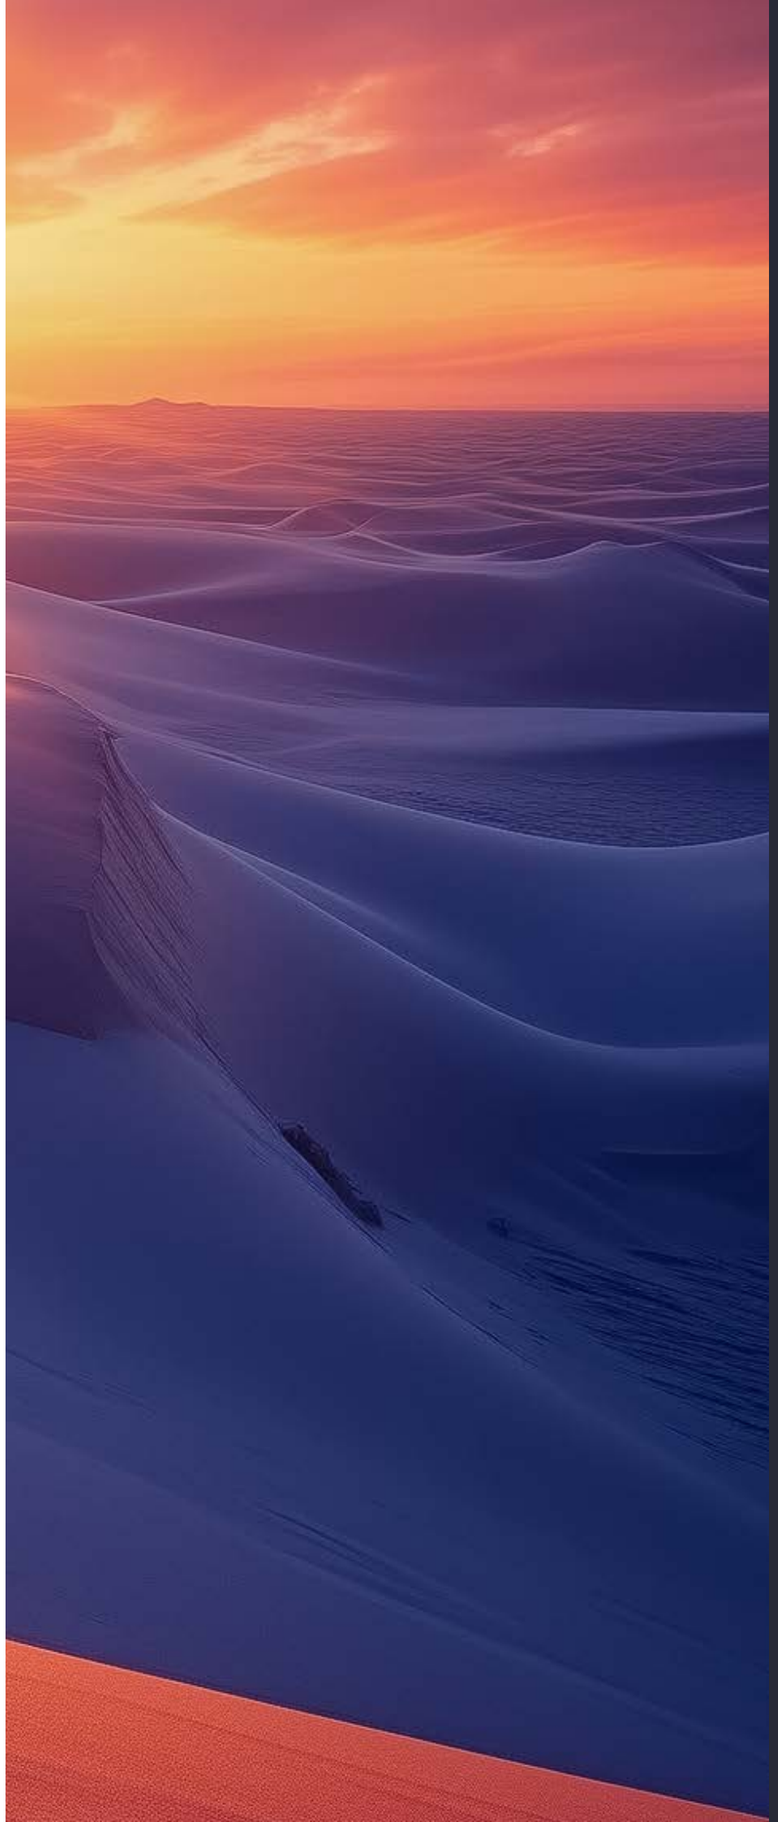

Start innovating now

Audit your decisions

Choose a function, like finance, supply chain, or marketing, and map the 10 most frequent decisions made each week. Then ask: which of these could be delegated to an AI agent, and which require uniquely human judgment? The results will surprise you.

Pilot a human plus agent workflow

Redesign a single, low-risk business process with an AI agent embedded in the loop. Don't aim for full automation, just real-time collaboration. Try demand forecasting, contract review, or pricing adjustments. Measure speed, quality, and trust.

Build your "AI from scratch" blueprint

Assign a cross-functional team to answer: "If we rebuilt this business in 2030, with no legacy and full agentic AI capability, how would we run operations, make decisions, and structure teams?" Document it. Then look for one idea you can implement now.

#AgenticAI
#BusinessReinvention
#AIDecisionMaking
#FutureOfWork
#DataPowered
#CapgeminiInnovation

Gen Garage

Where tomorrow's talent builds today's AI for good

Aishwarya Kulkrni

Program Manager – Gen Garage, Insights and Data, Capgemini

Gen Garage is redefining innovation by harnessing AI to build real-world solutions in areas such as disaster management, sustainable farming, and environmental risk mitigation. By fostering talent transformation and embracing cutting-edge technologies, we empower tomorrow's professionals to shape the future through hands-on impact. Step into Gen Garage, Capgemini's Insights & Data innovation hub — where visionary minds, guided by expert mentors, craft transformative AI-driven solutions at the intersection of talent and technology. Fueled by generative AI, machine learning, and automation, Gen Garage accelerates operational excellence and delivers innovations that drive efficiency, inclusivity, and sustainability. As businesses lean into data-powered insights, our solutions stay aligned with evolving needs — helping organizations stay future-ready, while making a difference today.

Gen Garage is innovating in different areas across a number of projects. Here are a few.

DisasterX: Faster and smarter response for disaster management, powered by AI

DisasterX is a cutting-edge AI-powered application designed to revolutionize disaster response by enabling real-time, adaptive decision-making in high-pressure environments. Leveraging agentic AI, DisasterX overcomes the challenge of delayed responses by autonomously analysing vast amounts of real-time data from multiple sources. This allows for optimized rescue efforts, efficient resource allocation, and enhanced recovery strategies. By continuously learning from past disaster scenarios, DisasterX improves prediction accuracy, minimizes human error, and accelerates response times – ultimately saving lives and reducing economic and environmental impact. Its autonomous capabilities ensure scalability and resilience, seamlessly adapting to both local and large-scale emergencies. As the rise of agentic AI reshapes automation, DisasterX stands at the forefront of intelligent, proactive disaster management, delivering greater efficiency and reliability in crisis situations.

KisanGPT: AI-driven insights for smarter, sustainable farming

KisanGPT is an AI-powered platform designed to revolutionize farming by providing real-time insights on crop health, weather forecasts, and sustainable agricultural practices. Using advanced language models, satellite data, and weather analytics, it offers personalized recommendations to help farmers optimize yields, conserve resources, and tackle climate challenges. The platform supports multilingual access and speech recognition, ensuring inclusivity for farmers across diverse regions. By integrating market trends, government policy updates, and best farming practices, KisanGPT enhances decision-making, boosts profitability, and promotes eco-friendly agriculture. This AI-driven solution not only improves efficiency but also fosters a more resilient and sustainable farming ecosystem.

The Green Horizon: AI-powered vegetation hazard management

This solution detects and manages vegetation hazards near power lines, preventing wildfires, outages, and safety risks. Using satellite imagery, machine learning, and weather forecasting, it provides real-time monitoring, predictive insights, and proactive risk mitigation. With an intuitive chatbot and geospatial analytics, it empowers organizations to optimize resources, reduce costs, and ensure safer, more sustainable infrastructure. By automating hazard detection and integrating user feedback, it enhances decision-making for utility companies and environmental agencies. This innovative approach not only improves operational efficiency but also supports long-term sustainability and infrastructure resilience.

Market trends, opportunities, and developments

Gen Garage strategically aligns its initiatives with prevailing market trends to address pressing societal and business needs.

The increasing investment in AI for disaster management presents a significant market opportunity for DisasterX to deliver innovative and data-driven solutions. With the rise of smart cities and the widespread adoption of IoT sensors in disaster-prone areas, vast amounts of real-time data can be leveraged for predictive analytics and rapid response. Gen Garage is at the forefront of this transformation, utilizing AI to enhance disaster preparedness and resilience. As climate change intensifies the frequency of natural disasters, the demand for intelligent, automated response systems continues to grow, positioning DisasterX as a key player in optimizing disaster mitigation and emergency management strategies.

KisanGPT taps into the growing demand for AI-driven agricultural solutions. By leveraging real-time analytics and precision farming techniques, Gen Garage maximizes market opportunities, helping farmers and agribusinesses adopt smarter, data-driven strategies. With advancements in AI and increasing support for sustainable farming practices, the platform positions itself as a game-changer in modern agriculture, driving innovation and long-term growth in the sector.

The Green Horizon initiative taps into the growing need for AI-driven environmental risk management. By integrating geospatial intelligence and predictive analytics, Gen Garage maximizes market opportunities, enabling utility companies and agencies to adopt smarter, data-driven strategies for sustainability and infrastructure resilience.

Gen Garage is where innovation gets hands-on – and where emerging talent learns by doing. By combining mentorship with real-world problem-solving, we're helping young professionals grow into AI changemakers while delivering solutions that matter. From climate-smart farming to disaster response, the Garage proves that AI for good isn't just a concept – it's a daily practice. The challenges may be big, but with the right mix of curiosity, code, and collaboration, we're building something that lasts. We'll have more cases to share in the next edition of the *Data-powered Innovation Review.*

> *Gen Garage accelerates operational excellence and delivers cutting-edge innovations that drive efficiency, inclusivity, and sustainability.*

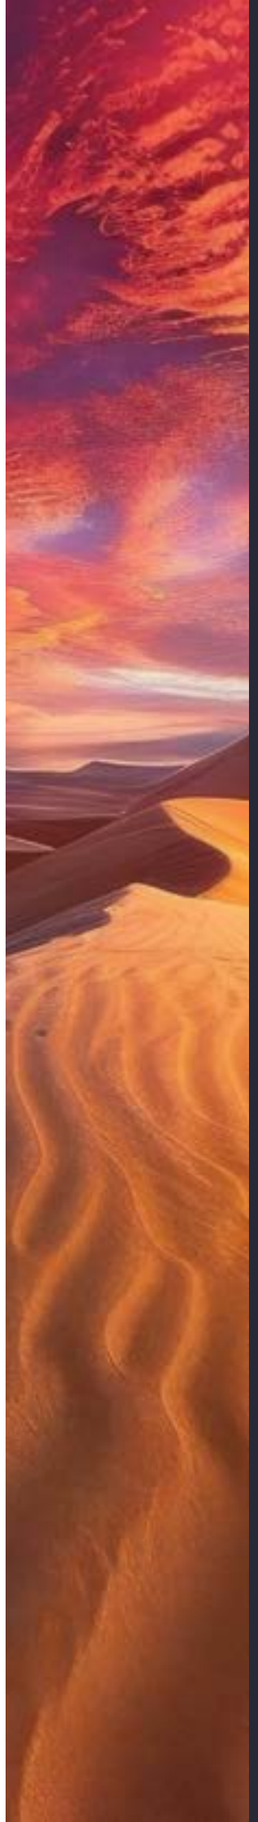

Start innovating now

Empower future talent

Get involved in innovation projects that enhance AI skills and leadership capabilities, preparing young professionals for real-world challenges.

Leverage AI for social impact

Adopt AI solutions to drive sustainability, inclusivity, and efficiency across industries, from disaster management to smart farming.

Stay ahead of market trends

Engage with cutting-edge technology and AI-driven insights to maintain a competitive lead in an evolving digital landscape.

#DataPowered #AI4Good
#DigitalInclusion #GenAI #InnovateForAll
#CuttingEdgeAI #NextGenNLP
#SmartMLSolutions #VoiceBots
#GenGarage #SustainabilityForGrowth
#TechForInclusion

Ø ETL

Less extract, more impact

Sjoukje Zaal

Chief Technology Officer, Insights and Data Europe, Capgemini

ETL used to be a puzzle only data engineers could solve – with enough time, code, and coffee. Now generative AI is quietly rewriting the rules. Traditional pipelines – slow, brittle, and endlessly maintained – are giving way to Ø ETL: adaptive, intelligent flows that go from prompt to pipeline in seconds. No scripts, no tickets, no heroic debugging sessions. Data integration is becoming faster, smarter, and far more democratic – turning engineers into orchestrators and putting AI to work where it actually makes sense.

A shift from manual to machine-driven

Extract, transform, and load (ETL) has always been about moving data from different sources, transforming it into a usable format, and loading it into a system where it can be analyzed. But with the explosion of data, increased regulatory pressure, and the move to hybrid and multi-cloud environments, this process has become much more complex.

Generative AI is changing the game. Instead of writing and maintaining endless scripts and workflows, organizations can now use AI models to automate ETL pipelines. These models understand the context of the data, learn from existing integration patterns, and generate optimized workflows on the fly.

This leads to significant benefits.

- **Speed:** AI can generate and update ETL logic in minutes, not days.
- **Consistency:** AI-driven pipelines are less prone to human error.
- **Adaptability:** They automatically adjust to schema changes or new data sources.

Beyond automation: intelligent integration

AI isn't just speeding things up – it's making data integration smarter. By applying natural language understanding, organizations can describe what they want in plain English, and the AI creates the integration pipeline.

For example, a demand planner can say: "Extract product inventory from Oracle, combine it with daily sales from Shopify, calculate stock turnover per SKU, and load it into Snowflake for reporting." With traditional ETL, manual SQL logic, batch jobs, and schema mapping must be created. AI will generate the pipeline on demand from the prompt.

This approach democratizes data integration. It removes the dependency on specialized engineers for every change and helps more people in the organization work with data directly.

Aspect	Traditional ETL	Ø ETL
Process flow	Extract → transform → load	Direct flow with real-time transformation
Data movement	Batch-based, moves data between systems	Minimal movement, transforms data in place or in real time
Development effort	Manual coding and configuration	Prompt-driven using generative AI
Responsiveness to change	Slow to adapt; changes require code updates	AI adapts to schema or logic changes automatically
Latency	High (due to batch processing)	Low (real-time or near real-time)
Scalability	Limited by infrastructure and manual processes	Scales dynamically with cloud-native infrastructure
Cost of maintenance	High due to constant monitoring and refactoring	Lower through automation and self-healing pipelines
Governance and lineage	Manual tracking and documentation	Built-in, auto-generated metadata and lineage

How this fits in a modern data strategy

Generative AI for ETL is a natural fit in environments where data fabrics or data mesh architectures are being implemented. Modern data strategies are shifting from centralized control to decentralized ownership. Concepts like data mesh and data fabric are driving this shift, giving teams more flexibility to manage, consume, and share data across systems. In these models, every domain owns its data products, but the organization still needs consistency, compliance, and efficiency at scale.

Supporting decentralization without losing control

In a data mesh, teams manage their own pipelines. Traditional ETL tools can't keep up with the constant change and complexity. AI-driven ETL supports this by giving each team a way to build and manage data flows independently – without starting from scratch or involving a central data engineering group every time.

Cross-cloud compatibility

Leading platforms are already moving in this direction.

- **Google Cloud:** With services like BigQuery Dataform and Cloud Data Fusion, Google supports declarative and visual data pipeline development. Generative AI models from Google's Vertex AI can integrate with these services to streamline data prep and transformation.
- **AWS:** Amazon's Glue Studio offers low-code/no-code pipeline development, and new AI integrations allow users to describe what they want in natural language. Combined with SageMaker and Bedrock, AWS is aiming to simplify the entire data lifecycle – from ingestion to modeling.
- **Microsoft Azure:** Azure Data Factory and Synapse Analytics are embedding AI directly into pipeline creation and monitoring. With Microsoft Copilot, users can ask for transformations, lineage, and integration logic using natural language.
- **Databricks:** With its Lakehouse architecture, Databricks is adding AI to simplify pipeline generation in notebooks and workflows. Unity Catalog, when paired with LLMs, supports context-aware data discovery and security enforcement.
- **Snowflake:** Its growing suite of AI features, including Snowpark and Cortex, allows SQL and Python users to automate parts of the data prep process. With Snowflake's native LLM support, the platform is well-positioned to offer AI-driven transformations at scale.
- **Open-source and hybrid platforms:** Tools like Apache Airflow, Dagster, and dbt are starting to explore AI plugins and extensions. These add automation and intelligence to open workflows, making it easier for developers to generate and maintain pipeline logic.

What's next?

We are moving toward a future where ETL as we know it may no longer exist. Instead, we'll see dynamic data integration powered by AI. The concept of "ETL pipelines" will be replaced by intelligent agents that continuously ingest, transform, and validate data in real time, guided by policies and context, not hardcoded rules.

This means that embedding generative AI into ETL processes across platforms gives organizations many advantages.

- **Time to value shortens:** Data products go live faster, helping teams act quickly.
- **Complexity reduces:** AI handles edge cases, schema drift, and exception handling in real time.
- **Data quality improves:** Built-in rules and real-time validation become part of the generated logic.
- **Business access increases:** More users across domains can work with data confidently, without needing to be engineers.

This isn't just a technological shift. It's a change in how we approach data, moving from pipelines built manually to systems that can learn, generate, and adapt automatically. Remember that AI isn't replacing data engineers – it's changing their role. The most successful organizations are those that help their teams adapt to becoming orchestrators and quality managers rather than code writers.

Maybe Ø ETL doesn't mean "no ETL" – but it definitely means no more business-as-usual. As AI takes over the heavy lifting, data engineers get to step back from pipelines and step up to strategy. The script is changing, the prompt is the new interface, and the future of data integration might just be zero-code, zero-friction, and all impact.

> *Maybe Ø ETL doesn't mean "no ETL" – but it definitely means no more business-as-usual.*

Start innovating now

Experiment with AI-driven data pipelines

Start small with a well-defined ETL use case where manual processes create bottlenecks. Try implementing generative AI to automate a non-critical data flow and measure the time savings and accuracy improvements.

Invest in data literacy and documentation

Improve your metadata management and data documentation practices. High-quality documentation significantly enhances how well AI tools understand your data relationships and can generate appropriate transformations.

Upskill your data teams

Help your data engineers transition from code writers to pipeline architects and quality experts. Create opportunities for them to work alongside AI tools while developing new skills in oversight and optimization.

#AIDataIntegration
#NoCodeETL #DataAutomation
#GenerativeAI #DataMesh
#CloudDataPlatforms

Born today, leading tomorrow

How AI teammates will transform the way we collaborate, create, and lead

Hemavathi Thiruppathi

Enterprise Architect, Insights and Data, Capgemini

Venkadesh Thirupathi

Enterprise Architect, Insights and Data, Capgemini

Meet your AI teammate – born today, built to grow, and ready to transform tomorrow. It's not just a tool, but a thinking partner that evolves with you. This AI never sleeps, always learns, and is already reshaping the way we work. Across industries like manufacturing, telecommunications, finance, and life sciences, AI boosts productivity by taking over repetitive, time-consuming tasks – freeing humans to focus on strategic and creative work. It can prioritize actions, suggest decisions, summarize meetings, and manage workflows like a true team member. To make this collaboration work, we need transparency, trust, explainability, responsibility, and emotional intelligence. AI teammates are changing the workplace, and they bring challenges and opportunities. Here is what it takes to succeed – with both the wins and the lessons learned along the way.

How will the workplace environment evolve in the future?

The workplace of the future won't just be powered by AI – it will be populated by it. As AI teammates become embedded in daily operations, we're not just automating tasks; we're rethinking how work itself gets done. This evolution unfolds across three key dimensions: reimagining workflows, evolving roles, and embracing cultural shifts. Together, they redefine what it means to be part of a team, human or otherwise.

Reimagining workflows

AI integrates into workflows, enabling faster, smarter execution. From automating data entry to generating real-time insights, AI reduces bottlenecks and enhances productivity.

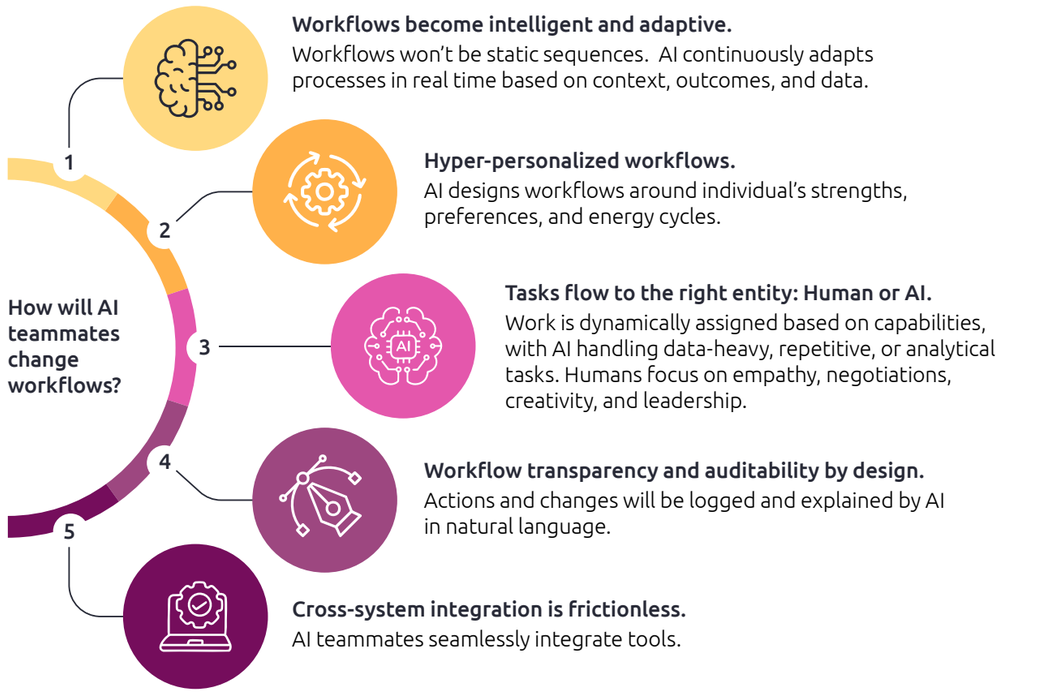

Evolving roles

Roles are shifting from task executors to value creators. Customer service agents now supervise AI chatbots, analysts generate AI-powered insights, and developers orchestrate AI systems. This shift promotes upskilling in areas like prompt engineering, data literacy, human-in-the-loop coordination, and ethical oversight. Emerging roles include:

- **AI Solution Lead:** Designs, prompts, and trains AI systems
- **AI Governance Lead:** Oversees critical decisions AI can't make alone
- **AI Ethical Assurance Lead:** Ensures fairness, bias detection, compliance, and transparency
- **AI Innovation Architect:** Experiments with AI to uncover new value streams
- **AI Transformation Lead:** Communicates AI's role to stakeholders, building confidence and adoption.

Cultural shift

Employees are expected to work with AI, not around it. Leaders help guide this shift by making AI easier to understand, promoting experimentation, and redefining success based on human-AI collaboration.

How will AI teammates change workplace culture?

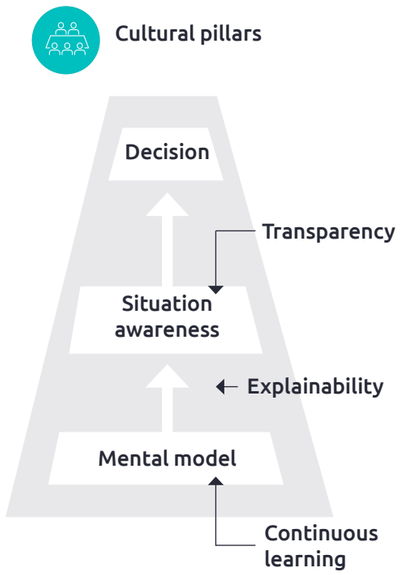

Human-AI collaborations is the norm
AI is not a separate function. It is part of every team. Employees collaborate with AI agents daily.

Decision-making is augmented and faster
AI powered insights guide decisions. Leaders rely on real-time data models and scenarios planning.

Culture of continuous learning
Upskilling is ongoing, aided by AI that recommends personalized learning.

Work-life harmony through AI-enabled flexibility
AI schedules and prioritizes tasks, optimizing work rhythms and preventing burnout.

Inclusive and adaptive teams
AI removes barriers of language, location, and accessibility, fostering diverse teams.

Primary challenges and solutions to tackle them

Category	Challenge	Solutions
Trust and transparency	Employees are skeptical of AI decisions, fear job loss	Transparent AI models with explainable outputs, human-in-the-loop frameworks, internal training on AI benefits and limitations
Skills gap	Workforce lacks AI literacy and collaboration skills	Upskilling programs, role redesign to blend human expertise with AI capabilities, AI onboarding like new team members
Resistance to change	Cultural unwillingness to adopt new workflows	Leadership-driven change management, AI champions within teams, gradual AI integration with feedback loops
Ethical concerns and biases	AI bias and fairness issues	Ethical oversight, bias detection mechanisms, compliance with ethical standards
Data privacy and security	AI needs large amounts of sensitive data	Role-based access, encryption, anonymization, strict compliance with regulations, federated learning, synthetic data
Integration with existing systems	AI tools may not align with legacy systems	APIs and middleware for smooth integration, cloud-based AI services, modular AI that plugs into existing workflows

Successes and setback

How AI teammates drive efficiency and business value-while also raising ethical concerns and the risk of hallucinations.

Education: Capgemini's AI system, developed with UNESCO and AWS, analyzes learning data to provide actionable insights, improving decision-making and learning outcomes in low- and middle-income countries.

Healthcare: The first baby conceived through robot-controlled IVF highlights AI's role, with AI autonomously selecting and preparing sperm under human oversight.

Human resources and recruitment: 7-Eleven's AI assistant, Rita, streamlines hiring by scheduling 85 percent of applicants within an hour.

Healthcare: Google's Med-PaLM 2 hallucinated clinical responses, misleading doctors. Overreliance risks patient safety, despite its role as an AI assistant.

Customer service: Air Canada's chatbot gave incorrect refund info. The court held the airline liable, exposing risks of AI teammates in critical customer interactions.

Media and publishing: *Sports Illustrated* used AI to publish fake-authored articles. Backlash followed, highlighting ethical issues in AI-generated content without proper editorial control.

Given all we've seen, adopting a human-first mindset is essential. To get the best from AI, we must prioritize teamwork, ethics, trust, and support for the people behind the processes. AI teammates may be tireless and proactive, but they still need guidance, context, and oversight. Like any valuable team member, they work best in partnership. True success lies in combining human creativity with AI's capabilities – not replacing one with the other, but elevating both together.

> *The future of work is not about humans versus AI. It is about humans who know how to work with AI.*

Start innovating now

Clear communication and training

Communicate clearly the purpose and goals of the AI Teammate, emphasizing its importance in accelerating productivity. Upskill the team on emerging technologies and train them to understand AI's strengths and limitations. This builds confidence, fosters trust, and empowers people to collaborate effectively with AI systems.

Define roles and responsibilities

Insist on a human-first mindset and clearly define task boundaries between AI and human team members to avoid confusion. Let AI handle repetitive work, while humans focus on creative thinking, decision-making, and interpersonal collaboration. This clarity improves productivity, leverages individual strengths, and enables a balanced and efficient team dynamic.

Nurture a collaborative culture

Establish an AI-human collaborative platform where teams can share knowledge, experiences, and feedback. This approach increases transparency, reduces bias, and drives continuous learning. It enables AI systems to evolve with the team's needs, ensuring alignment with real-world outcomes and fostering stronger, trust-based collaboration.

#Datapowered #AgenticAI #GenAI
#AITeammate #EmotionallyIntelligentAI

The fast and the frugal

Outrun the big models – without draining the grid

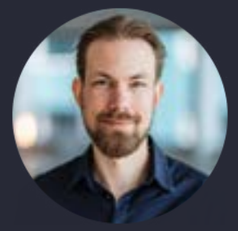

Joakim Nilsson

Knowledge Graph Lead, Insights and Data Sweden, Capgemini

Johan Müllern-Aspegren

Emerging Tech Lead, Applied Innovation Exchange Nordics and Core Member of AI Futures Lab, Capgemini

What if AI could be fast – without being furious to the planet? In the current AI arms race, bigger often means better… and a lot more power-hungry. But when every token costs energy, emissions, and compute, maybe it's time to question the size-over-smarts approach. Enter the frugal model: smaller, contextual, and good enough to hold its own against the LLM heavyweights, without draining the grid. Let's challenge the assumption that scale equals value, and show how leaner, fine-tuned models can achieve surprising results. Call it a new kind of AI performance: less drag race, more precision drift.

The hidden environmental cost of AI: One token at a time

As the AI gold rush accelerates, there's a quiet environmental tax we're all paying – one token at a time. A chat query may not feel like much, but beneath the surface lurk LLMs with a rather intimidating computational appetite. The hyperscalers' top-of-the-line LLMs powering today's smartest chatbots and emerging AI agents churn through electricity and water in quantities perhaps more befitting an aluminum smelter than a software service.

"Our GPUs are melting." – Sam Altman, CEO of OpenAI

Thus, changes in climate are not only driven by the usual suspects, like energy, transport, or agriculture, but increasingly by the digital technologies we embrace. Artificial intelligence, a technology seen by some as a tool to combat climate challenges, has its own environmental footprint. However, the potential of generative AI is undeniable. We shouldn't slow down, but we must get smarter in the way we are using it.

Make every token count

That means making every token count: reducing waste, and designing systems that are not just intelligent, but efficient.

Because LLMs don't just consume data, they consume power. Recent studies show that generative AI, including both model training and user inference, consumed an estimated six to nine terawatt-hours (TWh) of electricity in 2023, comparable to the energy use of a small country. By 2027, AI servers could draw as much as 134 TWh/year, roughly the energy needs of Sweden.

Agents on the rise

The real problem, however, may not only come from the models themselves – but the agents that will be built on top of them. Picture future companies run by a handful of humans and thousands of LLM-powered agents: optimizing code, outwitting other agents in stock trading, crafting legal frameworks, and winning the war for attention with AI-generated content. In this arms race, performance is currency, and currency is performance. The better your agents, the sharper your competitive edge and, today, the top-performing agents are driven by large, cloud-hosted LLMs. They're impressive, certainly, but they're also expensive financially as well as in terms of sustainability.

Watching Vin Diesel drift around

And the inefficiency adds up. As agents chain together models, expand context windows, and embed documents for every micro-task, we're witnessing token inflation on a grand scale. Like watching Vin Diesel drift around a carpark to reach a spot 10 metres away, we're using GPT, Gemini, and Claude top models to rephrase a sentence. So, we don't need to throw a library at a question that can be answered by a reference to an article.

When querying LLMs, verbosity is expensive. Every token adds cost, complexity, and carbon. Each processed token draws computational power. But not every token needs to be spelled out. Many can be implied, if the model has context. In order to give your agent context and enable downsizing, there are three main strategies.

Train a model from scratch. You train a foundation model from the ground up, using your custom dataset. It will be tailored entirely to your domain but can be extremely expensive and requires massive amounts of time and data.

Finetune an existing model. Start with an existing LLM and retrain the upper level on your domain-specific data. Change the "way" rather than the "what" in LLM responses. A great use case is to finetune a model on generating cypher queries on text prompts (as we will see soon).

Augmented retrieval. Keep the model frozen but supply external knowledge at runtime. That's where traditional RAG and its multidimensional cousin GraphRAG makes its entrance, giving the model a shared reference point without overloading it with detail. Rather than cramming background into every prompt, the model can now refer to entities and relationships already mapped in the graph.

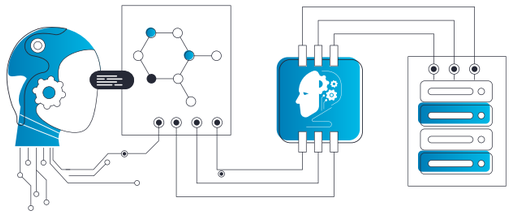

A leap of faith: Will a small, context-aware model cut it?

There's nothing new about this reasoning, but it takes a little leap of faith to switch from the smoke-and-flash of nitrous-fueled drag racers to the quiet grace of a machine engineered exactly for the task at hand.

This is more than just green tech evangelism: it's sound business logic. Using fewer resources is basically good business, delivering reduced latency, lower cost, and independence of data centers and clouds.

The central tension remains: will a small, context-aware model really be enough when it is competing against the full firepower of GPT-whatever running on a nuclear-powered server farm? The battle between efficient precision and brute force brilliance is about to play out here below.

Ladies and gentlemen, start your engines

To find out, we designed a small task to generate cypher queries from natural text, and pitted a finetuned local model of roughly 4 GB against the much larger gpt-4o-mini. To make things more interesting, we also invited the Llama-3.3-70b-versatile from Meta.

The challenge. All three models were fed the same text – a question regarding information in the underlying graph database or a request to update it. They need to generate a valid cypher query which will be executed, and the responses are then compared.

The underlying database consists of a limited set of startups, technologies, and founders – linked together based on data from an ecosystem register.

An example query used for benchmarking: "Which technologies are used across multiple startups?"

Meet the contenders

Provider	Model	Local/cloud	Size
Ollama	tomasonjo/llama3-text2cypher-demo	Local	Small
Groq	llama-3.3-70b-versatile	Cloud	Medium
OpenAI	gpt-4o-mini	Cloud	Large

Note that the tomasonjo/llama3-text2cypher-demo is finetuned to handle text-to-cypher. It is open source, based on the **llama3 model**, and can run fully locally on a laptop.

The cloud models are not particularly trained on cypher, but as it is part of the training material, they have a basic understanding of cypher.

Outcome. We ran a number of different queries. The results so far show that the local model and the gpt-4o-mini perform on par, with the llama-3.3-70b shows slightly less performance.

Note: The local model didn't just keep up – it outpaced gpt-4o-mini, slashing the generation time by more than 50 percent.

Our benchmark shows that when armed withrelevant context, local LLMs can match the performance of massive cloud-based GPT models – without the weight, the latency, or the energy bill. It's a direct challenge to the idea that bigger always means better. With the right architecture, local isn't just a fallback – it's a strategic edge. The race for smarter AI is not about who burns the most fuel. It's about who handles the corners best.

"Any intelligent fool can make things bigger, more complex, and more violent. It takes a touch of genius – and a lot of courage to move in the opposite direction."

– E.F. Schumacher

Results

Local Model - Generated Cypher & Results

Generated Cypher Query:

```
MATCH (s:Startup)-[:USES_TECH]-> (t:Technology)
WITH t.name AS technology, COUNT(s) AS startupCount
WHERE startupCount > 1
RETURN technology
```

Returned 1 records in 0.05 seconds

technology
APOC Library

GPT-4o-mini - Generated Cypher & Results

Generated Cypher Query:

```
MATCH (s:Startup)-[:USES_TECH]-> (t:Technology)
WITH t, count(DISTINCT s) as numStartups
WHERE numStartups > 1
RETURN t.name, numStartups
```

Returned 1 records in 0.11 seconds

Local and gpt-4-o-mini are on par.

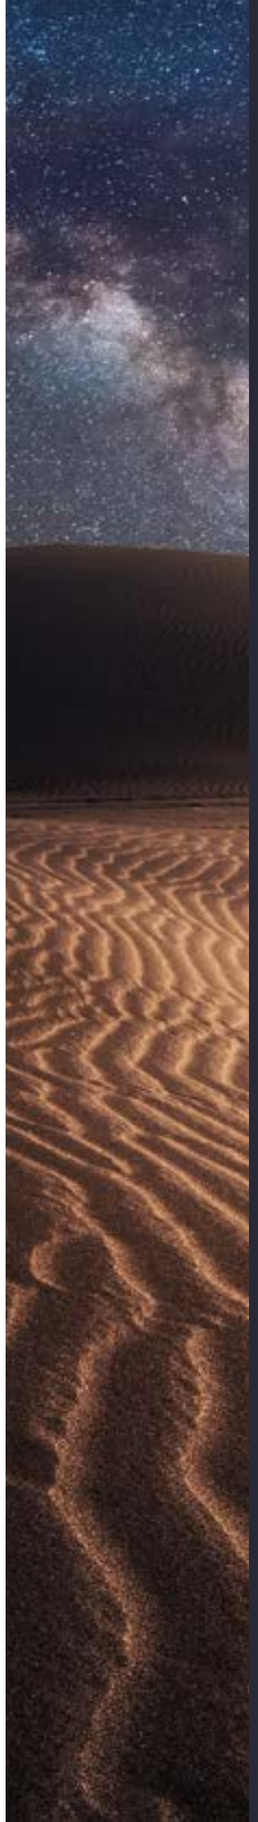

"LLMs don't just consume data, they consume power."

Start innovating now

Take a leap of faith

and go with smaller models; with the use of knowledge graphs, results become more context-aware, efficient, and environmentally sustainable.

Build processes

for quality assurance and automatic benchmark probing with large models.

Continuously measure

and communicate reductions in energy usage and carbon emissions resulting from data operations, fostering sustainable AI practices from the start.

#MultiAgents, #LLM, #LocalModels, #FineTuning, #Sustainability

The compliance upgrade

From policy overload to AI-powered safety and alignment

Dr Abi Giles-Haigh

Analytics and Innovation Director, Capgemini

Liz Henderson

Non-Exec Director, Board Advisor and Speaker, Capgemini

Regulations aren't going away – but the pain of keeping up with them might. From health and safety rules to environmental standards, most organizations struggle to stay aligned while managing mountains of policies, audits, and legalese. But a new generation of AI tools is transforming compliance from a reactive burden into a proactive advantage. AI and generative models are helping businesses bridge the compliance gap – making workplaces safer, processes smarter, and legal alignment less of a firefight and more of a feature.

The rising importance of regulatory compliance

Non-compliance can lead to substantial fines, reputational damage, and even risks to public safety. In this high-stakes landscape, integrating regulatory data into business processes has never been more critical. By embracing a data-driven approach, companies can transform legal complexity from a burden into an advantage, improving safety, efficiency, and agility.

Insights from **edition 8** of this magazine (The integral role of an AI bot, or agent, in your data governance committees: A view from the near future, Pg. 74) illustrate how AI can bridge the gap between unstructured legal documents and structured operations. This not only strengthens compliance but also builds resilience in an ever-changing regulatory environment.

The role of AI in workplace safety and beyond

Industries like water management face complex and evolving regulations similar to workplace safety. For instance, in England and Wales, the water and sewage sectors must comply with various Acts of Parliament and European directives. These layered policies demand continuous adjustments to remain compliant. However, today's digital landscape offers tools capable of handling unstructured legal language – no matter how intricate – and aligning it with structured business processes.

Consider **LegalSifter**, an AI-powered tool that combines machine learning with legal expertise to draft, review, and manage contracts more efficiently. In other industries, AI can help companies identify compliance gaps, flag safety risks, and proactively align processes with evolving regulations. McKinsey research shows that companies leveraging AI for compliance experienced a 30 percent reduction in non-compliance incidents, underscoring its tangible value.

Bridging the gap between regulations and processes

Business processes consist of related, structured activities designed to deliver a product or service. These processes exist at all levels of an organization and are typically represented as flowcharts outlining activities and decision points. By integrating regulatory data into these workflows, businesses can align technology, people, and processes under a cohesive, data-driven strategy. This not only ensures compliance but also creates safer, more efficient work environments.

The benefits of well-defined, AI-enhanced business processes include enhanced customer satisfaction, improved agility, and better employee safety, as evidenced by reductions in workplace injuries and lost working days. By utilizing AI to connect unstructured legislative data with business operations, organizations can improve safety standards, mitigate financial risks, and strengthen their overall resilience.

For example, directly connecting legislation to business processes can uncover potential vulnerabilities. Analyzing these connections helps businesses proactively address compliance issues and enhance safety. Without AI, manual interpretation of unstructured regulations leaves room for errors, increasing the likelihood of penalties, reputational harm, or community risks. A stark example is the **£160 million penalty** faced by Thames, Yorkshire, and Northumbrian Water in August 2024 after a sewage investigation. By using AI to align unstructured regulations with business processes, organizations can significantly reduce such risks.

Generative AI: The next frontier in compliance

Generative AI is transforming how businesses handle unstructured data, making compliance smarter and faster. By processing complex legal documents, AI can help design practices that not only ensure compliance but also enhance safety. Imagine a future where businesses can "discuss" legislation with conversational AI to verify whether their processes meet legal standards. Utilizing Retrieval Augmentation Generation (RAG) to link the policies can provide the relevant context to the conversation. Gen AI could also interpret laws and generate templates aligned with regulatory requirements, reducing costly fines and improving operational efficiency.

Moreover, generative AI empowers employees by increasing their understanding of relevant laws and offering decision-making tools. For example, AI-driven chatbots could provide conversational access to policy documents, enabling employees to explore legal requirements with ease. In HR, generative AI can help ensure compliance with working time directives by aligning schedules and processes with legal standards.

However, ethical considerations are crucial. Businesses must prioritize data privacy, mitigate bias in AI recommendations, and maintain accountability for decisions made using AI tools. Responsible deployment of these technologies will build trust and ensure equitable outcomes

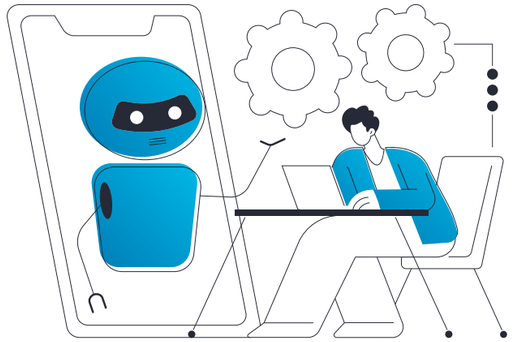

The case for data-driven business process mapping

Business processes consist of structured activities designed to deliver products or services. These processes, often represented as flowcharts, outline activities and decision points at all organizational levels. By integrating regulatory data into workflows, businesses can align technology, people, and processes into a cohesive, data-driven strategy. This ensures compliance while fostering safer, more efficient work environments.

AI-enhanced business processes lead to numerous benefits: enhanced customer satisfaction, improved agility, and better employee safety. These advancements are reflected in reduced workplace injuries and fewer lost working days. By connecting unstructured legislative data with operations, organizations can improve safety standards, mitigate financial risks, and build resilience.

Upgrade now

A compliance upgrade is no longer optional. In a world of growing policy overload, AI offers a smarter path forward – one where compliance is no longer a reactive burden but a proactive, data-driven advantage. By starting small, scaling responsibly, and embedding AI into the core of business operations, organizations can align safety, efficiency, and regulatory clarity. This is more than a tech shift – it's a mindset shift. And it's how the next era of compliance gets built.

Today's digital landscape offers tools capable of handling unstructured legal language – no matter how intricate – and aligning it with structured business processes.

Start innovating now

Pilot AI tools in high-risk compliance areas

Select one department and integrate an AI tool to automate compliance monitoring or document interpretation, then measure improvements in efficiency and error reduction over three to six months.

Map key business processes against regulatory requirements

Conduct a compliance-process mapping workshop with legal, operations, and tech teams, and feed the results into an AI platform capable of simulating compliance scenarios.

Develop an AI governance framework

Create a cross-functional AI ethics and compliance committee to oversee deployment, monitor data privacy practices, and evaluate algorithmic fairness across tools.

#IndustrialGenAI #Compliance #ReducingComplexity

Mind the gateway

Between the platform and the promise: Gateways keep Gen AI on track

Kumar Chinnakali

Portfolio Manager, Hands-on Contact Centre Management Architect, Insights and Data, Financial Services, Capgemini

Generative AI may be the express train to innovation, but without the right infrastructure, an enterprise could be left stranded at the platform. AI gateways are emerging as the crucial link between bold Gen AI ambitions and the operational realities of security, scalability, and control. They don't just route traffic – they enforce trust, optimize costs, and let teams build smarter, faster, and safer. So before you board the Gen AI hype train, make sure you mind the gateway.

Generative AI and large language models (LLMs) are transforming the way businesses operate, from automating customer interactions to powering smarter decision-making. But a key question often goes overlooked: **Is the infrastructure truly AI-ready?**

While traditional networking components do a decent job of securing general traffic, they simply don't understand or manage the complexities of AI workloads.

Evaluating AI-readiness goes beyond hardware and software capabilities. AI workloads demand refined control over data privacy, latency, and resource allocation. AI gateways fulfill this need by providing real-time semantic analysis, precise data masking, and intelligent caching.

These capabilities allow organizations to detect anomalies early and adjust in real time, preventing both data breaches and cost overruns. Such advanced filtering and control mechanisms are essential for maintaining a secure environment while scaling generative AI applications. By integrating these systems, enterprises can enhance compliance, operational agility, and the overall performance of their AI infrastructure.

Why AI gateways matter

AI workloads need more than firewalls and routers. Standard infrastructure doesn't know how to filter sensitive data from AI prompts, control costs tied to LLM usage, or monitor how AI models perform.

That's where AI gateways come in. Think of them as intelligent API gateways built specifically for AI systems.

- Protect sensitive data with semantic filtering and masking.
- Control runaway costs using rate limits and smart caching.
- Improve visibility with deep usage insights and metrics.

AI gateways create a **secure, scalable, and cost-efficient path** to Gen AI integration essential for today's enterprise landscape.

AI gateways are designed to implement advanced filtering and rate limiting that align with modern security protocols. They combine rule-based logic with machine learning insights to accurately identify sensitive information and manage high-volume AI traffic. This dual-method approach not only enhances security but also optimizes system performance.

Integration with observability tools provides granular insights into system behavior, enabling proactive adjustments. These measures are critical in preventing data leaks and cost escalations, ensuring that the AI models operate efficiently under varying load conditions. This thoughtful integration supports both immediate operational needs and long-term scalability.

Examples: Industry momentum

Leading platforms like **Gloo AI Gateway (KGateway)** and **Portkey AI Gateway** are already enabling enterprises to scale Gen AI securely. They tackle the common pitfalls of deployment, lack of observability, rising cloud costs, and shadow AI use through tight integration and smart automation.

Best practices: Design AI gateways for scale

Building a future-ready AI gateway? Keep these considerations in mind.

- **Separate concerns:** Keep AI logic out of business logic to reduce complexity.
- **Use smart proxies:** Tools like Envoy help with data masking, access control, and traffic shaping.
- **Go declarative:** Use YAML/JSON for managing policies to enable clear, versioned, and DevOps-friendly setups.
- **Open source first:** Lean on open-source tools to avoid vendor lock-in and stay flexible.

Implementing AI gateways represents a strategic evolution rather than a simple technical upgrade. The modular design enables iterative improvements and easier integration with existing systems. By adopting declarative configurations and leveraging smart proxies, organizations can maintain clear separation between AI-specific processing and core business logic.

This separation minimizes risks associated with unmonitored AI activity and uncontrolled costs. Additionally, integrating best practices like GitOps and Infrastructure as Code fosters a culture of continuous improvement and agile deployment. These measures collectively empower teams to innovate responsibly while keeping systems secure and efficient.

Market trends: The rise of AI platform engineering

Enterprises are shifting towards AI platform engineering, creating self-service systems for

developers to securely use AI services. AI gateways are a central component of this shift.

With GitOps, IaC, and automation-first strategies, enterprises can ensure governance without slowing teams down.

The result? Faster Gen AI adoption, fewer risks, and smarter use of internal resources.

Emerging paradigms and future directions

As enterprises embrace AI platform engineering, AI gateways become central to secure and scalable innovation. Their integration facilitates a seamless flow of data and control across diverse environments, bridging legacy systems with modern cloud-native architectures.

The real-time monitoring and policy enforcement capabilities ensure that even rapid iterations do not compromise security or cost-effectiveness. This forward-thinking approach fosters a self-service model where developers can confidently experiment while adhering to strict governance.

As AI continues to mature, the synergy between advanced infrastructure and intelligent gateways will be the cornerstone of sustainable digital transformation, driving both innovation and resilience. This integration not only streamlines processes but also sets a benchmark for industry best practices.

Why it matters

Generative AI isn't waiting at the station – it's accelerating fast. But without the right infrastructure, even the most promising initiatives risk derailing due to cost overruns, data leaks, or shadow deployments. AI gateways offer more than just safe passage. They're the intelligent infrastructure that keeps your AI ambitions securely on track. With fine-grained control, observability, and built-in governance, they ensure that innovation arrives not just quickly, but safely and sustainably. In a world where everyone is racing to deploy Gen AI, the real differentiator isn't speed, it's control. So before you scale, automate, or fine-tune that next model, remember: mind the gateway.

Traditional networking

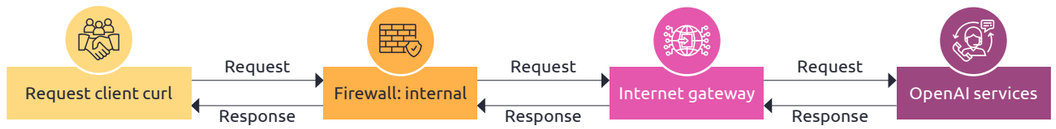

Direct versus gateway-mediated access to a large language model

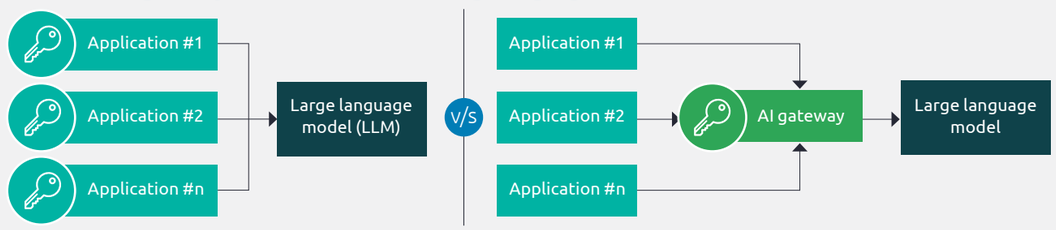

Integration of enterprise secrets management system with an AI gateway architecture

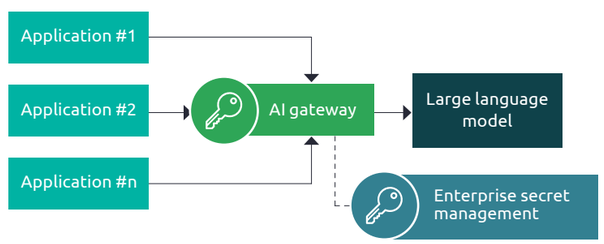

Content sanitization workflow with input and output guardrails within an AI gateway

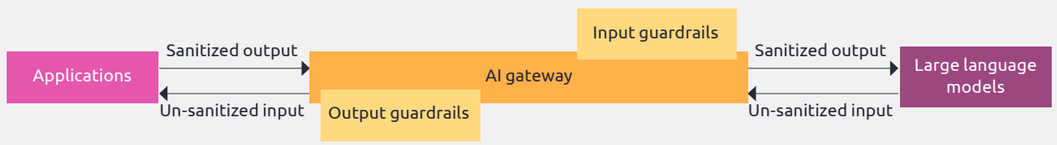

An AI gateway is more than just a security layer. It's the nervous system that makes safe, scalable AI integration possible.

Start innovating now

Assess your AI readiness

Start by evaluating if your current infrastructure supports semantic security, cost tracking, and observability for LLMs.

Introduce an AI gateway

Deploy a purpose-built AI gateway. Begin with pilot use cases and evolve your setup using modular design and declarative configuration.

Embed into your platform

Integrate AI gateways into your internal platforms using GitOps or IaC. This ensures self-service, speed, and governance all at once.

#AI4Enterprise #AIInfrastructure #AIGateway #EnterpriseAI #SmartScaling

Agents of confidence

Designing trust, autonomy, and alignment in an AI-native world

Robert Engels

AI CTO, AI Future Lab Lead,
Insights and Data, Capgemini

Mark Roberts

Deputy Head, Capgemini Generative AI Lab,
Capgemini Engineering

Agentic and autonomous systems aren't just redefining AI's capabilities – they're reshaping how we build trust in machines that act on our behalf. From thermostats to multi-agent decision networks, the shift is on: we're moving from tools that assist to agents that decide, adapt, and execute. But with great autonomy comes the pressing need for clarity, control, and context. This article breaks down a multidimensional framework for confidence in agentic AI systems – including autonomy, agency, authority, world models, and alignment – and highlights what it really takes to ensure these systems remain useful, reliable, and aligned with human intent. Because if the machines are making decisions, we might want to check what playbook they're using.

The rise of the agentic mindset

The AI world is shifting from assistance to agency. We're no longer just asking systems to predict next steps or summarize emails – we're expecting them to take action, make decisions, and operate semi-independently. This evolution brings us into the era of agentic AI: systems that are not just capable of interpreting information, but of executing plans with autonomy and authority.

It sounds like a superpower. But with every degree of freedom we give these systems, the demand for trust grows exponentially. We need to know not just what they do, but why, how, and within which limits. Confidence in automation isn't a matter of blind faith – it's a matter of well-placed design.

Beyond the co-pilot

The journey to agentic systems has been gradual. First came simple chatbots, then co-pilots that could make helpful suggestions. Then came autopilot systems that acted within narrow domains. Now, we are seeing the emergence of multi-agent systems: networks of independently operating AI agents collaborating in complex environments.

But more isn't always better. As systems grow in autonomy, we need more nuance, and better mental models, to describe, govern, and trust them. Giving AI the wheel doesn't mean we throw away the map.

Autonomy, agency, authority: The triple A of trust

Understanding agentic AI starts with three key attributes:

- **Autonomy** is the system's ability to make decisions independently
- **Agency** is its capacity to act on those decisions
- **Authority** defines the boundaries of what it is allowed to do

Think of a thermostat: high autonomy and agency, but limited authority. A secret agent? High across the board. These dimensions aren't binary – they're sliding scales that can be tuned depending on context, risk, and purpose. Getting that balance right is what separates a helpful AI teammate from an unpredictable digital loose cannon.

System design: It takes an ecosystem

Agentic AI isn't just about single actors. The real shift happens in multi-agent systems: swarms of intelligent agents coordinating, collaborating, or even competing.

Here, four design dimensions emerge:

- **Scale:** from single-agent to vast networks
- **Composition:** from homogeneous swarms to heterogeneous specialists
- **Complexity:** from simple workflows to emergent behaviors
- **Centralization:** from tight control to fully distributed decision-making

These variables drive not just system capability but governance requirements. Designing for trust means deliberately balancing sophistication with control. Think less swarm theory, more symphony, in which each agent must know its instrument and when to come in.

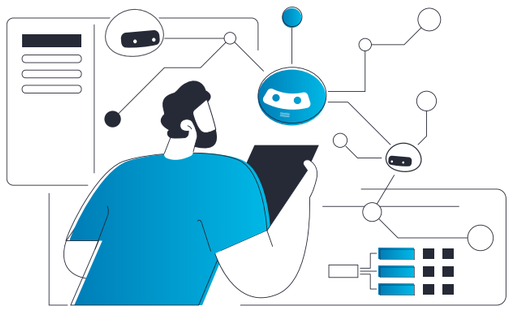

Payloads, not just personalities

What agents do matters just as much as how they work together. This is the idea of the payload view: assessing agents based on what tasks they can accomplish.

Agents vary along two more dimensions.

- **Specialist vs. generalist:** deep expertise versus broad adaptability
- **Deterministic vs. non-deterministic:** predictable behavior versus learning and evolving

A complex supply chain AI may need specialized, non-deterministic agents to adapt in real-time. A compliance bot? Best keep it deterministic. In other words, don't ask a jazz improviser to run payroll.

LLMs: Powerful, but not the agent

Contrary to popular belief, large language models (LLMs) are not agents. They excel at interpretation – translating between human language and machine inputs – but rarely possess real agency.

They can be part of an agentic system, especially as natural language interfaces, but they are not a substitute for agents that act. Confusing the two leads to overestimated capabilities and underestimated risks. Just because it sounds convincing doesn't mean it knows what it's doing.

Why world models matter

Agents rely on internal representations of their environment (world models) to make informed decisions. The richer and more accurate the model, the better the agent can act in alignment with goals.

Imagine a thermostat that not only senses temperature but also understands user habits, energy prices, and weather forecasts. Now apply that principle to business: from dynamic pricing to policy design, world models are the foundation for contextual, trustworthy AI.

The alignment imperative

The final, and arguably most important dimension, is alignment. What is the agent's purpose? And is it behaving in a way that truly reflects that purpose?

Misalignment is where the **paperclip-making nightmares** begin: an agent optimizing the wrong metric with catastrophic results. Avoiding this requires clarity of goals, careful training, and a robust feedback loop between system outputs and human intent. In short: tell your agents what to do, but make sure they also know why.

Design confidence, not just capability

Building agentic systems isn't just a technical challenge. It's a design opportunity to reimagine what it means to delegate, collaborate, and trust machines.

As businesses begin to embed these systems, the winners won't just be those with the most advanced technology. They'll be the ones who build AI that works reliably, safely, and transparently in service of people.

Confidence isn't the end goal. It's the starting condition for everything agentic AI could become. The rest? That's just execution. And hopefully not by accident.

Note: *chatGPT 4o, trained with Data-powered Innovation Review articles, summarizing the Business, meet agentic AI – Confidence in autonomous and agentic systems whitepaper by Capgemini's AI Futures Lab.*

> *The moment AI stops asking for permission is the moment we need to start demanding purpose.*

Start innovating now

Audit existing AI agents

for autonomy, agency, and authority levels.

Design small multi-agent pilots

to explore coordination models.

Define world models explicitly

and test them for coverage and realism.

Stress-test purpose alignment

across real-world scenarios.

#AgenticAI #TrustInAI
#AutonomousSystems
#AIAlignment #AIByDesign

Ananda Deshmukh
AI/ML, Analytics, Data Science, Insights and Data, Capgemini

Arne Rossmann
Innovation Lead, Insights and Data, Capgemini

Dr. Daniel Kühlwein
Head of GenAI, Insights and Data Germany, Capgemini

Dinand Tinholt
Vice President, Insights and Data North America, Capgemini

James Wilson
Generative AI Architect, Insights and Data, Capgemini

John Lewsey
Principal Solution Architect, Insights and Data, Capgemini

Liz Henderson
Executive Advisor, Insights and Data, Capgemini

Dr. Akhauri Prakashkumar
Chief Strategy Advisor, Insights and Data, North Central Europe, Capgemini

Christopher Scheefer
Group Sustainability Accelerator Lead for Data and AI, Insights and Data, Capgemini

Daniela Rittmeier
Head of Data and AI CoE, Insights and Data, Capgemini

Gurumurthy Sethuraman
VP, Strategic Initiatives, Insights and Data India, Capgemini

Joakim Nilsson
Knowledge Graph Lead, Insights and Data Sweden, Capgemini

Kumar Chinnakali
Managing Delivery Architect, Insights and Data, Capgemini

Marijn Markus
AI Lead, Managing Data Scientist, Insights and Data, Capgemini

Mark Oost
Vice President, Global Offer Lead AI and Generative AI, Insights and Data, Capgemini

Monish Suri
Global Google Partnership Lead, Insights and Data, Capgemini

Rajesh Iyer
Vice President, Global Head of ML and Generative AI, Insights and Data, Financial Services, Capgemini

Ron Tolido
Chief Technology Officer, Insights and Data, Capgemini

Dr. Sergey Patsko
Data and AI Deputy Group Offer Leader, Capgemini

Sjoukje Zaal
Chief Technology Officer, Insights and Data Europe, Capgemini

Weiwei Feng
Global Generative AI Portfolio Tech Lead, Insights and Data, Capgemini

Mark Roberts
Deputy Head, Capgemini Generative AI Lab, Capgemini Engineering

Nicolas Claudon
Chief Technology Officer, Insights and Data, South Central Europe, Capgemini

Robert Engels
AI CTO, AI Future Lab Lead, Insights and Data, Capgemini

Roosa Säntti
CEO, Insights and Data Finland, Capgemini

Shyamsree Nandi
Vice President, Insights and Data, Financial Services, Capgemini

Sudarshan Sahu
Team Lead, Emerging Technology Team, Data Futures Domain, Capgemini

Yashowardhan Sowale
Chief Technology Innovation Officer India, Architecture Head, India Domain Leader for AI, Insights and Data, Capgemini

Editorial team

Ron Tolido

Chief Technology Officer,
Insights and Data, Capgemini

Robert Engels

AI CTO, AI Future Lab Lead,
Insights and Data, Capgemini

Debleena Misra

Innovation Marketing Lead,
Insights and Data, Capgemini

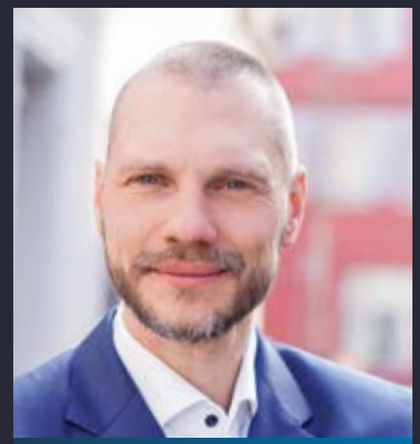

Arne Rossmann

Innovation Lead,
Insights and Data, Capgemini

About Capgemini

Capgemini is a global business and technology transformation partner, helping organizations to accelerate their dual transition to a digital and sustainable world, while creating tangible impact for enterprises and society. It is a responsible and diverse group of 340,000 team members in more than 50 countries. With its strong over 55-year heritage, Capgemini is trusted by its clients to unlock the value of technology to address the entire breadth of their business needs. It delivers end-to-end services and solutions leveraging strengths from strategy and design to engineering, all fueled by its market leading capabilities in AI, generative AI, cloud and data, combined with its deep industry expertise and partner ecosystem. The Group reported 2024 global revenues of €22.1 billion.

www.capgemini.com

Capgemini

Get the future you want