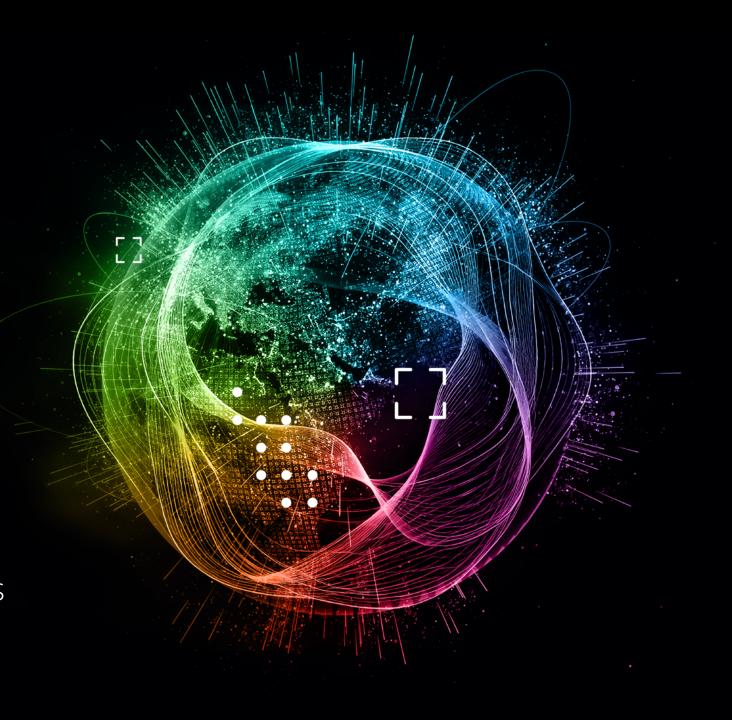
Deloitte.

The AI Dossier

A selection of high-impact use cases across six major industries

Aotearoa New Zealand edition

By the Deloitte Al Institute



About the Deloitte Al Institute

The Deloitte AI Institute™ helps organisations connect all the different dimensions of the robust, highly dynamic, and rapidly evolving Artificial Intelligence ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the "Age of With™."

The Deloitte AI Institute aims to promote the dialogue and development of AI, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, the future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, delivers impactful perspectives to help organisations succeed by making informed AI decisions.

No matter what stage of the Al journey you are in: whether you are a board member or a C-Suite leader driving strategy for your organisation – or a hands-on data scientist bringing an Al strategy to life – the Deloitte Al Institute can help you learn more about how enterprises across the world are leveraging Al for a competitive advantage. Visit us at the Deloitte Al Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet-ups and live events. Let's explore the future of Al together.





Foreword

Artificial intelligence is now an embedded part of business in New Zealand.

Over the past two years, we have seen Al move from experimental projects to tools that help organisations solve real problems, improve efficiency, and make better decisions. The technology is developing quickly, and its ability to understand complex information and take action is changing how we work.

This compendium brings together examples from six major New Zealand industries:



Consumer



Energy, Resources & Industrials



Financial Services



Government & Public Services



Life Sciences & Health Care



Technology, Media & Telecommunications

Each case shows how AI is being used to address specific business challenges, improve processes, and deliver results. We also highlight where agentic AI is involved to show how autonomous systems are being applied in practice.

With new technology comes new responsibility. New Zealand organisations are expected to use AI in ways that are transparent, fair, and secure. Throughout this compendium, we reference Deloitte New Zealand's 'Trustworthy AI Framework', which sets out practical steps for building systems that are reliable, explainable, and accountable. These principles are essential for building trust with customers, regulators, and the wider community.

Al is developing rapidly, and we are only beginning to see its potential. Our goal is to provide practical examples that help New Zealand organisations understand what Al can do right now, and to encourage thoughtful adoption that delivers real value. We hope these use cases inspire new ideas and support successful, responsible use of Al across the country.



Mike Horne
Chief Executive
Deloitte New Zealand

A word from our experts

With the rapid adoption of Al and the rise of agentic systems, establishing strong guardrails for responsible use is more important than ever. New Zealand's evolving regulatory landscape and OECD-aligned principles require organisations to embed fairness, transparency, and accountability into every deployment. As agentic Al systems take on more autonomous roles, continuous monitoring and scalable human-on-the-loop processes are essential to maintain trust and compliance.

As Kiwi organisations embrace generative and agentic AI, embedding trustworthy principles and strong governance is essential. Agentic systems introduce new challenges in transparency, auditability and risk management, requiring proactive oversight. By aligning with New Zealand's Al strategy and adopting scalable practices such as ethical design, robust data governance and integrated operating models, organisations can innovate responsibly. When ethics, resilience and explainability converge with business integration, Al becomes fit for purpose and trusted by design.

Generative and agentic Al are transforming assurance and advisory functions, from automating multi-step reconciliations to enhancing compliance checks. To realise these benefits, organisations must embed technologies in their platforms and pair Al-driven efficiency with skilled human oversight, ensuring outcomes are robust, reliable, and aligned with evolving regulatory standards.

Agentic AI is poised to accelerate data analysis and decision-making in New Zealand's M&A and corporate finance landscape. As autonomous systems take on higher-value tasks, balancing speed with governance and transparency will be critical. Organisations and advisors must ensure that Al-driven insights are explainable, defensible and auditable in order to support better-informed decisions. mitigate risk, and drive successful transaction and financial outcomes.

Agentic Al offers significant potential for New Zealand's tax professionals, enabling real-time compliance automation and proactive risk management. Embedding trust frameworks and robust governance early will be key to unlocking these benefits, supporting tax professionals as they navigate increasing regulatory complexity and deliver enhanced value to stakeholders.

As agentic AI becomes more accessible, New Zealand's SMEs have an opportunity to unlock productivity gains that were once reserved for large enterprises. What matters most for private businesses is not adopting the flashiest tools, but embedding AI in pragmatic, commercially grounded ways that strengthen decisionmaking, cash flow resilience, and customer experience. When applied with discipline and trusted governance, Al can help Kiwi businesses scale their ambition without scaling their overhead.

Al is redefining how New Zealand organisations empower leaders and teams to unlock new ways of working. As agentic AI capabilities emerge, we see a shift from simple automation to systems that can plan and execute tasks with minimal human intervention. To harness these benefits, organisations must invest in upskilling, robust governance, and transparent oversight, ensuring Al is aligned to strategic goals and public trust.



Mariette Van Niekerk Trustworthy Al Lead



Aravind SubramanianPartner, Al & Data



Silvio BruinsmaPartner, Assurance & Advisory



Daniel LawrenceDirector, Transactions



Phillip Claridge
Director, Tax



Dan HellyerPartner, Business Advisory



Dr. Amanda WilliamsonDirector, Deloitte New Zealand Al Institute

What is inside



Consumer

6



Energy, Resources & Industrials

4



Financial Services

66



Government & Public Services

94



Life Sciences & Health Care

124



Technology, Media & Telecommunications

154



The Consumer Al Dossier

















The Consumer Al Dossier

AI has already become an integral part of people's everyday lives, whether they realise it or not. And compelling new use cases for AI in the consumer space continue to emerge. From generative models that create rich, personalised content to autonomous agentic systems that can plan and execute tasks on a user's behalf - in areas like product design, pricing, and supply chain operations - Al is redefining how consumers discover, evaluate, and interact with brands. These tools are not just enabling better experiences; they are resetting consumer expectations across search, service, commerce, and entertainment.

For consumer companies, this shift creates both opportunity and urgency. Al can drive real-time customer engagement, intelligent automation, and more adaptive decision-making. Large language models and agentic Al systems are now capable of handling sophisticated business functions with minimal human intervention, opening new possibilities for operational efficiency and innovation.

However, leveraging these capabilities at scale requires more than technical integration. It demands unified data infrastructure, strong governance, and a willingness to reimagine core business processes. Leading companies are using Al not just to cut costs but to launch new products, redesign customer journeys, and compete on speed, relevance, and personalisation.

As regulation evolves and public scrutiny grows, sustainable advantage will come from deploying AI with transparency, oversight, and measurable impact. The winners will not be those who adopt AI the fastest, but those who align it best to strategic goals, operational realities, and consumer trust.

Al is redefining

brands.

how consumers

discover, evaluate,

and interact with

Note: The tags below each use case indicate its primary business function and whether Agentic AI is used.

Primary business function

Agentic Al

Dynamic pricing and inventory optimisation

Coordinating price and stock decisions in real time

Agentic AI systems can use multiple specialised agents to monitor a wide range of internal and external signals, then dynamically adjust prices, promotions, and inventory to optimise business performance.

ISSUE/OPPORTUNITY

In many retail environments, pricing and inventory decisions are made using fixed rules and periodic adjustments. This approach can leave money on the table when market conditions change quickly. It can also create costly overstocks when demand softens.

Businesses relying on traditional processes can't respond quickly enough to events like a competitor running out of stock, a sudden weather change, or a viral trend shifting demand. Also, by treating pricing and inventory

management as distinct processes, many retailers miss opportunities for joint optimisation. For example, an item might be discounted without considering replenishment timing, or stock might be held for too long at full price when a strategic promotion could accelerate sell-through.

Agentic AI can unify these activities, with specialised agents continuously collaborating to balance profitability, stock levels, and customer satisfaction.

HOW AI CAN HELP

Pricing optimisation

A pricing agent can continuously learn the price elasticity of each product and track competitor prices, adjusting in real time to capture revenue opportunities, avoid unnecessary markdowns, and react to changing market conditions.

Inventory management

An inventory agent can monitor stock levels across stores and warehouses, factoring in lead times and supply constraints to ensure replenishment decisions align with projected demand and pricing strategies.

Demand forecasting

A demand forecasting agent can analyse signals from internal sales trends, online search patterns, social media, weather forecasts, and local events to anticipate surges or dips in nearterm demand.

Promotions and bundling

A promotions agent can design targeted offers and product bundles (e.g., pairing slow-moving items with high-demand products), scheduling them based on real-time sales velocity and inventory.

Collaborative decision-making

All agents in the process share a common situational awareness and negotiate trade-offs. For example, if the demand agent forecasts a surge, the pricing agent might raise prices while the promotions agent delays discounts; conversely, if oversupply is detected, price reductions and targeted promotions might be implemented in specific regions or channels.













gs Sales

Agentic Al













Dynamic pricing and inventory optimisation

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

Because frequent price changes can be perceived as unfair or arbitrary, dynamic pricing agents should operate within clearly defined policies and thresholds to ensure consistent treatment of customers across channels and regions.



Robust and reliable

Bad data can lead to bad decisions. Agents should be designed with strong data validation and filtering processes to avoid reacting to false signals (such as misinterpreted social trends or inaccurate sales figures).



Transparent and explainable

Dynamic price and promotion changes can confuse customers and internal teams alike. To help address the problem, agents should provide clear reasoning for adjustments, including the data sources and logic used, so pricing and category managers can interpret and communicate the rationale.



Responsible and accountable

Rapid pricing and inventory actions can have strategic and reputational impacts. Al-driven decisions should align with the organisation's brand strategy, operating capacity, and regulatory requirements, with final oversight provided by qualified human managers.

POTENTIAL BENEFITS

Increased revenue and margins

Dynamic, coordinated decisions can capture additional profit during high-demand periods and optimise sell-through on slow-moving products.

Reduced waste and overstock

By aligning pricing and replenishment strategies, excess inventory – especially perishables – can be cleared before it becomes unsellable.

Improved customer experience

Timely, relevant promotions increase customer satisfaction and loyalty while maintaining trust in pricing fairness and reducing stockouts.

Al-orchestrated product design

Automated, end-to-end product design powered by AI agents

Agentic Al systems can orchestrate the entire product design lifecycle – from market sensing to concept creation, product development, and iteration – continuously adapting to market changes in real time.

ISSUE/OPPORTUNITY

Traditional product design in the consumer industry is often a linear, stage-gated process that can take months or even years from concept to launch. Also, from hundreds of ideas, often just one or a few options are commercialised. Although this limited approach helps manage complexity, it also slows innovation and impairs an organisation's ability to respond quickly to shifting consumer tastes or competitive moves.

The challenge is compounded by siloed functions, with design, sourcing, marketing, and supply chain often operating independently on different data systems and timelines. As a result, valuable insights from sales data, customer feedback, or social trends may not inform product development until it's too late.

Agentic AI can make the entire product design process dynamic and continuously adaptive, reducing time-to-market, unlocking new levels of creativity, and enabling better and faster alignment with what consumers actually want.

HOW AI CAN HELP

Market sensing and opportunity identification

A market sensing agent can analyse real-time data from trend reports, social media, consumer sentiment, and competitive intelligence to identify unmet needs and emerging product opportunities.

Concept generation and feasibility analysis

A concept agent can create diverse and innovative product ideas informed by market insights while a feasibility agent evaluates each idea against sourcing options, production cost, manufacturing timelines, and regulatory constraints.

Design development and prototyping

A design agent can produce technical specifications and digital prototypes, enabling rapid iteration and deep product visualisation without the time and cost needed to create physical samples.

Validation and dynamic iteration

A validation agent can test designs against historical performance, customer feedback, and simulated market conditions, while a coordination agent can orchestrate updates across product lifecycle management, marketing, and supply chain systems to adjust plans in real time.







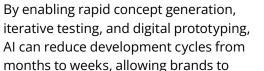








POTENTIAL BENEFITS



respond quickly to market opportunities.



Increased innovation

Faster time-to-market

Generating new and diverse ideas more quickly - in greater volume - expands the creative possibilities for new product development.



Innovative design that aligns with real-time market shifts increases the likelihood new products will be a hit with target customers.

Lower development costs

Digital prototyping and early-stage feasibility analysis reduce the need for costly physical samples and late-stage redesigns.

Al-orchestrated product design

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Because Al agents might produce concepts that are technically or commercially unviable, outputs should be validated in controlled simulation environments and reviewed through human-in-the-loop processes before advancing to production.



explainable

Transparent and Since design decisions can have major cost, brand, and regulatory implications, agents should provide clear reasoning and evidence for their recommendations, including source data and assumptions.



accountable

Responsible and Products must comply with safety and regulatory requirements Also, legal questions remain over intellectual property (IP) rights for Al-generated outputs. IP protection and ownership rights can be complex when Al is involved in the creative process. To address such issues, Al-driven design activities should align with brand standards and legal constraints, with final approvals retained by qualified human decision-makers.







Next-generation store operations

Autonomous in-store coordination to optimise retail execution

Agentic AI systems can coordinate in-store activities by continuously monitoring conditions and taking automated actions to achieve smooth, efficient, customer-responsive operations.

ISSUE/OPPORTUNITY

Running a high-performing retail store involves hundreds of large and small decisions each day: allocating staff to handle peak traffic, restocking shelves when inventory runs low, responding to customer requests, and ensuring that promotional displays are set up correctly. In many cases, these actions are handled reactively, based on direct observation by a manager or sales associate, rather than being driven by data in real time.

This reactive approach can lead to problems such as stock-outs, bottlenecks at checkout, haphazard execution of merchandising plans, and missed sales opportunities – operational frictions that can quickly erode revenue, profitability, and customer satisfaction.

Agentic AI can help stores become highly efficient, semi-autonomous systems – where human associates focus on value-added service and strategic priorities while AI handles routine operational tasks.

HOW AI CAN HELP

Continuous store sensing

A store sensing agent can monitor real-time data streams from cameras, IoT sensors, POS systems, and digital twins to track foot traffic, queue lengths, inventory levels, associate availability, and local events.

Dynamic task allocation

A task management agent can reprioritise and assign tasks such as restocking, returns processing, online order pickup, or promotional setup based on current demand and available labour.

Automated compliance monitoring

A compliance agent can use computer vision and sensor data to monitor planogram adherence, promotion execution, and safety hazards, triggering immediate corrective actions as needed.

Coordinated multi-agent oversight

A store manager agent can oversee all other agents, resolving conflicts, optimising labour deployment, and coordinating with upstream systems such as workforce management, ERP, and order management platforms.













Operations

Agentic Al















Next-generation store operations

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Because retail environments often involve incomplete or questionable sensor data, agents should be designed to function effectively under imperfect conditions, with human escalation protocols in place when Al output is uncertain or fails.



explainable

Transparent and To avoid blind reliance on Al agents, managers and staff should have access to clear explanations of the rationale behind Aldriven decisions.



Responsible and accountable

Store operations must comply with a wide range of standards, including labour laws, safety requirements, and company policies. All Al-driven actions should align with these standards, with ultimate accountability retained by human supervisors.

POTENTIAL BENEFITS

Improved labour productivity

Associates can spend more time on highvalue tasks and less time on low-value activities and manual monitoring.

Higher sales conversion and customer satisfaction

Automated store operations can help optimise in-stock rates, checkout lines, and service levels.

Manager bandwidth for strategic decision-making and leadership

Store managers can focus on performance improvement, training, and coaching, rather than spending their days fighting operational fires.

Autonomous supply chain operations

Using AI agents to improve efficiency in global automotive supply chains

Agentic AI systems can improve the efficiency and resilience of automotive supply chains by using specialised agents to forecast demand, optimise planning, detect disruptions, and autonomously adjust operations.

ISSUE/OPPORTUNITY

Automotive supply chains are complex and vulnerable to disruptions from shifting demand, supplier delays, logistics bottlenecks, and external forces such as pandemics, policy changes, and weather. Traditional supply chain processes rely heavily on periodic data reviews and manual adjustments, which often cannot keep pace with sudden changes in demand and supply. These limitations can lead to higher costs, supply delays, and increased operational risk.

With tariffs, global market volatility, and various sustainability pressures (including electrification) reshaping the industry, automakers need supply chains that are dynamic, predictive, and capable of adapting in real time. Agentic Al provides a pathway to autonomous supply chain operations that can be more flexible, efficient, and resilient.

HOW AI CAN HELP

Data readiness and transformation

A data readiness agent can perform quality checks and identify exceptions, while a data generator agent can transform raw inputs into structured data for optimisation.

Optimisation and demand mapping

A suggestion optimisation agent can run AI/ML models to autonomously identify the best-performing options, while a demand mapping agent can align demand signals with the correct product configurations.

Validation and explainability

A validation/explanation agent can review outputs, ensure consistency, and provide transparent reasoning to supply chain managers for greater trust in the system's recommendations.



























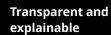
Autonomous supply chain operations

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Inaccurate recommendations can disrupt production or logistics. Al agents should be validated regularly against real-world outcomes and monitored to provide ongoing reliability.



Because supply chain managers need to understand how recommendations are generated, agents should be designed to provide clear explanations of their optimisation logic and how the underlying data was used.



Responsible and accountable

Given the risk of AI agents taking inappropriate or inconsistent action, humans should have the final responsibility for supply chain adjustments.

POTENTIAL BENEFITS

Greater supply chain resilience

Proactive detection of bottlenecks and disruptions – coupled with real-time adjustment – can minimise costly delays and maintain production and supply continuity.

Faster, data-driven decisions

Dynamic demand forecasting and optimisation helps supply chains respond quickly to shifting market demand and challenging operational conditions.

Autonomous warranty adjudication

Using AI agents to automate warranty claims processing

Agentic AI systems can streamline the adjudication of automotive warranty claims by using specialised agents to assess claim filings, flag anomalies, generate documentation, and support human adjudicators.

ISSUE/OPPORTUNITY

For automakers, warranty adjudication is a crucial function that directly affects costs, customer satisfaction, dealer relationships, and compliance. Today, the process often involves multiple handoffs, manual reviews, and inconsistent action. This hampers efficiency and speed, increases the risk of undetected fraud, damages relationships with customers and dealers, and drives up administrative costs.

Manual adjudication also makes it difficult to conduct a comprehensive and consistent review of claims.

Limited time and resources mean that only a subset of claims can be deeply reviewed. Potential errors, fraudulent claims, and incomplete filings could fall through the cracks. To mitigate such problems, automakers need a more efficient, scalable, and consistent way to manage warranty claims while preserving fairness and transparency.

HOW AI CAN HELP

Data validation and fraud detection

Al agents can review incoming claims for completeness, identify unusual patterns, and flag potential fraud, waste, or abuse before the claim progresses.

Customer and claims history analysis

Agents can cross-check claims with customer and vehicle history, uncovering relevant information that strengthens the adjudication process.

Documentation and reporting

An agent can automatically generate detailed reports for human adjudicators, reducing manual work and providing decision-makers with the necessary context to make informed decisions.

Decision support and denial drafting

Agents can propose denial reasons with clear justifications and then draft denial letters for human approval, helping adjudicators operate more efficiently and improving consistency across claims.



























Autonomous warranty adjudication

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Errors in root cause analysis or improper routing can lead to incorrect claim outcomes. Actions of Al agents should be validated against historical claim datasets and continuously monitored for performance accuracy.



explainable

Transparent and Warranty decisions can have a big impact on customers, dealers, and regulatory compliance. As such, AI agents should provide clear reasoning for why a claim was flagged, routed, or recommended for denial.



accountable

Responsible and Given the significant financial and reputational risks associated with warranty adjudication, final decisions should be left to human adjudicators, with Al agents providing decision support rather than operating as fully autonomous systems.

POTENTIAL BENEFITS

Increased efficiency and coverage

Automated warranty adjudication allows for broader and deeper claim analysis with less time and effort - improving both efficiency and accuracy.

Improved consistency and fairness

Standardised Al-supported processes reduce variability in claim outcomes, helping to make adjudications more consistent, transparent, and fair.

Al assistant for vehicle buying and leasing

Guiding consumers to the right car with personalised, multi-agent assistance

Agentic AI systems can streamline the vehicle buying and leasing process through specialised agents that evaluate numerous purchase options and provide hyperpersonalised recommendations to consumers.

ISSUE/OPPORTUNITY

Car buying and leasing is a complex and expensive decision that involves comparing different vehicle models, feature availability, financing structures, total cost of ownership, and dealer inventory. Consumers often find this process intimidating and confusing, which leads to delays, dissatisfaction, or switching to competing brands when their preferred model is unavailable.

For OEMs and dealers, missed sales opportunities during this critical decision window represent lost revenue and weakened customer loyalty. The challenge is compounded by lack of visibility into production pipelines or limited ability to match customer preferences with available inventory, leaving dealers struggling to balance consumer demand with real-world availability.

HOW AI CAN HELP

Personalised vehicle matching

A central vehicle search and advisor agent helps customers identify models that align with their preferences, budgets, and usage needs, whether buying, leasing, or exploring certified pre-owned (CPO) options.

Comprehensive financial analysis

A buy agent analyses total cost of ownership – including loan payments, depreciation, maintenance, and taxes – while a lease agent evaluates lease terms, monthly payments, and conditions to help customers compare financing options with full transparency.

Inventory and production visibility

An OEM agent analyses vehicles in production pipelines and offers booking options, helping OEMs capture demand even when current dealer inventory does not meet customer criteria.

Streamlined communication and support

A communication agent delivers supporting documents, sends summaries, and ensures smooth integration with dealership systems, reducing customer effort and follow-up calls.



























Al assistant for vehicle buying and leasing

MANAGING RISK AND PROMOTING TRUST



explainable

Transparent and Because purchase and lease decisions involve major financial commitments, agents need to provide clear explanations of cost breakdowns, assumptions, and trade-offs in their recommendations.



Robust and reliable

Errors in inventory matching or financial analysis can erode customer trust. Al agents should be validated against realworld dealership and OEM data and continuously updated to ensure accuracy.



accountable

Responsible and Al agents can have a significant influence on consumers' carbuying decisions. As such, their outputs should be positioned as guidance tools, with customers and dealer staff retaining final responsibility for understanding and confirming selections.

POTENTIAL BENEFITS

Increased sales conversion

By identifying various inventory and production pipeline options, OEMs can help reduce the number of customers they lose to competitors when the preferred model is unavailable.

Enhanced customer experience

Personalised recommendations and simplified comparisons improve decisionmaking and reduce the stress of navigating the complexities of financing and leasing.

Reduced dealer workload

Automated handling of routine inquiries reduces a dealership's call volume and required manual effort, allowing sales staff to focus on higher-value interactions with customers.

Marketing content assistant

Content generation

Al can be used to enable the creation of efficient, consistent, and personalised content across a range of modalities.

ISSUE/OPPORTUNITY

Companies face a significant challenge in managing and optimising marketing content. With hundreds of websites for brand portfolios, each in dozens of languages, companies struggle to allocate enough time and resources to create customer group-specific product descriptions, images, video, and even audio. Enterprises also wrestle

with consistency across descriptions, imagery, ads, and other media, and the materials may not always be optimised for the necessary purposes (e.g., product descriptions for search versus e-mail). Companies need a method to provide a seamless and personalised brand experience across different ecosystems and touchpoints.

HOW AI CAN HELP

Next-gen content generation

With AI, the enterprise can create product descriptions, imagery, video, and more much faster and more consistently than with existing tools and processes.

Personalisation at scale

Al models can draw from multimodal data (e.g., text, image, geospatial data) to create personalised and contextually relevant content. The model can be used to catalogue content and adapt content and user flow based on language, region, and customer behaviour trends.

Assisting compliance

Due to the consistency AI enables across modes, languages, and contextual factors, the enterprise can enhance regulatory compliance for materials across different geographies, cultures, and topics.















Marketing content assistant

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

While tasked with producing superior marketing materials, Al systems may invent inaccuracies, which will lead to poorer customer engagement and outcomes.



Fair and impartial

Biases in the data (e.g., due to incomplete datasets) could lead to unequal quality of content in the face of different geographical or cultural factors.









Catering to the customer

By tailoring content and the user experience based on language, region, and customer preferences, the enterprise can drive customer satisfaction and loyalty.



Personalised content can promote higher engagement, traffic, and conversions through tailored and relevant marketing experiences.

Cost efficiency

Using AI for content creation allows the enterprise to develop and maintain content at scale without the costs associated with commensurate human labour.









Planning for promotions

Reimagined trade promotions

Al can be used to prepare promotion plans, negotiation materials, pre-works, and pitch-decks.

ISSUE/OPPORTUNITY

When it comes to planning and negotiating trade promotions, Consumer Packaged Goods (CPG) organisations draw from a multitude of data sources and there is often not enough time to filter through all relevant information. What is needed is a way to more rapidly consult data sources to enhance trade pricing negotiations

by predicting outcomes, customising strategies, and tailoring selling stories. At the same time, there is also a challenge in understanding complex transactional data from retailers, which holds valuable insights for the design of successful promotion plans (i.e., what, where, and how to promote).

HOW AI CAN HELP

Supporting employees

Al can be used to prepare negotiation materials by combing through older campaigns or deals, sorting the relevant information, and generating suggestions. This helps equip the human employee with materials like pre-works (e.g., consolidated material from prior years) and pitch decks, supporting their negotiations.

Predicting outcomes

Al can help optimise trade shelf spacing and investment allocation by predicting outcomes and conducting scenario building and storytelling. It can also be used to build scenarios with cultural customisations for negotiation processes with retailers.

Optimisation support

With AI, users rapidly analyse EPOS data and transactional information to provide insights that help optimise the design of promotional programs, setting the right price points, promotion mechanics, and anticipating sales uplift to inform production processes of the expected demand.













The Al Dossier | Consumer

Planning for promotions

MANAGING RISK AND PROMOTING TRUST



Safe and secure

Because price, margin information, and negotiation strategies are consumed by the model, it must be secured to prevent the leakage of sensitive commercial data.



Fair and impartial

The data used to train and fuel the model may be dated, leaving new target groups and small but growing customer segments potentially underrepresented. As a result of this latent bias, the model may be challenged to provide commensurate accuracy for all groups and segments.









Driving efficiency

By using AI to augment preparing and sorting materials, the organisation promotes efficiency in trade promotion processes.



Leveraging AI can help improve allocation of resources across price, promotion, and negotiation strategies.

Data-driven decision-making

Using AI to create materials for trade negotiations enables human workers to access much more information and make more informed, data driven decisions.









Data access for all

Data-empowered business users

Al can help guide business users to key insights in consumer behaviours by enabling them to combine data from various sources through natural language queries, and by summarising issues to action without the help of dedicated analysts.

ISSUE/OPPORTUNITY

Everyone in the business should be consumer-focused, but while the marketing function may have access to customer data, business stakeholders in product design, trading, retail operations, supply chain, and other functions may only encounter slices of customer information. Currently, enterprises need dedicated analysts

to pull SQL queries and curate data for decision-making, which creates a barrier to customer information and insight. Data is held across different silos, and existing interfaces are only built to answer pre-populated questions. The result is that most business users cannot fully leverage the enterprise's models and data, and cross-functional

HOW AI CAN HELP

Greater access to insights

An AI system can help stakeholders across all business functions better understand the consumer by simplifying data mining and analysis with user friendly interfaces and natural language queries. This allows users to ask questions relevant to their work and extract actionable insights without compromising functionality.

Bringing down data barriers

The system can aggregate data from various sources and domains (e.g., purchasing patterns, customer service, website and browsing data, marketing campaign response) to provide comprehensive insights into consumer behaviours. Reaching across data silos, the system can automatically identify outliers and summarise issues to guide decision-makers to areas requiring attention.

























Data access for all

MANAGING RISK AND PROMOTING TRUST



Safe and secure The Al model is exposed to sensitive and proprietary enterprise data, which creates a risk of potential data leakage. To mitigate this risk, the enterprise may look at restricting data access to the Al provider, as well as carefully determining what consumer data should be exposed to the model.



Robust and reliable

For business users to make confident decisions informed by Al, they need to be able to trust the outputs. To this end, data inputs must be accurate and up to date, and outputs should be validated and monitored.



Transparent and explainable

Business users require sufficient context to interpret consumer data, and while analysis conducted by a data expert inherently contains a level of "human in the loop," when using an Al model, business users need the capacity to understand context and outputs.

POTENTIAL BENEFITS

Agile decision-making

Business users are empowered to make more informed decisions about product launches, sales, and other customer-related initiatives both quickly and efficiently.

Time and resource efficiency

Simplifying data access and analysis for business users can accelerate time to insight without additional burdens on data analysts and the technical workforce.

Seeing is believing

Virtual try-on

All can be used for style transferring, which allows consumers to see a digital rendering of clothes and other products on their own bodies, in their homes, and elsewhere.

ISSUE/OPPORTUNITY

In the clothing and make-up industry, consumers typically try on products to determine whether they want to purchase and keep it. Yet, this traditional method of selecting products is challenged by online shopping,

where the consumer relies on pictures and product descriptions to inform their decision. This can lead to high return rates and affiliated costs to the company, as well as customer dissatisfaction.

HOW AI CAN HELP

Accurate style transferring

By analysing images or videos of the customer and the desired style, AI can create realistic representations of how the clothing or product would look in the real world.

Virtual mix-and-match

Al allows customers to more easily explore a wider range of style options, clothing combinations, and accessories.

Greater personalisation

By considering factors such as body shape, skin tone, and personal style, AI can suggest suitable products that align with the customer's preferences.

























Seeing is believing

MANAGING RISK AND PROMOTING TRUST



Private

By working with and augmenting consumer photos and videos, the model is exposed to sensitive or personally identifiable information, which is subject to privacy regulations and standards. Leveraging Al for style transferring requires the enterprise to ensure user data is safely stored, transferred, and used.



Transparent and explainable

When consumers input an image of themselves or their surroundings, they need to understand how that media is used by the enterprise, how consumer-machine interactions are tracked and recorded, and whether there are any privacy risks to the consumer when using the style transferring application.



Fair and impartial

If the training set is unbalanced and therefore biased, renderings for virtual try-ons may be more accurate or realistic for one demographic group over another, potentially impacting customer satisfaction and regulatory compliance.

POTENTIAL BENEFITS

Customisations for the customer

Catering to the customer buying experience with a simpler way to explore product offerings promotes customer satisfaction.

Reduced return rates

When customers can better see and imagine how a product looks before making a purchase, it helps reduce the likelihood of mismatched expectations, product dissatisfaction, and returns.

Simpler sales

Making it easier to choose which product to buy by virtue of a simpler method for exploring options can support sales growth.

Trend analysis and insights

Al can be used to analyse data from virtual try-on experiences to gather insights on customer preferences, popular styles, and emerging trends.

Code assist for developers

Augmented developer

Al can be used to supplement the work of software developers by helping create and maintain multiple applications and platforms.

ISSUE/OPPORTUNITY

To give customers a seamless digital experience, enterprises are challenged to develop and maintain applications across different platforms. Yet, developers and other highly skilled professionals are in high demand and short supply. To overcome the talent

gap, AI can be used to supplement a developer's effort by automating aspects of code creation and maintenance so the developer can focus on more complex code writing and validating AI outputs.

HOW AI CAN HELP

Offloading lower-level work

Al can augment the completion of repetitive tasks, such as the deployment and maintenance of code across different platforms (e.g., iOS, Android, webapps).

A developer assistant

Al can be used in the development of the code itself, serving as an assistant supporting software developers in writing and maintaining code. It can also promote consistency across platforms and applications, such as by converting functional code to different environments.

























Code assist for developers

MANAGING RISK AND PROMOTING TRUST



Safe and secure

Code created with AI may include vulnerabilities that may be difficult to identify during development and even after deployment. Given the importance of cyber security, enterprises need to ensure generated code does not introduce security risks.



Robust and reliable

Al is susceptible to errors, and when using it for development tasks, human validation is necessary to mitigate the risk of bugs or vulnerabilities in code as it is created and maintained for multiple applications.

POTENTIAL BENEFITS

Efficient deployments

Using AI can help developers efficiently deploy and maintain code across platforms.

Digital consistency

Using AI helps developers maintain a consistent experience across multiple platforms by ensuring each environment functions at the same level of quality, thanks to automation (e.g., code conversion) that augments developer capacity and capabilities.



Customer support on demand

Customer assistant

Al-enabled virtual agents can improve the customer experience by providing realtime, personalised support and creating new ways of interacting with customers.

ISSUE/OPPORTUNITY

After purchase, customers may seek information or support around a product or service. While traditional call centres have implemented basic Al capabilities to automate responses to customer inquiries, the automation is often limited in its capacity to interpret

customer questions and respond in a conversational and helpful way. The need is to accurately and proactively respond to customer inquiries and online trends in an efficient and effective manner.

HOW AI CAN HELP

A conversational agent

Al can enable new ways of engaging with customers, using speech-to-text and natural language inputs to generate empathetic and personalised conversations for aftersales support and handling customer complaints.

Better use of human capital

Because generative AI can provide instant, personalised responses to customer queries, offer relevant solutions, and engage in conversations, customers can gain faster response and resolution, and organisations can free up human agents to focus on more complex customer issues.



























Customer support on demand

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

The quality and accuracy of customer interactions impact the customer experience and brand impression. If an Al-enabled customer assistant fails to provide accurate and personalised advice or product instructions, it could degrade (rather than enhance) the quality of the customer interaction.



explainable

Transparent and Customers should have the opportunity to gain a clear understanding of what the model can and cannot do. Also, to promote transparency and positive engagements, enterprises should set customer expectations for the virtual assistant.

POTENTIAL BENEFITS

Enhanced customer experience

Personalised and accurate support and troubleshooting contribute to a positive brand reputation and improve customer relationships and loyalty.

Increased efficiency

By using AI to automate various aspects of customer engagement, a higher volume of customer interactions can be accomplished simultaneously, improving response times and creating the capacity to scale with customer demand.



A virtual shopping assistant

Product recommendations

Al can be used to create personalised product recommendations based on customer preferences and behaviour.

ISSUE/OPPORTUNITY

Suggesting the right products to customers can increase sales, and hyper-personalised product recommendations are often the most effective at driving a sale. Data-based product recommendations are already possible today, but they often lack a conversational, natural language tone.

What is more, recommendations may lack a hyper-personalised quality as they are based on broader customer segments and purchase history, as opposed to individual customer search criteria and feedback.

HOW AI CAN HELP

Hyper-personalised recommendations

Based on customer input and preferences, Al can generate tailored recommendations, making the buying process more personalised and convenient. In addition, the interactive and iterative approach to product recommendations that Al enables can yield more targeted suggestions than current search engine capabilities.

Image as input/output

Consumers can enter an image of preferred styles (e.g., a celebrity in a designer outfit), and the Al model can identify products and make recommendations based on the image.



























A virtual shopping assistant

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

Latent bias in training and testing data may lead the model to express a preference toward some products or product combinations when making recommendations. Ongoing monitoring, data updates, and human validation can contribute to continuous improvement and bias mitigation.



Private

The model may be exposed to customer data throughout the course of an interaction, and that personal information may be subject to regulatory protections. Important considerations include how the customer data is stored, transferred, and used, as well as how the data is consumed and used by the model itself.

POTENTIAL BENEFITS

Enhanced customer experience

Delivering personalised and accurate support, guidance, and troubleshooting helps create a positive brand reputation and improves customer relationships and loyalty.

Increased efficiency

Using AI to automate selected customer engagement activities can improve efficiency and scalability while improving customer satisfaction.

Next-level market intelligence

Market research

By harnessing Al's capacity to read and summarise vast amounts of relevant material, companies can expedite market research and gain concise insights for effective decision-making in new markets.

ISSUE/OPPORTUNITY

When researching entry possibilities in new markets or customer groups and identifying new target segments, enterprises face a variety of challenges. Things like a lack of market data, unfamiliar customer preferences, cultural and economic differences,

competitive analysis difficulties, regulatory complexities, high market entry costs, potential brand perception challenges, and uncertainties about demand and market acceptance all impact the speed and quality of market research.

HOW AI CAN HELP

Market intelligence

Al can help simulate market scenarios, generate synthetic data to fill data gaps, predict customer preferences based on existing patterns, offer cross cultural insights, aid in competitor analysis, suggest compliance strategies, optimise market entry costs, simulate brand perception scenarios, and provide demand forecasting to reduce uncertainties.

Information synthesis

Al enables rapid market research by efficiently reading and summarising extensive volumes of pertinent material, presenting the information in a readily understandable format for market research teams.

Novel market segmentation

Al generated data may reveal new and previously unidentified market segments within the target market. This can open up additional opportunities for niche marketing and product customisations.

Richer personas

Rather than relying on basic surveys and focus groups for understanding consumer likes and dislikes, AI can identify specific customer preferences and create detailed profiles. Using AI, market research teams can even create fictional yet plausible customer personas based on the market's unique characteristics, helping the company better understand their potential customers' behaviour and preferences.



























Next-level market intelligence

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

Al models may learn from biased datasets, leading to biased outputs that do not accurately represent the actual market.



Robust and reliable

Given Al's potential to hallucinate and produce inaccurate outputs, Al-generated insights should be verified with real-world data and traditional research methods to ensure accuracy and reliability.



Responsible and accountable

While AI can complement market research, it should not replace traditional research entirely, as it may miss qualitative nuances and human expertise.



Transparent and explainable

To trust Al outputs, users require the ability to understand which samples and research methods were used to generate recommendations and insights.

POTENTIAL BENEFITS

Cost-effective research

Al can reduce the costs associated with traditional market research methods by generating large datasets and simulating scenarios.

Risk mitigation

By simulating market responses, CPG companies can identify potential risks and challenges in the new market before making substantial investments. This helps reduce the chances of product failure and financial losses.

Integrated business planning

Al consolidation of forecasting and planning across the enterprise

Al can help an organisation consolidate real-time sales, demand, and supply data across all functions, creating a single source of truth to drive faster, more strategic decisions in finance, supply chain, marketing, and sales.

ISSUE/OPPORTUNITY

Today's companies have a wide variety of systems for planning and forecasting. However, the individual outputs from those disparate systems often conflict with each other and don't provide a unified view of what's really going on. Different teams – finance, supply chain, marketing, and sales

- create their own forecasts using siloed data and inconsistent approaches. The potential results? Mismatched projections, inefficiencies, delayed decision-making, and significant operational waste.

HOW AI CAN HELP

Real-time consolidation

Al can consolidate real time inputs from sales, inventory, marketing trends, and supply chain metrics to produce dynamic forecasts.

Sophisticated analysis

Al enables trend recognition, historical pattern analysis, and early alerting on supply demand gaps, while also facilitating scenario planning and pricing strategy refinement all through a unified dashboard.

Actionable insights

The system can continuously update itself as new data flows in, signalling demand shifts or regional product affinities and providing decision-makers with actionable insights.



























Integrated business planning

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Given the system's critical business impact, resilience is key. Al models should be capable of updating in real time and integrating multiple data streams reliably and accurately. Extensive pilot testing can help fine-tune model accuracy before scaling.



Transparent and explainable

A user-facing dashboard that clearly shows inputs, trends, and recommendations can help business leaders understand how forecasts are generated, what assumptions are at play, and what real-world data is influencing outputs – reducing blind reliance on the system and promoting human-Al collaboration.



Safe and secure

To help mitigate security breaches and operational disruptions, robust security protocols should be embedded in both the technology infrastructure and data flows, with IT overseeing access controls, data encryption, and integration with existing ERP systems.

POTENTIAL BENEFITS

Unified forecasting with less redundancy

Al can help minimise conflicting forecasts across departments, creating a single source of truth for the entire enterprise.

Improved collaboration and decision-making

Cross-functional teams are able to operate from the same real-time data set, improving alignment. Also, leaders spend less time on data consolidation and cross-checking, enabling them to make better-informed decisions more quickly.

Greater supply chain efficiency

Integrated business planning powered by Al enables better inventory and warehouse management, which can reduce supply disruptions, shortages, and waste.

Social media content generation

Automated, multimodal content creation that is trend-aware and always on

Al is now being used to autonomously produce social media content – text, images, hashtags, and videos – that aligns with brand identity and capitalises on viral trends in real time.

ISSUE/OPPORTUNITY

Social media is a key channel for communicating with customers and shaping brand perceptions, and an important driver for awareness, engagement, and sales conversion. But creating personalised, high-quality content at speed and scale – while maintaining brand consistency and legal compliance – is a difficult balance.

Large enterprises often rely on global agencies to support social media content across dozens of brands and channels. This approach can be very costly, timeconsuming, and limited by human working hours. Also, in a media environment where trends can shift in an instant, traditional methods likely cannot scale or respond quickly enough to keep pace with opportunities in real time.

HOW AI CAN HELP

Detecting and analysing trends and events

Al can help detect and analyse influencer trends and brand affinity across a wide range of social media platforms 24/7, identifying opportunities to shape consumer expectations in real time. Retrieval augmented generation (RAG) capabilities provide real time access to social data, such as trending hashtags, viral video clips, and current events.

Model-agnostic orchestration

Content creation tasks can be dynamically routed to the most cost effective or best performing Al models, optimising output while reducing compute costs.

Generating multimodal creative content

Al offers the ability to autonomously generate creative content across modalities, while remaining contextually and culturally aware. Key capabilities include: (1) LLMs for generating social media copy, product descriptions, captions, and hashtags; (2) multimodal image models for visual asset generation, including pack shots, brand imagery, and marketing visuals; and (3) short form video generation.













Social media content generation

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Retrieval-augmented generation can reduce hallucinations and improve model performance over time. Fallback models and safety nets can mitigate failures or inappropriate content generation under unpredictable conditions.



Fair and impartial

The content generation pipeline should be evaluated regularly for potential cultural, social, or representational biases. Human oversight can ensure that outputs reflect brand values.



Private

No personal user data should be used in the generation process; models should be trained and tuned on anonymised or public datasets. Data residency and usage should comply with regional regulations, including the EU Al Act.

POTENTIAL BENEFITS

Always-on, real-time responsiveness

Traditional content workflows often require long lead times for ideation, approval, and execution. With AI, brands can respond almost instantly to real-time events, seasonal trends, or cultural moments by generating relevant content in minutes, enabling more agile and timely brand engagement.

Scalable content production at low marginal cost

Al enables brands to produce high volumes of personalised, platform-specific content – text, images, and video – without requiring a linear increase in resources. The system can support hundreds of product lines and campaigns with minimal incremental effort, greatly improving operational leverage.

Cost reduction through automation and budget reallocation

Augmenting external creative agencies and internal content teams with Al-generated outputs can reduce operational costs. It can also free up budget that can be reallocated toward more strategic initiatives, such as paid media, analytics, or customer experience improvement.

Data-driven personalisation and targeting

Al systems can tailor content for audience segments based on behaviour, geography, platform norms, or product affinity. This micro-personalisation allows brands to deliver relevant content to niche audiences, increasing engagement and conversion potential.



By automating repetitive or timeintensive content generation tasks, marketing and creative professionals can focus more on high-value work such as strategy, brand storytelling, or campaign optimisation. This reallocation of effort can lead to improved job satisfaction and better use of talent.

and visual identity With proper tuning an

Consistent brand voice

With proper tuning and governance, Al-generated content can more reliably align with predefined brand guidelines, helping to ensure a unified voice across markets, languages, and touchpoints. The system can learn and reinforce tone, terminology, and aesthetic standards consistently.















The Energy, Resources & Industrials AI Dossier

















The Energy, Resources & Industrials Al Dossier

Al is emerging as a critical enabler of transformation across the energy, resources, and industrial sectors. In industries defined by asset intensity, operational complexity, and margin pressure, Al offers the ability to sense, predict, and act with speed and precision. From optimising energy production and forecasting demand to predicting equipment failures and supporting engineers in the field, Al is helping organisations operate more efficiently, safely, and sustainably.

Recent advances – particularly in industrial-grade machine learning, computer vision, and autonomous control systems – are expanding the frontiers of what's possible. Al agents can now manage assets and oversee field operations with minimal human input. Meanwhile, generative models are being used to simulate physical systems, accelerate engineering design, and fine-tune operational planning.

These capabilities are arriving at a critical moment. Energy, resources, and industrials companies face mounting pressure to decarbonise, digitise, and build resilience in the face of global volatility. Al can support these goals by unlocking more value from existing assets, enabling predictive and

prescriptive analytics, and surfacing hidden efficiencies in large-scale operations. Integrating real-time data from sensors, satellites, and enterprise systems, Al can offer decision-makers a more holistic, adaptive view of their operations.

However, realising this potential requires more than model development. It demands robust infrastructure, workforce upskilling, data governance, and alignment with regulatory frameworks – particularly as Al applications begin to intersect with safety-critical systems. Companies that take a disciplined, value-focused approach to Al adoption are already seeing meaningful returns, from reduced downtime and energy waste to faster innovation cycles.

As capabilities mature, Al is set to become not just a tool for operational excellence in the energy, resources, and industrials sectors, but a cornerstone of competitiveness in an increasingly automated and resource-constrained world.

In industries defined

by asset intensity,

margin pressure, Al

offers the ability to

sense, predict, and

act with speed and

complexity, and

operational

precision.

Note: The tags below each use case indicate its primary business function and whether Agentic AI is used.

Primary business function

Agentic Al



Al-driven predictive maintenance

Avoiding downtime through autonomous, multi-agent diagnosis and intervention

Agentic AI systems can monitor industrial equipment health, anticipate failures, diagnose root causes, and proactively schedule maintenance – keeping equipment running smoothly and reducing maintenance costs.

ISSUE/OPPORTUNITY

Industrial equipment failure can trigger substantial costs; yet, traditional maintenance is often reactive or rigidly scheduled, leading to unexpected breakdowns or wasteful overmaintenance. Critical factors such as

labour shortages, weather events, and the move toward electrification are all increasing the need for Al-powered predictive maintenance of industrial equipment and more dynamic, datadriven asset maintenance planning.

HOW AI CAN HELP

Continuous sensor monitoring and anomaly detection

Al agents can analyse vibration, temperature, pressure, and other IoT sensor data in real time, leveraging predictive analytics to flag deviations from baseline performance and predicting impending failures.

Root cause analysis and diagnosis

When anomalies are detected, specialised agents can assess historical failure logs, maintenance records, and environmental conditions to pinpoint likely failure modes.

Automated work order generation and scheduling

Other agents can generate detailed work orders and schedule tasks based on production cycles, resource availability, and cost constraints.

Simulation and reinforcement learning

Multi-agent reinforcement learning systems can simulate inspection intervals and failure scenarios to reduce maintenance expenses and downtime.

Human-centric integration and continuous improvement

Al agents can collaborate with human maintenance teams: presenting findings in clear, natural language, validating outcomes, helping to prioritise alerts and recommend next steps, and refining models over time based on new data and outcomes.

























Al-driven predictive maintenance

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Predictive accuracy should be validated across varied asset types and environmental conditions. Also, Al agents should be tested against historical failure cases and simulated breakdown scenarios, with fall back mechanisms for human review in uncertain situations.



Transparent and explainable

Explainable AI outputs improve adoption and help build trust. Agents need to provide transparent reasoning for their recommendations and actions (e.g., "Vibration on bearing exceeds historical threshold during peak load," or "Leaf spring failure consistent with past incidents"), supported by traceable data sources.



Safe and secure

Industrial systems are vulnerable to cyber threats, and the consequences of a breach can be severe. Agent platforms should include intrusion detection, secure communication with edge devices, and resilience against malicious sensor tampering or spoofing.



Responsible and accountable

Although Al agents can provide valuable decision support, ultimately human technicians and maintenance managers are responsible for critical decisions and actions. As such, clear escalation protocols need to be established for ambiguous or high-risk alerts.

POTENTIAL BENEFITS

Less unplanned downtime

Real-time, automated detection can trigger repairs early when needed, turning potential disruptions into planned maintenance and reducing productivity losses.

Lower maintenance costs

Focusing on condition-based needs can reduce unnecessary maintenance, minimising spare-part inventory and technician labour.

Extended asset lifespan and operational efficiency

Continuously monitoring the condition of industrial equipment enables more precise upkeep and longer service life. Data-driven insights improve scheduling and reduce waste, boosting overall productivity and sustainability.

Autonomous drone-based infrastructure inspection

Conducting unmanned, Al-guided inspections of physical assets

Autonomous drones, guided or enhanced by agentic AI, can inspect physical infrastructure such as power lines, pipelines, and transmission towers - capturing and analysing quality imagery at lower cost and risk.

ISSUE/OPPORTUNITY

Traditional infrastructure inspections in energy, mining, utilities, and industrial environments typically rely on human teams using scaffolding, helicopters, or ropes. In addition to being costly, dangerous, and slow, these manual inspection methods are often disruptive to ongoing operations. What's more, they can miss subtle problems such as hidden defects and weather-related damage.

Autonomous drone systems, especially docked or "drone in a box" variants. enable frequent, hands free inspections that reduce costs, minimise risk to human personnel, and improve overall quality and speed.

HOW AI CAN HELP

Automated flight and mission control

Drones launch autonomously from preprogrammed docks, fly designated routes, capture high-resolution imagery, and then return for charging and data offload - all without human pilots.

Al-driven defect detection

Onboard analytics can automatically process visual, thermal, or LiDAR data to highlight problems such as cracks, corrosion, vegetation encroachment, and structural anomalies.

Agentic inspection orchestration

An orchestration agent can oversee fleets of drones, schedule inspection missions dynamically, monitor inspection results in near-real time, and then automatically trigger follow-up maintenance workflows when anomalies are detected. This greatly reduces the need for manual coordination.

Extended-range and precision navigation

Advanced AI systems now support beyond-visual-lineof-sight (BVLOS) operations over long distances using perception-aware controllers to maintain accurate positioning and avoid obstacles.













Field Services

Agentic Al













POTENTIAL BENEFITS

Safety improvements

Human workers no longer need to climb towers, pilot helicopters, or enter hazardous zones for routine inspections, reducing the risk of injuries.

Cost and operational efficiency

Autonomous drones reduce inspection time and logistics overhead, significantly lowering costs and enabling more frequent monitoring without disrupting operations.

Proactive maintenance and uptime

Higher-frequency missions and more consistent data allow earlier detection of issues. This enables maintenance to be done before failure occurs, reducing unplanned downtime and enhancing asset reliability.

Autonomous drone-based infrastructure inspection

MANAGING RISK AND PROMOTING TRUST



Inaccurate or inadequate inspections can lead to safety risks and failures. Drone systems need to be validated against varied situations and adverse environments, not just routine test missions and historical defect cases. Also, escalation paths should be established for manual review of ambiguous findings.



Transparent and explainable

Al should provide tangible evidence (such as high-resolution images, heat maps, or annotated photos) accompanied by clear explanations that enable maintenance teams to confidently verify issues.



Private

Some drone inspections may inadvertently capture sensitive data. Agents need to enforce data minimisation and restrictedaccess policies to protect private or proprietary information during data collection and transfer.



Safe and secure

Since drones operate near power lines and potentially dangerous industrial equipment, systems must include breachresistant communication, secure command chains, and rigorous safeguards against signal spoofing or unauthorised control.















Autonomous field operations management

Improving safety and efficiency by using AI agents to support field operations

Agentic AI systems can improve field operations by managing task coordination and automating frontline decision-making, enabling field workers to focus on complex, high-value activities rather than routine logistics.

ISSUE/OPPORTUNITY

Field operations such as utility maintenance, site monitoring, asset inspection, and emergency response are typically labour-intensive and highly fragmented. Workers in remote locations often face incomplete data, manual task assignments, shifting priorities, and

limited coordination. These challenges can slow response times, increase safety risks, and make it harder to scale up operations. Al agents can help address such issues by autonomously handling coordination, data collection, and routine decision-making.

HOW AI CAN HELP

Task management and response automation

Al agents can identify operational issues such as inspection gaps, maintenance alerts, and compliance breaches. while also providing real-time support to field teams, bridging the gap between manual and autonomous operations. They can generate task lists, assign them to the right field crews, and follow up on execution, reducing coordination delays and helping to avoid missed maintenance cycles.

Context-aware adaptation

Agents can respond in real time to changes such as weather shifts, unexpected hazards, and shifting regulatory requirements. They can also reschedule tasks, redirect resources, and escalate safety alerts as needed.

Specialist collaboration

Multi-agent Al teams can mimic expert field units. For example, a maintenance agent can diagnose faults; a scheduling agent can reroute technicians; and a compliance agent can confirm actions comply with regulatory requirements. By working together, Al agents can provide specialised yet cohesive action.

Continuous learning and feedback

Agents can learn from experiences such as safety incidents, unexpected site constraints, and real-world repair times. This feedback loop improves their task allocation and decision logic, boosting long-term effectiveness and responsiveness.













Autonomous field operations management

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Field operations take place in unpredictable environments. Agents should be tested against variable and adverse conditions, such as network outages and natural disasters. Fallback paths should include human intervention for ambiguous cases.



Transparent and explainable

Technicians, supervisors, and regulators need to understand Al decisions. Agents should provide clear rationales such as "rescheduled inspection due to stage-1 weather alert," backed by task logs and field data.



Safe and secure

Safety in field operations is of paramount importance. As such, agentic systems must include cyber security safeguards such as encrypted communications, device authentication, anomaly detection, and protection against unauthorised commands.



Responsible and accountable

Governance structures should define escalation paths, signoff mechanisms, and oversight protocols for Al-generated plans. Although agents can provide useful support, ultimate responsibility for critical decisions and actions should remain with human field supervisors and operations managers.

POTENTIAL BENEFITS

Faster field response

Al agents can close the gap between detection and action by assigning tasks and dispatching crews within minutes of an alert, eliminating unnecessary delays.

Higher operational uptime

Proactive task assignment and adaptive scheduling can help ensure that critical infrastructure remains up and running, minimising downtime and safety disruptions.

Scalable field capacity with fewer resources

Al agents can help smaller teams effectively manage larger service areas, expanding operational reach without requiring proportional staffing increases.



Intelligent commercial operations

Driving smarter bidding, pricing, and customer engagement with AI agents

Al can streamline and accelerate commercial operations by using agents to automate contract pricing, bid preparation, demand forecasting, and customer engagement workflows.

ISSUE/OPPORTUNITY

Commercial teams in the energy, resources and industrial sector face complex pricing challenges across multiple product lines, geographic regions, and regulatory regimes. Traditional methods are labourintensive and revolve around spreadsheet models and fragmented data sources. This highly manual approach slows down responses to competitive pressure and limits the ability to grow.

Businesses in the sector need systems that can accelerate commercial workflows, integrate real-time signals, and personalise offers to customers at scale. Multi-agent Al offers a potential solution, providing a path to more flexible bidding and pricing through automated analysis and continuous adaptation.

HOW AI CAN HELP

Pricing and bid preparation agents

A pricing agent analyses input costs, market trends, competitor rates, and regulatory constraints. A bid preparation agent then drafts customised proposals based on customer histories, relevant contract terms, and compliance guidelines.

Demand forecasting and scenario modelling

A demand agent pulls in realtime signals – such as weather forecasts, energy prices, and consumption patterns – then simulates demand under different conditions to inform pricing and bid decisions.

Customer engagement coordination

An orchestration agent coordinates the bid, pricing, and demand agents to ensure cohesive action. It sequences tasks, manages version control, and forwards AI output to human commercial teams for review and execution.

Contract rollout and monitoring

After the sale, specialised agents monitor contract performance and market deviations, flagging margin erosion that needs to be addressed and triggering renewals when favourable conditions arise.

























POTENTIAL BENEFITS

Faster bidding cycles at lower cost

Al automation can streamline and accelerate bid preparation, reducing costs and improving responsiveness and competitiveness.

Higher margin capture

Dynamic pricing based on real-time data and scenario modelling helps maximise margins.

Scalability

Al agents help commercial teams efficiently manage larger customer portfolios, enabling faster business growth without proportional staffing increases.

Better strategic alignment

By freeing commercial teams from routine administrative tasks, AI agents allow people to focus more attention on strategic market expansion, negotiations, and relationship building.

Intelligent commercial operations

MANAGING RISK AND PROMOTING TRUST



Pricing errors can reduce margins and trigger regulatory issues. Agents must be tested on historical scenarios and stress-tested for market shocks.



Fair and impartial

Agents must avoid bias toward certain customers or regions. Pricing models should be audited to ensure fairness, and human commercial teams should have the ability to override automated suggestions.



Safe and secure

Given the sensitivity of pricing data and customer value, agentic systems must include secure authentication, intrusion detection, and validation of external inputs to guard against manipulation or cyber interference.



accountable

Responsible and Escalation protocols should be in place for high-value or sensitive proposals, with humans retaining final responsibility for commercial offers and contract decisions.



Materials design

Al empowers materials designers to explore a wider design space, optimise material properties, and expedite the discovery of new materials.

ISSUE/OPPORTUNITY

Developing new materials is challenging, costly, and time-consuming. One reason is that the chemical space is vast and complex while the number of chemically feasible molecules is unknown. Also, the

materials discovery, development, and optimisation process present different complexities at each stage, increasing the time required to reach a final design.

HOW AI CAN HELP

Streamline experimental process

Using AI to determine the most efficient experimental procedures for probing or optimising materials can streamline the experimental stages of development by removing redundant experiments and undertaking those that are cost- and time-optimised.

High-entropy alloy (HEA) engineering

Traditional techniques for developing HEAs with excellent physical, chemical, and mechanical properties are time-consuming and costly, making AI modelling a promising alternative development pathway.



























Expediting experiments and design

MANAGING RISK AND PROMOTING TRUST



Safe and secure Intellectual property or a similar competitive advantage could be compromised by using AI in materials design, as models trained on proprietary or sensitive data could potentially reveal valuable insights or design strategies to competitors.



Responsible and accountable

Companies should be mindful to identify and mitigate unintended negative ramifications of materials designed with the support of AI, such as long-term environmental impacts from materials that cannot be manufactured in responsible and sustainable ways.

POTENTIAL BENEFITS

Fuelling innovation

Al applications have the capability to rapidly generate and prioritise a wide range of virtual materials with diverse compositions and structures. This virtual screening process allows researchers to identify potential candidates for specific applications or material properties much more quickly than traditional experimental methods.

Bringing down costs

Through efficiency savings and the rationalisation and/or elimination of experiment consumables, the organisation can reduce development costs.

Enabling discovery

Al maximises the likelihood of discovering materials with superior properties by leveraging its ability to efficiently explore and navigate a vast design space of potential materials.

Understanding the ore

Minerals processing optimisation

Al can make the process of chemical separation of minerals from ore more cost- and time-efficient, safer, and more environmentally sustainable.

ISSUE/OPPORTUNITY

In mineral processing, chemical additives must be matched to the exact contents of the ore to separate as much material as possible from waste minerals without destroying it. The process is complicated due to the fact that modelling and testing each compound is time-and

effort-intensive, complex mineralogy and interrelationships between minerals can hinder recovery, and environmentally hazardous chemicals are often necessary to process certain compounds.

HOW AI CAN HELP

Ore characterisation and mapping

Al models can be trained on large datasets of mineral samples to generate synthetic samples that mimic the characteristics of real-world ores. Comprehensive databases can be built for mineral identification, classification, and prediction of ore properties, permitting insights into the behaviour and composition of different ores without testing on known processing assays.

Process optimisation

Models that simulate the physical and chemical processes involved in mineral processing can help optimise factors like grinding parameters, flotation conditions, and separation techniques. This can improve efficiency, reduce energy consumption, and enhance mineral recovery rates.

























Understanding the ore

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Al models may struggle to generalise mineral samples and processing scenarios that are significantly different from the training data. The model might not capture the full range of variations and unique characteristics of novel ores, which could lead to suboptimal processing recommendations. Also, if Al models cannot interpret complicated physical and chemical qualities like particle size distribution, mineral composition, and processing conditions, the model may generate suboptimal strategies or overlook critical factors.

POTENTIAL BENEFITS

Accelerated exploration

The cost and time needed to characterise ore and develop a processing workflow can be significantly reduced, and cost and efficiency trade-offs can be optimised to maximise mineral recovery while minimising operational costs.

Eco-friendly operations

Keener insights into mineralogy using Al can help reduce the amount of environmentally damaging additives and resources needed for processing without sacrificing production volume or efficiency.

Occupational health

Optimised processing can help reduce human exposure to toxic chemical additives and fine particle dust, which contributes to a safer work environment.

Optimise the design

Site design generation

Al can support the development of site plans by automating aspects of the design process, providing designers with new possibilities and reducing the associated time and cost.

ISSUE/OPPORTUNITY

Site planning is a multi-stage, iterative process to optimise cost, efficiency, and safety, but it is also an expensive and time-consuming exercise involving numerous stakeholders and third-party specialists. Site planning can require surveys in remote, sometimes hostile

locations. Forecasting near- and longterm impacts involves assessing a multitude of factors, and site-specific activities such as topological and geological surveying can be labour intensive and expensive.

HOW AI CAN HELP

Automated layout generation

Designers can use AI to analyse site constraints, design requirements, and input from engineers to quickly generate layout options for site plans that consider factors such as zoning regulations, operational use, and user preferences.

Design optimisation

Al can help optimise site plans by analysing parameters like solar orientation, traffic flow, and accessibility to suggest optimal infrastructure placements. This can help improve energy efficiency, support better space utilisation, and enhance the user experience.

Efficient documentation

and annotation

By analysing design elements and structures in the generated plans, AI can automatically annotate the plans with relevant information, such as dimensions, materials, and specifications. This automation could save designers considerable time and effort, allowing them to focus on higher level design tasks.

























Optimise the design

MANAGING RISK AND PROMOTING TRUST



Responsible and accountable

Al for design optimisation may focus primarily on efficiencies, such as cost reduction or time savings, while potentially neglecting other important considerations, such as environmental sustainability, community impact, or long-term adaptability. The model should be configured to balance multiple objectives and prioritise trade-offs to achieve better overall outcomes. Using Al for site planning also raises legal considerations around intellectual property, ownership of Al-generated designs, liability for design flaws, and privacy restrictions for sensitive or proprietary data.

POTENTIAL BENEFITS

Acceleration with automation

Using AI for site planning can accelerate the completion of time-consuming processes.

Discovering new solutions

With AI quickly creating a variety of site designs, the planning process can include a greater diversity of designs and promote innovative planning solutions.

Reducing risk

Al can simulate and analyse potential hazards and safety risks in site plans. Algenerated planning would consider factors such as weather events, traffic patterns, and emergency response routes. It could propose alternative design options to proactively minimise risks to safety and reduce potential property damage in case of unforeseen events.

Enhancing employee safety

Personalised OHS training

Al can be used to develop personalised and immersive occupational health and safety (OHS) training materials that allow trainees to be safely exposed to realistic scenarios and thereby reduce or better respond to real OHS incidents.

ISSUE/OPPORTUNITY

Traditional OHS training may only cover some potential scenarios, and it lacks practical opportunities to apply new skills and knowledge. Workers need to be prepared for emergency scenarios but cannot practice managing these scenarios in a real-world setting due to the cost and risk involved.

HOW AI CAN HELP

Virtual reality (VR) training

Combined with VR, AI can be used to develop virtual training environments that replicate operational conditions. With realistic scenarios that simulate OHS incidents, trainees can navigate hazardous situations, identify risks, and improve their OHS awareness and response capabilities in a safe setting.

Customised training content

Al can be used to customise training materials based on specific job roles, site conditions, or regulatory requirements. This technology can analyse large volumes of data, such as incident reports, OHS guidelines, or compliance standards and generate tailored content, including videos, interactive modules, or quizzes.

























Enhancing employee safety

MANAGING RISK AND PROMOTING TRUST



Safe and secure

Real-life emergencies can be highly stressful and traumatic, and replicating these scenarios virtually could imperil the psychological safety of trainees. The final design of simulations should be reviewed by human trainers to remove potentially harmful visualisations.



accountable

Responsible and Al-generated training materials should be continuously monitored to identify any potential issues, inaccuracies, or outdated information. Regular updates to the training content should be made to reflect the latest safety guidelines, regulations, and best practices.



Fair and mpartial

Al-generated training materials should be designed to be accessible to all types of learners, including individuals with disabilities. Organisations should consider providing closed captions for videos, adjustable training scenarios to accommodate different skill levels, and alternative formats

POTENTIAL BENEFITS

Safety through preparedness

Increased training engagement and readiness for emergencies supports workforce safety and fewer OHS incidents.

Customised training

A personalised approach to OHS training helps address the specific needs of workers, ensuring they receive relevant and targeted instruction.

Dynamic compliance

Changes in legislation, regulation, and policies can be quickly reflected in training materials by using AI to make updates.

Peering below the surface

Hydrocarbon reservoir exploration

Al can be used to optimise exploration success rates, reduce costs, and mitigate risks associated with hydrocarbon reservoir location and characterisation.

ISSUE/OPPORTUNITY

Oil and gas exploration involves a high degree of uncertainty and risk. Advanced technologies and extensive data analysis are needed to navigate the subsurface and accurately locate and characterise reservoirs. Extracting oil and gas from underground reservoirs requires advanced drilling techniques and technologies,

and harsh environmental conditions. deep water, and complex logistics make offshore exploration difficult. This makes exploration a capital-intensive and timeconsuming process involving multiple stages of seismic surveys, analysis, drilling, and testing.

HOW AI CAN HELP

Seismic data analysis

To overcome incomplete, low volume, or poor-quality seismic data, Al can support enhanced data analysis and interpretation. Al can be used to generate new data samples that resemble the patterns and characteristics of the existing seismic data, address missing or incomplete seismic data, improve data quality through denoising or resolution enhancement, and more effectively interpret complex data patterns.

Reservoir characterisation

By analysing data sources such as well logs, core samples, and production data, Al can create models that simulate the more complete behaviours of hydrocarbon reservoirs. This enables a better understanding of the reservoir dynamics, which helps optimise production strategies and improve recovery rates.

























Peering below the surface

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

False positives or misinterpretations may result in costly and time-consuming drilling operations that do not yield productive reservoirs, making human expertise crucial to validating insights and decision-making. Also, Al models could overlook critical factors or geological nuances that human geoscientists would recognise and could fail to contextualise the data when generating outputs. Without a contextual understanding, Al models and interpretations may lack accuracy or fail to capture the full complexity of reservoirs.

POTENTIAL BENEFITS

Informed investments and decisions

A deeper, more complete understanding of the characteristics of hydrocarbon reservoirs reduces uncertainty and enables better investment decisions.

Amplifying exploration

Improved data quality supports more accurate subsurface modelling, imaging, and structure characterisation, which leads to an increased ability to accurately locate hydrocarbon reservoirs.

Smarter strategy

With an earlier and more complete understanding of reservoir characteristics, less time is needed to optimise production strategies.











A smart eye in the sky

Smart summaries for drone surveying

Al can assist in summarising large volumes of drone footage and enable querying to enhance productivity and efficiency.

ISSUE/OPPORTUNITY

In the mining industry, drones are increasingly used for tasks such as mapping, management of dam tailings, safety management, blast assessment, environmental monitoring, and haul road optimisation. In the case of Optical Gas Imaging (OGI) to detect gases and volatile organic compounds leaking from vessels (e.g., pipelines), unmanned drones mounted with OGI cameras have proven useful for surveying a variety of equipment

over vast areas. Using drones in this way permits frequent scans and reduced costs associated with fugitive gases. Yet, while advanced AI solutions (e.g., volumetric monitoring) have been developed for applications using drone footage, manual inspection of drone footage is still required for environmental monitoring, security review, safety assessment, and retrospective analysis.

HOW AI CAN HELP

Smart summaries

Combined with computer vision solutions, Al can create smart assistive summaries in natural language from thousands of hours of drone footage. Assistive smart summaries can be based on a pre-determined template requested by the user, where observations are generated about elevations, topology, lighting, vegetation, and other factors. Summaries can also be queried in natural language so questions can be asked without the assessor manually reviewing all footage.

Querying the footage

When using OGI to detect leaks, there may be instances where a leak is irreparable but still must be managed. With AI, specific sites can be efficiently reviewed and monitored using simple natural language queries.













A smart eye in the sky

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Al models may struggle to interpret environmental indicators, assess ecological impacts, or consider local conditions and regulations. Training data availability and quality in particular can impact the Al model's ability to generalise and handle diverse environmental scenarios. Inadequate or biased training data may result in limited or skewed analysis and summaries.



Private

Drone footage may contain sensitive information, including personally identifiable information, facial images, or confidential business information, and the footage may also be captured on private properties or areas with restricted access. In using Al to analyse and summarise the footage, unsecure data handling and access can raise privacy concerns as well as legal and regulatory implications.

POTENTIAL BENEFITS

Supplementing human expertise

Querying smart assistive summaries helps ensure critical observations are not missed due to human error or cost and time constraints.

Faster time to insight

Replacing manual drone footage inspection with assistive summaries saves significant time and effort.













Resilient logistics and planning

Supply chain optimisation

Al can support supply chain optimisation by leveraging its ability to simulate, model, and generate data-driven insights.

ISSUE/OPPORTUNITY

Global supply chains are highly interconnected with many dependencies and multiple stakeholders. The inherent complexity creates challenges to efficiency, resilience, and cost avoidance,

making supply chain intelligence a critical component of supply chain management. What is needed is a way to rapidly analyse data from internal and external sources to identify patterns and areas for improvement.

HOW AI CAN HELP

Supply chain intelligence

Al can help identify and simulate potential disruptions or risks in the supply chain. By assessing port congestion, shipment routes, and tier-n supplier mapping, Al can be used to predict risks and their corresponding impact on operations, then recommend actions to mitigate those risks. This allows supply chain managers to proactively implement mitigation strategies, develop contingency plans, and improve overall resilience.

Scenario analysis and optimisation

Supply chain managers can use Al to run what-if scenarios in a digital twin environment that reflects the real-world supply chain. By simulating the impact of changes in demand patterns, production capacity, inventory strategies or supplier reliability, supply chain managers can improve risk assessments and proactive decision-making based on real-time conditions.

Supply chain planning

Al enables supply chain professionals to use natural language to interact with advanced planning solutions. Questions concerning all supply chain areas, such as planning, inventory, supply assurance, order management, and global logistics, can be easily asked, helping even less experienced users navigate complex topics and data.

Supplier assessment

Al can assist in supplier evaluation and relationship management by analysing financial reports, performance metrics, customer feedback, and other data and then generate insights and predictions around supplier performance, risk factors, and opportunities for collaboration. This helps supply chain professionals make informed decisions when selecting, negotiating with, and managing suppliers.















Resilient logistics and planning

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Supply chain management involves complex trade-offs, strategic considerations, and tacit knowledge that AI models may not fully capture. AI outputs may also fail to balance ethical considerations or long-term strategic goals. As such, human judgment and validation is central to the interpretation and augmentation of AI outputs.



Fair and impartial

When using AI for supplier evaluation, negotiating, and contracting, bias in the data or model could lead to unfair recommendations and discriminatory practices. By taking into account factors such as fair contract terms, social responsibility, and ethical sourcing practices, organisations can promote decision-making processes that are fair and transparent.

POTENTIAL BENEFITS

Resilient supply chains

Enhancing supply chain resilience allows the organisation to respond quickly to changing market dynamics and permits greater agility to take advantage of emerging opportunities based on real-time insights and recommendations.

Enhanced performance

By prioritising alerts that require human intervention and differentiating between noise and disruption, the organisation can drive greater efficiency in the supply chain.

Optimising efficiency

Making optimised decisions across the supply chain, from supplier selection to fulfilment optimisation, helps reduce costs, minimise waste, and improve overall operational efficiency.













64

Enabling a better grid

Grid and energy efficiency optimisation

Al can be used to better understand the state of the grid and factors that could support more efficient energy consumption, minimising losses and improving overall grid efficiency.

ISSUE/OPPORTUNITY

Energy grids are massive and intricate systems with interconnected components operating in a dynamic and uncertain environment. Maintaining a balance between energy supply and demand is crucial for grid stability, but it is challenged by the difficulty in predicting and managing fluctuations in energy demand. The integration of intermittent

renewable energy sources (e.g., solar) further complicates the supply-demand balancing act as these depend on weather conditions. Regulatory frameworks, policies, and market structures also constrain the ability to balance technical optimisation.

HOW AI CAN HELP

Promote informed customer behaviour

Energy companies can incentivise consumers to adjust their energy consumption based on their specific energy use patterns using conversational chatbots powered by generative Al. Al models can analyse historical data and customer preferences to recommend personalised strategies to reduce energy usage. When there is an immediate need to reduce peak loads to improve grid stability, Al applications can be used to alert customers about what they can specifically do to help. What is more, conversational chatbots can be used as an educational tool for consumers to understand and optimise their energy usage.

Document and map digitisation

Al can be used to digitise documentation, infrastructure maps, and records of energy use, as well as for image-to-image translation or image restoration (e.g., by removing noise, adjusting brightness, and enhancing

contrast). This improves the quality of the documents and yields searchable documents that can be used to train existing Al classification and forecasting tools.

Grid layout and expansion

Al can assist in designing optimal configuration and expansion plans for the energy grid. Al models can generate optimised grid designs that minimise transmission losses and maximise efficiency by considering factors such as population density, existing infrastructure, and energy demand projections.

Energy trading and market analysis

Al models can simulate the behaviour of electricity markets under different scenarios, such as regulation changes or the introduction of new technologies. This can help energy companies optimise their trading strategies and make more informed investment decisions.













Enabling a better grid

MANAGING RISK AND PROMOTING TRUST



Private

Using AI in customer behaviour analysis and chatbot interaction involves handling sensitive customer data. Risks include data breaches and unauthorised access to customer information and chat logs, and risk mitigation requires robust security measures, customer data protection, and adherence to privacy regulations.



Safe and secure Al models are vulnerable to adversarial attacks, where malicious actors manipulate inputs to deceive or exploit the system, for example, to influence energy trading decisions or disrupt grid operations. Robust security measures and regular testing are necessary to mitigate such risks.

POTENTIAL BENEFITS

Diversifying energy sources

Al supports the integration of variable renewable energy sources while maintaining stability and reliability.

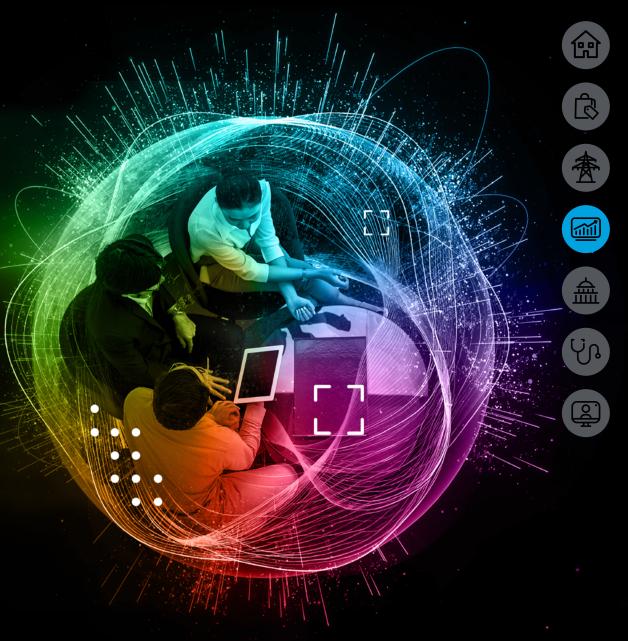
Dynamic demand response

Using AI for improved visibility of the grid's current state allows companies to better respond to fluctuations in demand.

Ongoing optimisation

As more trends, data and documents are digitised and analysed over time, Al enables continuous improvement in efficiency optimisation and managing

The Financial Services Al Dossier















The Financial Services Al Dossier

As AI capabilities mature and regulatory clarity improves, AI is moving from experimental to essential – powering financial services businesses that are smarter, faster, and more responsible.

In an environment where trust, precision, and speed are paramount, Al is enabling financial services firms to detect risk earlier, serve customers more effectively, and compete with greater agility. Whether it's fraud prevention in real time, hyperpersonalised marketing, predictive trading, or automated claim reporting, Al is shifting how institutions create value and manage complexity across banking, capital markets, and insurance.

Al is already having a transformative impact on knowledge work, helping people do their jobs more efficiently and effectively and supporting human decision-making with faster, deeper analysis of complex data patterns. It is also helping behind the scenes to triage and manage cyber threats, develop and debug new computer programs, and even generate synthetic data to train smarter Al models.

Agentic systems go even further, enabling AI to handle tasks such as algorithmic trading while keeping workers in the loop for oversight. Al's powerful capabilities can help financial institutions address rising pressure to modernise legacy systems, reduce costs, and respond to evolving customer expectations. However, they also bring heightened scrutiny around fairness, transparency, and accountability. Firms that want to harness the full power of Al will therefore need to invest not only in improved model performance but also in foundational elements such as governance, data quality, and organisational readiness.

Al's powerful capabilities come at an opportune time to help financial institutions address rising pressure to modernise legacy systems, reduce costs, and respond to evolving customer expectations.

Note: The tags below each use case indicate its primary business function and whether Agentic AI is used.

Primary business function

Agentic Al

67

Al-powered risk management and regulatory compliance

Using AI agents to provide 24/7 risk and compliance monitoring

Agentic AI systems can act as an always-on compliance and risk management team, with specialised agents continuously monitoring risks across a financial institution's operations.

ISSUE/OPPORTUNITY

Financial services firms must continuously monitor for risks – ranging from fraudulent transactions to compliance violations – across a vast landscape of operations.

At scale, with millions of customers and complex regulatory demands, risk management can be costly, labour intensive, and subject to error. Missed risks can lead to massive penalties, reputational damage, and systemic vulnerabilities. Yet, human teams are limited in their ability to

keep pace with the overwhelming volume of transactions, communications, and rule changes that require their oversight.

Traditional monitoring approaches can be siloed and reactive, identifying issues only after damage has been done. Multiagent Al offers a proactive, comprehensive solution that can scan for risks, interpret new regulations, and test institutional resilience under different scenarios.

HOW AI CAN HELP

Risk detection and monitoring

Agents that specialise in different risk areas (such as credit card fraud, anti-money laundering (AML), or cyber security) can automatically detect and monitor risks. They can also work together, sharing alerts across domains. For example, if a cyber security agent detects a breached account, it can warn a transaction monitoring agent to increase scrutiny on that account.

Regulatory change monitoring

Other specialised agents can parse new regulations using natural language processing, interpret rules and guidelines, map new requirements against company controls to identify compliance gaps, and then draft reports or suggest control changes for humans to implement.

Stress testing and scenario analysis

A stress testing agent can simulate scenarios such as interest rate jumps or credit crunches, coordinating with agents representing different departments (e.g., loans, trading, deposits) to assess institutional safety.

























Al-powered risk management and regulatory compliance

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Financial firms face massive penalties and reputational damage for compliance failures. As such, multi-agent systems should be validated against historical fraud cases, past audits, and synthetic stress scenarios to ensure their outputs are dependable and comprehensive.



Safe and secure

Pulling data from various databases to feed AI agents requires significant IT work and raises potential security concerns. Given the sensitive nature of risk management data, agentic systems must be secured against unauthorised access or accidental disclosure of private customer information.



Transparent and explainable

Risk management is overseen by humans who are personally liable and may be wary of trusting "black box" agents. Explainability is key. Agents need to provide clear audit trails (for example, explaining why a transaction was flagged, and including references to rules or past cases). This clarity makes Al outputs easier for human officers and regulators to trust and validate.



Fair and impartial Biases in data sources or training models can result in unfair or inaccurate risk assessments, raising concerns for customer trust, brand reputation, and regulatory compliance.

POTENTIAL BENEFITS

Higher assurance and lower compliance costs

Proactive and comprehensive monitoring reduces the likelihood of missed issues. helping firms avoid massive penalties while reducing the high labour costs associated with traditional compliance processes.

Reduced risk exposure and faster detection

Always-on agentic systems can sift through transactions, emails, and news at speed and scale, reducing a firm's exposure by identifying risks much sooner.

Improved audit and regulatory relations

Al agents specialising in compliance can speed up audits, doing in hours what might otherwise take human teams weeks to complete.











Ultra-personalised financial advice and wealth management

Using AI agents to deliver highly personalised financial planning

Agentic AI systems can serve as an automated financial advisor, providing tailored wealth management by analysing an individual's full financial picture and adapting plans to changing needs and market conditions.

ISSUE/OPPORTUNITY

Traditionally, financial advisors segment clients into broad categories – by age, income, or risk tolerance – and provide standardised advice. This limits personalisation and overlooks the nuances of an individual's financial situation. At the same time, wealth management firms often reserve high-quality, customised advice for affluent clients because it is expensive and resource-intensive to deliver at scale.

As client expectations shift toward greater personalisation and cost-efficiency, firms are seeking ways to deliver more customised service at lower operational cost – without compromising regulatory compliance or fiduciary responsibility.

HOW AI CAN HELP

Market monitoring

A market scout agent can track real-time developments such as interest rate changes, stock prices, and macroeconomic indicators to identify relevant opportunities and risks.

Personal financial analysis

A client profile analyst agent can aggregate financial data (e.g., bank accounts, credit card spending, loans) and behavioural patterns (e.g., risk-taking versus saving) to build a detailed and dynamic personal profile for each client.

Planning and strategy generation

A planning agent can synthesise market and client inputs to recommend bespoke financial strategies – for example, rebalancing investment portfolios in response to market changes or personal life events such as a home purchase or retirement.

Compliance integration

A compliance agent can help ensure recommendations adhere to fiduciary standards and regulatory guidelines, making the Al outputs safe and trustworthy for both clients and advisors.

Continuous fine-tuning

Unlike static robo-advisors, multi-agent Al systems can provide adaptive guidance, refining recommendations as circumstances evolve, rather than offering static one-time plans. They can identify portfolio risks in real time, flag anomalies for review, and automate activities such as threshold-based rebalancing, factor-based investing, and tax loss harvesting.

Ultra-personalised financial advice and wealth management

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Because errors could cause significant client losses, agents should be validated against historical financial data and tested under varied market and life-event scenarios to ensure they generate sound, resilient advice.



Transparent and explainable

Al-driven decisions – such as rebalancing actions or investment recommendations – should be accompanied by clear rationales and audit trails, enabling accountability and regulatory review.



Fair and impartial

Advisory models should be monitored to ensure they do not steer clients toward biased product selections, specific asset classes, or vendor-preferred solutions. Advice must be applied equitably across client segments.



Private

Because Al agents handle highly sensitive financial data, strong safeguards must be in place to protect information, including anonymisation and strict access controls.



POTENTIAL BENEFITS



Multi-agent AI makes high-quality, personalised wealth management affordable for mid-market clients. expanding access beyond the traditionally affluent segment.



Custom investment strategies ultrapersonalised by AI to fit an individual's life stages, goals, and preferences can improve customer satisfaction and financial outcomes in a cost-effective way, boosting retention and conversion.



Al-driven wealth management can help a firm establish a premium position in a commoditised market by offering a hightouch yet affordable client experience that is difficult to achieve in other ways.

Greater scalability and efficiency

Automation of portfolio management tasks enables firms to serve more clients at lower cost, freeing human advisors for complex decisions.













Al agents for algorithmic trading and market simulation

Enhancing trading strategies and insights with multi-agent collaboration

Agentic AI systems enable trading firms to use specialised agents to simulate artificial markets and to execute diverse trading strategies, enabling smarter, faster trades and richer insights into market dynamics.

ISSUE/OPPORTUNITY

Financial firms have long been at the forefront of harnessing intelligent technologies such as algorithmic trading and AI to achieve a winning edge. Now, multi-agent Al is taking the game to a whole new level

Traditional algorithmic trading transformed financial markets. However, behind closed doors, multi-agent systems are already in production and demonstrating better adaptability to market changes outperforming single-strategy approaches in various timeframes. In fact, it is now common for different algorithmic strategies (i.e., agents) to be running concurrently and even "competing" for capital allocation based on performance.

The latest innovation is to leverage more explicit agent frameworks and interagent communication, with multiple and varied trading or simulation agents acting independently yet sharing information to optimise outcomes. Looking ahead, we expect to see greater adoption of multiagent reinforcement learning for strategy development. We may also see exchanges using agent-based AI to monitor market stability or simulate the impact of rule changes. Without such tools, firms risk being outpaced in highly competitive markets (and regulators may risk missing early warning signs of instability). However, in many cases, the biggest innovations could remain hidden from view for competitive reasons.

HOW AI CAN HELP

Market simulation

In artificial market simulations. agents representing different trader archetypes (e.g., retail investors, institutional traders, market makers) can enable researchers, investment firms, and regulators to observe emergent phenomena and test different strategies and scenarios before they go live.

Specialised trading agents

A firm can simultaneously deploy live agents that apply different strategies in the market – such as short-term arbitrage, mediumterm trend following, or options hedging - sharing signals to avoid conflicts and improving portfolio resilience.

Coordinated action

Agents can operate semiautonomously under the supervision of a coordinator agent or human risk manager, ensuring portfolio alignment and preventing overexposure. In some cases, different Al agents might even collude or negotiate with each other to create a competitive advantage.

Reinforcement learning and communication

Advanced systems can use multiagent reinforcement learning, where agents learn from each other through trial and error in simulated environments. Some of these systems might leverage LLMs to enable agents to communicate and explain their reasoning, making the simulation results easier to interpret.













Operations

Agentic Al

MANAGING RISK AND PROMOTING TRUST

Al agents for algorithmic

behave reliably under pressure.

trader types, avoiding skewed outcomes.

prevent leakage of proprietary trading logic.

trading and market simulation

Because unpredictable agent interactions can create market

instability, systems should be stress-tested against diverse

scenarios, including flash-crash conditions, to ensure agents

Agents should be monitored to prevent behaviours that could

manipulation. Simulations should be designed to reflect diverse

unfairly favour certain counterparties or drift into market

Trading data and strategies are highly sensitive and commercially valuable. Systems should ensure strict data security and access controls to protect intellectual property and

Diversified trading performance

Multi-agent systems can function like an

Agents can operate in milliseconds, dynamically hedging exposures and adjusting positions to manage risk in volatile conditions. This increased speed can help a firm boost its trading volumes, validate analyses in real time, and potentially drive greater profitability while mitigating risks.

Agent-based simulations could provide regulators and institutions with cuttingedge tools to identify systemic risks or opportunities, enabling better preparation for disruptive events.

POTENTIAL BENEFITS

"Al investment committee," combining different strategies to improve returns across market regimes.

Faster reaction times

Systemic insight



















Fair and

impartial



Robust and reliable

Safe and secure

Given the adversarial nature of markets, safeguards should be established to prevent agents from being deceived or manipulated by other firms' Al. Strong cyber security and resilience measures are essential to maintain safe operations.

Al agents for credit underwriting

Adaptive, data-driven underwriting through multi-agent collaboration

Agentic AI systems can transform credit underwriting through specialised agents that analyse applicant data, monitor market context, assess risk, and maintain compliance – creating highly personalised lending decisions.

ISSUE/OPPORTUNITY

Traditional credit underwriting often relies on generic segmentation (e.g., age, income, or credit history) and static risk models. This approach can be slow, inflexible, and exclusionary, frequently locking out those with thin credit files or non-traditional income sources. At the same time, lenders struggle with outdated workflows that fail to adapt dynamically to changes in markets or individual circumstances.

As lenders work to expand their businesses responsibly, they need more agile underwriting solutions capable of assessing varied data and real-time conditions and delivering accurate decisions across a broad customer base.

HOW AI CAN HELP

Multisource data aggregation

A data gathering agent can compile structured and unstructured information (e.g., bank statements, tax filings, e-commerce history, location signals) to build an accurate and comprehensive borrower profile.

Dynamic scoring and simulations

Another agent can apply adaptive scoring models and run repayment scenario simulations, guiding approval, denial, or escalation recommendations based on risk capacity and context.

Regulatory alignment

A compliance-focused agent can confirm all assessments meet suitability requirements, fiduciary standards, and traceability benchmarks. This solution layer supports audit readiness and decision traceability.

Workflow orchestration

Agents working in concert – each handling discrete steps in the credit workflow, from data gathering to scoring and reporting – can provide continuous fine-tuning of underwriting decisions as new data or market conditions emerge.













ß











Al agents for credit underwriting

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Credit decisions can have life-changing implications. Agents must be validated continuously against historical performance and varied economic conditions. Simulated stress tests should regularly challenge outcomes to ensure resilience across credit cycles



Transparent and explainable

Regulators and borrowers require clarity on underwriting decisions. Agents should log their reasoning (e.g., which income streams or risk patterns influenced a score), presenting auditready explanations that demystify the model's logic.

POTENTIAL BENEFITS

Faster, smarter approvals at lower cost

Multiple data sources, specialised agents, and adaptive models enable faster, more precise underwriting without increasing manual workloads.

Financial inclusion

By incorporating non-traditional signals and actively monitoring for bias, Al agents can extend credit responsibly to underserved populations.

Operational scalability and quality

Multi-agent systems can boost lender efficiency and client satisfaction by enabling scalable, high-volume credit processing with consistent quality and regulatory compliance.













Transformation with speed and confidence

Code assistant for digital transformation

Al can enable banks to increase digitisation at a faster pace through code assistants.

ISSUE/OPPORTUNITY

Many financial institutions are pursuing cloud and data transformations, which are essential steps in preparing the organisation for using AI tools of many kinds. In some cases, legacy hardware is retired as data is curated and shifted to the cloud, freeing up humans for

more valuable work while bringing down the costs associated with on-premise infrastructure. Yet, these kinds of transformations are significant undertakings that can bring long lead times and high costs. There is also a risk of failure and error.

HOW AI CAN HELP

Supercharge your human capital

Al can be used as a component of cloud and data transformations to empower developers working across the enterprise on applications, data engineering, machine learning, and front-end development.

A helping hand in code development

As organisations explore new digital and cloud capabilities, development teams can accelerate and simplify their work by using Al as a force multiplier when writing, debugging, and documenting code, as well as translating ideas to code.

A shorter path to software

end code.

Part of success in transformation hinges on how quickly new enabling software can be deployed. There are opportunities to use Al in software development to shorten the lifecycle and more quickly reach a stable and deployable version, such as by helping rapidly write APIs, ETL, data pipelines, or even front-













Transformation with speed and confidence

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Partial automation of programming-related tasks requires the system to be reliably available and accurate. If availability cannot be guaranteed to an acceptable extent, the benefits of automation must be weighed against the risk of erroneous or buggy code.



Responsible and accountable

The training data for foundation models may create legal risks related to intellectual property or copyright infringement. If the training data contains copyrighted material, the organisation deploying the model needs to evaluate whether the presence of intellectual property in the training set could lead to legal challenges against the enterprise. Also, while the use of Al can accelerate the work of developers, without a human in the loop (e.g., to validate and debug code), critical failures may occur. Shoring up accountability may involve documenting and communicating standards and expectations for employees using



Safe and secure

By using an Al system, proprietary code bases may be exposed to third parties, raising questions around the security of the data and controlled access to it. An inadvertent breach of confidential intellectual property could have significant enterprise impacts.

POTENTIAL BENEFITS

Lower transformation costs

By shortening the software development lifecycle, the organisation can reduce overall costs for digital and cloud transformation.

Lower the bar to digital entry

Using AI opens the door for financial services organisations of all sizes, capabilities, and technology maturities to digitise and move to the cloud in a way that was previously out of reach for many.

Business intelligence at your fingertips

Enterprise-wide data search and access

Make business intelligence via enterprise data search accessible to all through natural language interfaces.

ISSUE/OPPORTUNITY

For many banking and insurance companies, data is stored in a multitude of locations, from local hardware to cloud storage solutions. This makes it difficult to effectively query different databases and retrieve relevant information quickly and efficiently. Complicating the matter, multiple mergers and acquisitions over time may have compounded the diversity

of data locations and databases, which hinders data mining for insights. Financial services organisations are informationintensive enterprises, and without the capacity to easily query all data, the result is poorer or incomplete insights that can increase enterprise risk while threatening customer satisfaction.

HOW AI CAN HELP

An enabling interface

Al capabilities can be built on top of an existing solution to facilitate the communication of queries from the user to the search layer. This serves as the interface between search layers and databases, allowing users to easily mine all enterprise data, as well as generate structured analytics reports.

Speed to insight

With the ability to query and analyse disparate data sources using AI as an interface, the enterprise can move past traditional business intelligence techniques and dramatically reduce the time required to generate insights while increasing the workforce accessibility to business intelligence.













Business intelligence

When it comes to governance and control, granting more data

access to a wider segment of the workforce can create a more

complex challenge of restricting who in the organisation is

Given the known challenges with model reliability and the

insights influencing decision-making and leading to negative

When dealing with sensitive and proprietary information, the

organisation must contend with securing the data, removing

information, either due to faulty function or a targeted attack.

or obscuring it in training and testing sets, and evaluating the model to determine whether it could "leak" protected

potential for hallucination, banking and financial services organisations face the risk of inaccurate or false Al-derived

permitted to access sensitive business data.

ramifications even at the market level.

at your fingertips

MANAGING RISK AND PROMOTING TRUST









Al as an interface between search and data enables business users to query databases depth programming experience. This gives intelligence without additional burdens on

A new level of data-

Real-time access to all of an enterprise's data can help organisations become even more insight driven, which supports improved growth prospects through access

POTENTIAL BENEFITS

Lower technical hurdles

and obtain tailored results without inmore of the workforce access to business IT and data science teams.

driven decisions

to the right insight at the right time.















Responsible and

accountable

Robust and

reliable



Fixing the missing data issue

Synthetic data generation

Generate synthetic data for model training, anomaly detection, and identifying cyber and deception attacks.

ISSUE/OPPORTUNITY

Missing data is a significant challenge for financial services organisations. Datasets may be incomplete, data transfers restricted, and potential anomalies underrepresented in the data. Using synthetic data can help overcome these challenges. In cloud transformation, data transfers may be delayed due to risks and regulations around data

governance, and use of synthetic data can enable a smoother and more efficient transformation. Meanwhile, machine learning anomaly detection systems (such as those for identifying fraud, waste, and abuse) must be trained on data from previous events. Yet, the dearth of data from these rare events can make anomalies harder to assess.

HOW AI CAN HELP

Improve model training

Generative AI can be used to quickly create synthetic data to supplement machine learning model training data, which is then used to aid and accelerate digital and cloud transformations. In this way, generative AI complements the enterprise's wider AI initiatives, fuelling (rather than replacing) other AI deployments.

Amplify anomaly event detection

The rarity of anomaly events can make it difficult to train machine learning systems to detect instances of fraud, waste, and abuse. Creating synthetic data with generative Al gives ML systems a larger suite of examples that lead to a greater capacity to find patterns and anomalies in the data.

Harden the organisation's

cyber posture

Just as synthetic fraud data can be used to train models to identify fraudulent activity, synthetic adversarial data can be used to train models to detect and mitigate cyber security risks and deceptive behaviours targeting virtual assistants.

















POTENTIAL BENEFITS

Faster path to the cloud

Generative Al-created synthetic data can accelerate digital and cloud transformations by making the transition smoother and more efficient.

Tackling fraud

Use synthetic data to train machine learning systems on rare or unknown events, such as novel types of fraud.

Security confidence

Adversarial synthetic data contributes to the enterprise-wide imperative to fortify digital assets against cyber threats.













Fixing the missing data issue

MANAGING RISK AND PROMOTING TRUST



Fair and imparti<u>al</u> A significant risk when generating synthetic data is that historic biases can creep into the generated data, perpetuating those biases (e.g., in the case of certain communities or socioeconomic groups being underrepresented in the data because those groups have had lower levels of banking activity in the past). This bias is not necessarily intentional but is still harmful.



Robust and reliable

Synthetic data created with generative AI can be limited in its scope and scale, and it should not be presumed to be accurate or perfectly reflective of real-world data. An overreliance on synthetic data may inject problems with data reliability, which can hamper the validity and usefulness of outputs and model training.

Getting to know your customer

Research-based report generation

Generate and summarise reports on new customers to inform employee decisions for customer onboarding.

ISSUE/OPPORTUNITY

New customers are the lifeblood of a growing financial services enterprise, but onboarding customers can be a highly manual and time-consuming process.

Know Your Customer (KYC) standards and rules require institutions to develop meta

reports on customers based on economic outlook, equity research, adverse media, and new prospect due diligence. These are high human involvement tasks, with deskbased meta-research consuming valuable time and resources.

HOW AI CAN HELP

Condensing results for easier consumption

Al can be used to summarise and filter results from existing search engines to inform meta reports, as well as to summarise information for the customer relationship manager.

Research and analysis to inform reports

Generative AI, alongside other machine learning models, can be used to conduct preliminary data searches and meta-analysis, potentially accelerating the KYC process.

















POTENTIAL BENEFITS

Timely insights

Faster and more efficient search and analysis can give decision-makers more up-to-date information and insights that enable better, more timely decision-making around customer onboarding.

Cost reduction

By streamlining and simplifying the report generation process, costly labour hours can be redirected to more valuable work.

Efficiency

The end user saves time and effort by more easily accessing and consuming relevant information.

MANAGING RISK AND PROMOTING TRUST

Getting to know

your customer •



Robust and reliable

When using AI to perform research and analysis, there is a risk that it could skip or misconstrue highly relevant information, which could skew the conclusions in the meta-analysis and hamper sound decision-making. If a new customer is erroneously assessed to be higher risk and a relationship manager declines the opportunity to engage the customer, the consequences are missed revenue and diminished customer engagement.



Private

When dealing with a customer's financial or personally identifiable information, the enterprise faces legal and regulatory standards for data privacy. When using AI, the organisation should take steps to ensure sensitive information does not inadvertently leak through model outputs, as well as govern who has access to the model, the underlying data, and the customer data it references.









Enhanced AI support for customers

Financial guardian

A personal virtual assistant powered by Al can cater to daily needs of customers.

ISSUE/OPPORTUNITY

Customer service and engagement is vital for financial services organisations, particularly as they transform from a product-focused to a customer-focused business. As a part of that, there is a pressing need for the enterprise to rapidly and accurately answer both common and

complex customer queries and do so with the timeliness customers expect. Yet, increased digitisation in financial services has reduced access to representatives who can answer customer questions; this is at a time when financial services customers seek a hyper-personalised experience.

HOW AI CAN HELP

An empathetic, personalised interface

Generative AI can be paired with other models to create a customer interface that delivers a hyper-personalised experience, such as by training the generative AI model to provide answers or insights with empathy. It also overcomes some of the challenges around more traditional chatbots that can lack an empathetic tone.

A more capable digital agent

Using generative AI can improve the usefulness and accessibility of a chatbot interface. With the capabilities of large language models (LLMs), a personal digital assistant can summarise contracts and answer nuanced questions, and the customer may enjoy a range of interface options, including text, audio, and imagery.













Enhanced AI support for customers

MANAGING RISK AND PROMOTING TRUST



Responsible and accountable

While a generative Al-enabled solution may provide valuable answers and recommendations most of the time, there is a risk that too much confidence may be placed in the validity of the outputs, both by the organisation and the customer. Generative AI is not an infallible oracle, and an overreliance on the Al solution may have a detrimental impact on customer actions, which can in turn increase financial risk.



Private

When confidential or personally identifiable information is inputted via the digital interface by the customer, the financial institution is obligated to follow the laws and rules that dictate how that sensitive information can be transmitted, stored, and accessed. Failing to do so could raise legal peril and potentially subject the enterprise to greater cyber risks.



Transparent and explainable

End users require a clear understanding of how their information will be processed and that they are interacting with a machine. At the same time, the enterprise needs to be able to interpret outputs and understand how and why the generative AI model created a given output.

POTENTIAL BENEFITS

Customer satisfaction

Greater accessibility and more timely answers can lead to a more personalised and satisfying customer experience. This can drive an increase in net promoter score, reflecting increased customer retention and loyalty.



While customer loyalty and brand reputation fuel business growth, integrating generative AI into virtual assistants can further enhance operational efficiency. Alenabled systems can serve more customers at scale while the human workforce focuses on resolving the most complex issues or attracting new customers.

Financial inclusivity through hyper-personalisation

When customers feel that banking applications are delivering a valuable, personalised experience, it encourages greater customer engagement and interest in service offerings, supporting both the customer's financial wellbeing, as well as that of the enterprise.













Customised marketing for the individual

Hyper-personalised sales and marketing assistant

Regulatory-compliant marketing material generation across different geographies.

ISSUE/OPPORTUNITY

Financial services marketing operations are increasingly coming under regulatory scrutiny for issues such as mis-selling and misinformation. Part of the challenge for multinational organisations is that cultural differences as well as varying levels of customer understanding about products

may create regulatory risk for enterprises in a given geography. To overcome this, organisations are investing significant manual labour to maintain a compliant marketing function, which is both time consuming and costly.

HOW AI CAN HELP

Customised materials for different audiences

Al can be used to create marketing materials that contain the appropriate tone, language, and cultural references, while also supporting consumer understanding of the product to maintain regulatory compliance.

Personalised sales at scale

With AI, financial services organisations have the ability to create marketing materials that are customised to individual customers – and to do so at scale.



























Customised marketing for the individual

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

For Al-generated marketing to be valuable, organisations must be able to rely on the validity of the output. Al can be prone to hallucinations, and when tasked with creating marketing that touts one product over another, there is a risk the model will return false statements. This injects potential regulatory violations that could result in fines and other penalties. To shore up reliability with Al, outputs should be validated by a human with the subject matter understanding to do so.



Fair and impartial Datasets may contain latent bias of which the organisation is unaware. This could be due to how the data was acquired, recorded, and curated, and the challenge is compounded when operating in multiple geographies. Organisations should consider the datasets used to train and fuel generative Al systems and whether unknown bias could lead to marketing materials that fall short because they fail to reflect important geographical and cultural differences.

POTENTIAL BENEFITS

Individual-level marketing

Hyper-personalised marketing takes the enterprise to a new stage of customer engagement and enticement that is infeasible through manual effort alone.

Confident compliance

Using AI for marketing development can help ensure the content remains in line with regulatory expectations across many geographies, thereby reducing regulatory risk.

Driving marketing ROI

Personalised marketing can support new sales, strengthen the customer relationship, and reduce the cost of marketing operations while also improving the timeliness of outreach and engagement.

Ensuring the integrity of claims

Automated claims reporting

Al can be used to automatically generate reports based on descriptions or pictures of the relevant subject.

ISSUE/OPPORTUNITY

During claims processing for property and casualty insurance, claims agents must decide whether a destructive event was insured, and if so, the amount and cost of the damage. These processes are complex and can be time consuming, and claims agents have few tools to support their decision-making.

HOW AI CAN HELP

Virtual damage rendering

Al can be used to help visualise the damage by replicating it virtually. The basis for the replication and visualisation can be customer conversations, damage documents, photos, official reports, and other relevant media. In visualising the data in this way, the claims agent is empowered to make better decisions when assessing the degree and cost of damage.

Automated claims reporting

With AI, claims reports can be generated automatically based on photographic evidence.















POTENTIAL BENEFITS

Cost reduction

Faster claims processing and more accurate damage assessments can reduce labour costs and claims payments.

Customer satisfaction

By processing claims faster and with less administrative burden, customer satisfaction improves due to quicker, more streamlined adjudication and payment.

Identifying fraud

More expedient claims processing has the advantage of identifying potential fraud more quickly, helping to ensure the integrity of claims and payments.















MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Damage visualisation requires a high degree of accuracy, and erroneous AI outputs could lead to claims being paid incorrectly, potentially leading to overpayment (a detriment to the organisation) or underpayment (a detriment to the customer).



explainable

Transparent and If claims agents use AI to automate aspects of claims processing but are unable to articulate to customers how the Al model derived its outputs or contributed to the cost and damage assessment, customers may not accept the outcome of the claims process.

A virtual bank experience

VR-enabled retail banking centres

Virtual reality customer agents powered by generative AI can change the retail banking experience and interactions.

ISSUE/OPPORTUNITY

Banks seek to provide customers with multiple methods of interacting with their accounts, services, and offerings. As digital services grow, customers may prefer to conduct banking transactions remotely through virtual transactions. However, this approach still requires a human agent, and employees are the most expensive

cost component in customer service – particularly when three levels of customer support are necessary. Chatbots can help automate virtual transactions, but existing chat tools are limited to specific, preprogrammed dialogue and options.

HOW AI CAN HELP

Bring the bank to the customer

With a generative Al-enabled virtual space, customers can use a VR headset to conduct business with the financial institution and interact with a service representative from the comfort of their own home in a way that is convenient for the customer.

Hyper-personalised service

A generative AI agent can provide conversational, tailored responses to questions about customer accounts and financial needs. This helps the enterprise cater to its customers' desires for a personalised experience while also avoiding the costs that come with adding more human customer service workers.

In a virtual space, customer data can be accessed in real time by the conversational agent to provide faster, higher quality service and offerings.

Speed and quality of service

























A virtual bank experience

MANAGING RISK AND PROMOTING TRUST



Transparent and explainable

Given a generative Al-enabled chatbot's capacity for conversational outputs, there is a risk that customers may not realise they are interacting with a machine. To promote trust, customers should not only be informed that the chatbot is not human, but they should also understand how their inputs and information are stored, accessed, and used.



Fair and impartial The datasets used to train and inform the chatbot may contain latent biases, such as underrepresented customer groups or semantic deficiencies in some languages but not others. As a result, the model may simply not work as well for some customers, leading to a variety of negative customer impressions and complaints.



Responsible and accountable

With issues around reliability and data quality, if the chatbot outputs erroneous data or recommendations, a human stakeholder needs to be accountable for the outcome. Organisations can promote accountability by keeping a human in the loop and documenting roles and responsibilities.

POTENTIAL BENEFITS

Customer-focused experience

A generative AI-enabled virtual banking experience connects customers with tailored digital representatives that can converse in the customer's preferred language in a timelier manner.

Liberated human capital

With conversational chatbots satisfying most customer inquiries, the human workforce can focus on the most complex or value-driving customer service activities.

Cost avoidance

Virtual agents allow the bank to serve more customers without expanding the human workforce, which helps limit customer service costs.



Focused cyber

Al-driven vulnerability management that focuses on real threats

Al can help filter, analyse, and prioritise security alerts, allowing organisations to focus their resources on true vulnerabilities – which can help reduce their cyber security risk.

ISSUE/OPPORTUNITY

Financial institutions and other large enterprises rely on multiple security tools that generate millions of alerts daily – many of which are false positives or non-actionable.

Security teams and developers spend time triaging and remediating vulnerabilities that pose little or no real risk. This inefficiency delays responses to real threats, creates operational fatigue, and increases the likelihood of breaches – threatening business continuity and brand reputation.

HOW AI CAN HELP

Evaluating alerts

Al models can evaluate alerts by calculating breachability (likelihood the vulnerability can be tapped in a real-world scenario) and exploitability (likelihood it can be actively taken advantage of to inflict harm).

Setting priorities

An Al-driven system can use enterprise-specific context such as workload risk profiles, application architecture, and network exposure to separate high-risk from low-risk vulnerabilities. Al automatically identifies which threats demand attention and which can be deprioritised, providing a data driven, risk-based prioritisation process.













Focused cyber

MANAGING RISK AND PROMOTING TRUST

POTENTIAL BENEFITS

Less noise

Moving from thousands of daily alerts to a handful of real threats enables security and engineering teams to focus their time and effort more effectively.



Improved security and reputation

Decreasing security incidents strengthens a company's standing with customers, investors, and regulators.



Cost savings and improved efficiency

Better resource allocation allows security and engineering teams to spend less time and money protecting the organisation from cyberthreats.



Enhanced productivity and morale

Minimising time spent on false positives reduces friction between the security and engineering teams, improving productivity and morale.



Robust and reliable

The vulnerability management system should be continuously retrained and tested against live alert data and historical incidents to validate its prioritisation logic. A human-inthe-loop approach provides important oversight on critical decisions, especially during early-stage deployment.



Safe and secure The models should be hosted in secure environments (e.g., onprem or restricted cloud instances), with strict access controls and audit trails. Systems must be designed to prevent data leakage and be regularly tested for robustness.

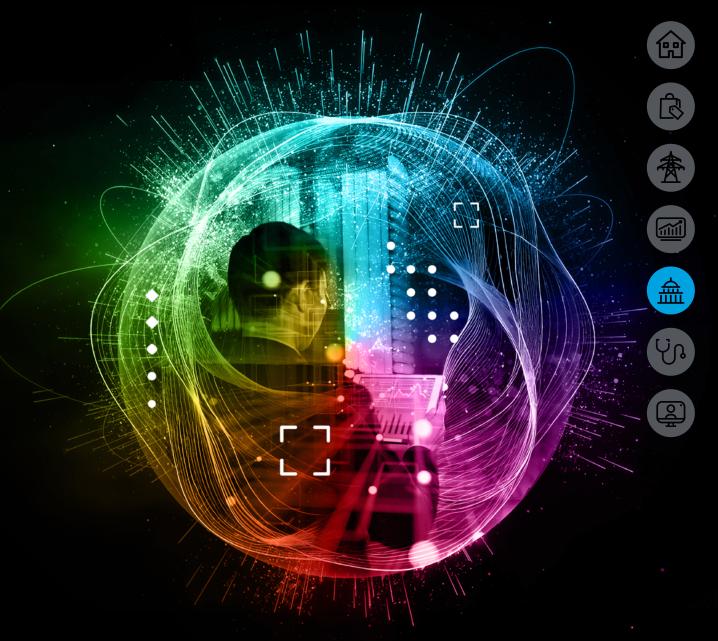


Transparent and explainable

Al outputs should include the rationale for prioritisation with traceable decision paths that help analysts understand and trust the recommendations.



The Government & Public Services Al Dossier

















The Government & Public Services Al Dossier

Al is becoming a critical tool for driving mission outcomes in the government and public services industry. From frontline service delivery to policymaking and infrastructure management, Al systems are reshaping how government and public services organisations operate, deliver services, and engage with constituents.

In an environment often constrained by legacy systems, budget pressures, and complex regulatory mandates, AI offers the ability to modernise processes, improve responsiveness, and unlock new levels of efficiency and insight.

Al is already being deployed for a wide range of use cases across the sector, such as drafting policy documents, analysing and summarising global legislation and policies, simulating urban planning scenarios, hyperpersonalising education, and engaging with the public through virtual public servants.

Increasingly, agentic AI enables more autonomous support across many areas – including regulatory oversight, permitting, benefits, intelligence, and health documentation – augmenting, not replacing, human expertise.

Such capabilities, when paired with secure data integration and worker oversight, can reduce administrative burdens, streamline operational workflows, and enhance the precision and speed of government decision-making.

Al systems are

reshaping how

public services

operate, deliver

services, and engage

with constituents.

organisations

government and

As government and public sector use cases mature and the regulatory landscape evolves, AI is poised to become not just a tool for modernisation, but a catalyst for more adaptive and data-driven governance. However, the adoption of AI in government and public services also brings unique challenges.

Public institutions must meet high standards for transparency, fairness, and accountability – often in environments where privacy concerns, procurement complexity, and uneven digital maturity can limit deployment. As such, success depends not just on technical innovation, but on cross-agency collaboration, trust-building with the public, and alignment with legal and ethical frameworks.

Note: The tags below each use case indicate its primary business function and whether Agentic Al is used.

Primary business function

Agentic Al

B











Al-supported regulatory examinations and inspections

Augmenting regulators with multi-agent systems for fast, more accurate review

Agentic AI systems can support regulators with specialised agents that analyse data, review large volumes of documents, and uncover insights to streamline examinations and inspections.

ISSUE/OPPORTUNITY

Regulatory examinations and inspections are essential to ensuring compliance in industries such as financial services, food safety, and transportation. These processes require regulators to review vast amounts of historical and current market data, pull from multiple document sources, and interpret complex regulations. The work is labour-intensive, time-consuming, and prone to bottlenecks, which can delay the detection of compliance risks.

As regulatory responsibilities expand and the business environment becomes increasingly complex, regulators need tools that can streamline and accelerate analysis without sacrificing quality or rigor. Agentic Al offers the ability to automate routine examination tasks, highlight anomalies, and provide real-time support to examiners.

HOW AI CAN HELP

Quantitative analysis

Al agents can extract formulas, run calculations, and compare data against regulatory benchmarks to identify areas of potential concern.

Document review

The systems can process large volumes of reports, filings, and other documents, uncovering key insights and flagging areas that warrant attention from human examiners.

Dynamic examiner support

A conversational interface allows regulators to interact with the system, request clarifications, and guide the direction of analysis, ensuring a human stays in the loop.

Cross-domain applicability

The framework currently being used for financial services examinations can be adapted to other types of regulatory inspections, including food safety, tax compliance, and transportation.













Al-supported regulatory examinations and inspections

MANAGING RISK AND PROMOTING TRUST



Private

Because examinations can involve sensitive institutional and personal data, agentic systems should employ strict data protection protocols and comply with all relevant privacy regulations.



Robust and reliable

Flawed analyses can lead to missed risks or incorrect enforcement actions. To avoid such problems, Al agents should be validated against known regulatory outcomes and tested continuously for accuracy.



accountable

Responsible and Since regulatory judgments can have serious legal and economic consequences, Al agents should be used primarily as decision-support tools, with human regulators retaining final authority over important decisions and actions



Fair and impartial

To ensure equitable oversight, agents should be designed to apply standards consistently across institutions and sectors, minimising the risk of bias in how examinations are prioritised or conducted.

POTENTIAL BENEFITS

Faster examinations at lower cost

Automated analysis and document review can shorten the time required to complete inspections, helping regulators do their jobs efficiently and effectively.

Improved risk detection

By examining vast amounts of quantitative and qualitative data, Al agents are more likely to detect problems or anomalies that could go unnoticed.

Scalable oversight

With AI agents handling routine analysis, regulators can expand their oversight capacity without a corresponding increase in staff, helping to maintain compliance in complex and high-volume domains.

Al-driven permitting

Accelerating and simplifying government permitting processes

Agentic AI systems can streamline permitting workflows by using specialised agents to scan applications, extract key data, check compliance, and provide real-time feedback to applicants and reviewers.

ISSUE/OPPORTUNITY

Permitting is a critical function for government agencies at the central and local government level, covering areas such as land use, energy, waterways, and housing. These processes can be slow, opaque, and labour-intensive, requiring reviewers to manually assess large volumes of documentation against complex and sometimes overlapping regulations. Small errors or missing

information in applications can create lengthy delays, increasing costs for both applicants and agencies. With government entities facing larger workloads and limited staffing, applicants and regulators alike are seeking opportunities for efficiencies. New approaches can help accelerate permitting while ensuring compliance with local, state, and federal requirements.

HOW AI CAN HELP

Automated document review

Al agents can scan and extract relevant data from submitted applications, checking for completeness and identifying missing or inconsistent information.

Compliance validation

Specialised agents can evaluate applications against permitting requirements, standard operating procedures (SOPs), and applicable regulations, flagging potential compliance issues early in the process.

Applicant feedback and support

Using AI and other advanced technologies, agents can provide real-time feedback to applicants, helping them correct errors or supply missing documentation before delays occur.

Reviewer assistance

Al agents can support human reviewers by summarising key findings, verifying that evaluations align with regulatory standards, and highlighting areas requiring human attention.













POTENTIAL BENEFITS

Reduced costs and permitting timelines



By automating document review and compliance checks, processing times can be reduced from months to days, lowering costs and accelerating the permitting process.

Improved consistency and fairness

Standardised evaluations conducted or supported by AI agents can help all applicants be assessed against the same criteria, reducing variability across reviewers.

Reduced applicant frustration

Real-time feedback improves transparency and helps applicants correct issues earlier, reducing costly delays and improving trust in the process.













Al-driven permitting

MANAGING RISK AND PROMOTING TRUST



To maintain public trust, agents should be designed to apply requirements consistently across all applicants.



Robust and reliable

Errors in interpreting regulations can lead to incorrect approvals or major delays. To improve consistency and reliability, agents should be trained on verified SOPs and continuously updated as requirements evolve.



Private

Since applications often contain sensitive personal and corporate data, agentic Al systems should be designed with strong data protection measures and full compliance with privacy regulations at all three levels of government.



Responsible and accountable

Permitting decisions can have important legal and financial ramifications. To avoid errors in accuracy or judgment, agentic Al systems should be primarily used for decision support, with final approval authority retained by human reviewers.













Al-enhanced benefits eligibility

Improving access to government benefits programs

Al systems can streamline the application and determination processes for government benefits programs and public assistance.

ISSUE/OPPORTUNITY

Determining eligibility for public benefits such as health services, food stamps, unemployment insurance, workers comp, and childcare can be a time-consuming process for applicants and government agencies alike. Traditional public assistance portals may have redundant data entry, unclear requirements, and complex navigation, which can frustrate applicants (especially those with limited digital literacy) and delay access to benefits.

Meanwhile, government caseworkers manage lengthy manual intake processes and may have a lack of support in interview preparation and execution. These challenges increase agency workloads and can lead to incomplete applications, inconsistent reviews, and a higher likelihood of errors

HOW AI CAN HELP

Conversational intake

Conversational agents can replace static application forms with guided, human-like interactions that clarify questions and adapt dynamically to applicant responses.

Automated document intelligence

Document analysis agents can extract and validate data from submitted records, reducing manual effort and maintaining accuracy in eligibility assessments.

Interview assistance

Al agents can support human caseworkers by preparing interview outlines, validating applicant responses in real time, and highlighting missing or inconsistent information.

Onboarding and on-the-job assistance

Al can accelerate new caseworker onboarding by generating tailored training materials and insights from historical cases, helping agencies address workforce challenges more efficiently. Al agents can also support human caseworkers by preparing interview outlines, validating applicant responses in real time, and highlighting missing or inconsistent information.

End-to-end applicant support

The specialised agents can work together across the intake and review processes to provide applicants with a seamless experience that is less intimidating and confusing, helping to ensure applications are complete and compliant.













Al-enhanced benefits eligibility

MANAGING RISK AND PROMOTING TRUST



explainable

Transparent and In order for applicants and caseworkers to trust Al-driven recommendations, agents should provide clear explanations of eligibility determinations, including references to applicable rules and policies.



Robust and reliable

Conversational agents might misunderstand user inputs, especially in cases of language barriers or low digital literacy. Agents should be trained on multiple languages and continuously monitored for performance.

POTENTIAL BENEFITS

Faster access to benefits

Automating intake and eligibility can reduce processing times, helping applicants receive support more quickly.

Cost savings and improved productivity

Using Al agents to streamline processes and automate manual tasks can reduce costs and enable caseworkers to focus more attention on complex cases and applicant support.

Improved applicant experience

Guided interactions, real-time feedback, and simplified workflows can reduce frustration and confusion, especially for applicants with limited digital literacy.

Smarter intelligence

Improving intelligence reporting via AI agents and multimodal data integration

Agentic AI systems can autonomously collect, synthesise, and crossreference intelligence data from various sources, producing faster, more comprehensive insights.

ISSUE/OPPORTUNITY

Intelligence analysts, including those working with open-source intelligence (OSINT), face the daunting task of searching, collecting, and interpreting vast amounts of multimodal data – from social media and public records to satellite imagery, communications, and financial reports. This manual process is time-consuming and limited by the capacity of human analysts to track patterns across large datasets.

As threats evolve and the volume of public and classified data continues to grow at breakneck speed, agencies need tools that can help accelerate the production of intelligence reports while improving accuracy, adaptability, and depth of insight. Agentic Al offers the potential to automate core intelligence gathering activities so human analysts can focus on validation and higher-level reasoning.

HOW AI CAN HELP

Automated data discovery

Al agents can continuously scan for and retrieve information from multiple structured and unstructured data sources, both classified and public – including briefings, news media, and other reports – reducing the need for manual searches and data collection.

Contextual linking of insights

If analysis of an old report produces new leads – such as a location or personal association – Al agents can automatically launch related queries to build a more complete picture of personal networks and activities.

Multimodal data integration

Specialised AI agents can process different types of inputs (e.g., text, imagery, sensor data) and combine the results into comprehensive, integrated insights.

Adaptive intelligence reporting

As the various agents collaboratively draft and refine reports, they can highlight key findings and update them dynamically as new data becomes available.

























Smarter intelligence

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Automated agents might misinterpret ambiguous or misleading data, potentially affecting the accuracy of intelligence outputs. This is especially true for open-source information, which can be intentionally misleading or fake. Al agent actions should be validated against verified sources and continuously monitored to reduce errors.



Safe and secure

Given the sensitive nature of intelligence queries, special care must be taken to prevent adversaries from influencing the model (or gathering their own intelligence from what is queried).



Transparent and explainable

Providing adequate transparency in autonomous information exchange can pose significant technical and operational hurdles, especially given the complexity of integrating data from multiple sources. Al agents should document their reasoning and data sources so the outputs are easier to trust and audit.



Responsible and accountable (

Given the life-and-death consequences of intelligence information, Al agents should be viewed as decision-support tools, with human analysts and leaders continuing to have final responsibility for decisions and actions.

POTENTIAL BENEFITS

Faster examinations at lower cost

Automating data collection and integration accelerates intelligence reporting and reduces manual effort, which can enable agencies to respond more quickly and efficiently to emerging threats.

Deeper insights

Cross-modal, autonomous querying and analysis can uncover hidden connections that might be missed with traditional manual processes, providing a more comprehensive and dynamic understanding of complex domains.

Improved analyst capacity

By offloading routine search and synthesis tasks, Al agents allow human analysts to focus on higher-level reasoning, strategic decision-making, and collaboration with colleagues.



Automated clinical documentation

Reducing provider burden by automating patient visit notes

Agentic AI systems can automatically generate structured clinical documentation from patient visits, reducing providers' workload, improving record quality, and enabling them to spend more time with patients.

ISSUE/OPPORTUNITY

For providers in large government health care systems, documenting patient visits can be an administrative burden. Traditional processes require physicians to spend hours after each visit manually writing SOAP (Subjective, Objective, Assessment, Plan) notes and updating electronic health records (EHRs). This

not only disrupts the flow of care but often requires after-hours "pajama time" spent on paperwork, which contributes to provider burnout. Also, incomplete or inconsistent documentation creates risks for care quality, regulatory compliance, and continuity across providers.

HOW AI CAN HELP

Automated transcription and generation

A generator agent listens to the audio of patient visits, converting it into structured SOAP notes and EHR-ready data, which minimises the need for manual entry.

Validation and compliance

A validator agent reviews the generated notes for clinical accuracy, completeness, and regulatory compliance, confirming that records meet organisational and legal standards.

Workflow orchestration

An orchestrator agent manages the entire documentation process, routing tasks among the specialised agents and providing smooth, efficient workflows.

Seamless integration

The agentic Al system formats the validated documentation into common formats (e.g., DOCX, HTML) for direct integration into EHR systems, reducing friction and manual effort for providers and staff.

























POTENTIAL BENEFITS

Lower costs and reduced provider burnout

Automating documentation can reduce clinical staff workload, reducing costs and freeing providers from after-hours paperwork, which can improve their worklife balance.

Improved record quality and compliance

Standardised, Al-validated notes can improve the completeness and consistency of patient records, enhancing care continuity and compliance with regulations.

Data-driven insights

Structured documentation enables advanced analytics and quality improvement initiatives, helping government health systems improve care delivery at scale.

Automated clinical documentation •

MANAGING RISK AND PROMOTING TRUST



Fair and impartial Documentation agents should be designed and tested to work effectively across varied patient populations and care contexts.



Robust and reliable

Because errors in transcription or orchestration could compromise documentation quality, Al agents should be validated regularly and continuously updated to reflect evolving medical and compliance requirements.



Transparent and explainable

For providers to trust that Al-generated documentation accurately reflects patient encounters, Al agents must be able to explain how notes were created and why specific items were flagged for correction.



Private

Some patients might be uncomfortable having their medical visits recorded and then analysed by Al. Every patient should be given the choice to opt out. Also, Al agents must strictly comply with applicable New Zealand privacy and health regulations and other relevant regulations, encrypt audio and text records, and minimise retention of personally identifiable information.

Virtual public servant

Citizen engagement

Al can enable virtual assistants that provide personalised responses to citizen questions about public services.

ISSUE/OPPORTUNITY

Government organisations perform a range of functions, from supporting public health to promoting tourism. Data about government and public services, however, is often stored in a variety of formats and locations (e.g., on-prem, cloud), challenging interoperability. When citizens contact agencies to inquire about

services and resources, human agents are challenged to rapidly access and summarise information to satisfy citizen questions. This is a time-consuming, labour-intensive endeavour for the organisation, and it may not meet citizen expectations for fruitful engagement.

HOW AI CAN HELP

A digital agent for engagement

An Al-enabled virtual assistant can serve as the interface between citizens and government information, helping with questions and transactions via empathetic, natural language.

Reaching across datasets

The virtual assistant can distill and summarise information from myriad sources on a variety of topics to answer questions in a multitude of languages regarding service requirements and appointment options.

























Virtual public servant

MANAGING RISK AND PROMOTING TRUST



accountable

Responsible and While virtual assistants may be valuable for providing information, they may not be suited to providing true insight and advice. Agencies need to guard against over-reliance on Al solutions and the potential for citizens to take some action based on faulty or improper Al output.



Safe and secure

A model tasked with providing accurate information may be a target for cyber criminals seeking to access sensitive information or manipulate the model and its underlying data. Many government agencies contend with cyber security regulations and standards, making model security a priority.



Robust and reliable

Model accuracy and timeliness depends in part on the data sources it can access. If information is outdated or incorrect, it creates a risk of erroneous outputs. Human stakeholders responsible for updating information have a direct impact on model reliability and user trust.

POTENTIAL BENEFITS

Promoting citizen engagement

When public services are more accessible due to more efficient and robust technology, user engagement and citizen satisfaction in government offerings increase.

Increasing accessibility

A virtual assistant powered by AI can interact with citizens in their preferred language and ultimately help bring down social barriers to engaging public services.

Citizen satisfaction

Government agencies operate in service to the public, and providing fast access to information about services promotes a positive public perception of government function.

Insights for all

Knowledge management

Al can serve as an interface to help public sector organisations become insightdriven by making data more accessible.

ISSUE/OPPORTUNITY

From the census to transportation and procurement, government agencies collect and release huge amounts of open datasets. By encouraging the use, reuse, and distribution of open datasets, government organisations can promote data-driven innovation and

citizen-centric services if combined with an agency's internal datasets. For public industry stakeholders to become truly insight-driven, they require the means to interrogate all relevant data, even if they lack a technical background in data science or related fields.

HOW AI CAN HELP

Greater accessibility

Al can provide a natural language interface that allows non-technical users to access and understand data that might otherwise only be accessible to technical users.

Democratising insights

Rather than placing all of the burden for data analysis, interpretation, and visualisation on a technical team, an Al interface reduces the effort by allowing more stakeholders to work with the data and derive their own insights.

























Insights for all

MANAGING RISK AND PROMOTING TRUST



Safe and Secure

An Al model that taps into a variety of datasets can make it difficult to control which data is accessed by which stakeholders in which organisation, raising important considerations for model security and governance.



Private

When dealing with sensitive and proprietary information that is subject to varying laws and regulations across jurisdictions, organisations are called on to ensure the Al model does not leak, inadvertently divulge, or inappropriately access sensitive or restricted data.



explainable

Transparent and To accurately interpret data and Al outputs, the end user needs to understand which data was referenced for the output, which could not be accessed, and the potential biases in the available data.

POTENTIAL BENEFITS

Scaling data access

An AI solution that can access a variety of datasets and data types allows public servants to draw conclusions from a broader set of knowledge and information.

Fostering collaboration

When more public servants can access insights and knowledge, it promotes insight-driven action across agencies, helping to fuel greater collaboration between a larger set of stakeholders.

Faster insights

Al can help accelerate the process of identifying and consuming relevant information, driving speed and efficiency.

Simulating urban planning scenarios

Urban planning/future of cities

Al can be used to help urban planners in the ideation and design of novel urban concepts.

ISSUE/OPPORTUNITY

The scale and speed of urbanisation brings a host of challenges, such as lack of affordable housing, overburdened transportation systems, traffic congestion, lack of drinking water, sanitisation issues, and environmental quality. The challenge for city officials and urban planners is to imagine the future of cities by overcoming creative hurdles and developing city designs that are resilient, sustainable, and human centric.

HOW AI CAN HELP

Generating 3D city models

Using AI, thousands of 3D images can be rapidly created to help guide and refine a city design. Such 3D images form part of the design brief for urban planners and the master city plan.

Simulate natural disasters

Al can simulate natural disasters like earthquakes, floods, or hurricanes to evaluate the vulnerability of city infrastructure and plan for resilient urban infrastructure.

Planning for the future

By simulating population growth and demographic trends, Al can develop scenarios for urban expansion and plan for adequate infrastructure, housing, transportation, and public services that accommodate urban growth.



























Simulating urban planning scenarios

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

While AI models may create interesting or attractive designs, they require human review and validation to ensure they meet urban planning requirements and can be feasibly built in the real world.



explainable

Transparent and A lack of contextual knowledge of urban planning may lead Al to develop improbable scenarios, and analysts need to be able to understand how and why the model produced an output in order to confirm and validate it.

POTENTIAL BENEFITS

Super-charge creativity

Using AI to rapidly create a plethora of designs and scenarios helps city officials imagine the future of cities and plan for upcoming challenges.

Faster ideation and iteration

With a faster method to create design iterations, urban planners can accelerate the design and decision-making processes.

Improved decision-making

Using AI in city planning enables decisionmakers to model various scenarios and optimise urban designs for better resource utilisation, sustainability, and quality of life for residents.

Education 2.0

Hyper-personalised education

Al can be used to hyper-personalise digital teachers that can adapt to student learning needs and curricula.

ISSUE/OPPORTUNITY

The demand for schoolteachers can often exceed supply. While the available teachers contend with larger class sizes, they also need to accommodate students with different learning styles and educational needs. Yet, because of

the one-to-many nature of traditional schools, teachers are challenged to deliver the kind of personalised learning support and instruction that students need to be successful.

HOW AI CAN HELP

A digital, adaptive teacher

Al can serve as a virtual instructor, drawing from resources and lesson plans to hyper-personalise the learning experience. The model can check the student's work and comprehension and adapt lessons and learning strategies according to the student's individual weaknesses, strengths, and preferences.

A force multiplier for teachers

When personalised digital teachers can work with students one-on-one to master new skills and knowledge, the human instructor can focus on higher level planning, interacting with students, evaluation, and student support.























Education 2.0

MANAGING RISK AND PROMOTING TRUST



accountable

Responsible and While digital teachers can offer valuable advantages in adaptive learning, the model should not be expected to satisfy all of the important lessons teachers impart, such as social lessons around collaboration, conflict resolution, and empathy. The human element in teaching is essential, and educational institutions need to take a responsible approach to integrating Al-enabled teachers.



Robust and reliable

Because Al is susceptible to outputting inaccuracies and hallucinations, there is a risk that a virtual teacher could teach incorrect facts or produce poor learning strategies.



Private

Student data is subject to education regulations, making model security and data privacy a priority when deploying digital teachers.

POTENTIAL BENEFITS

Catering to the student

Employing adaptive learning with AI can promote knowledge retention and understanding by tailoring teaching approaches to the student's learning style.

Remedy the talent gap

Leveraging AI helps overcome teacher shortages, allowing more students to access quality education.

Removing barriers

An Al-enabled teacher is not restricted to a physical classroom. With online access, digital teachers could be accessible to students in any environment or geography, helping to bring down barriers to attending school.

Digitising policymaking

Policy creation assistant

Al can be used to search large volumes of policy documents and output natural language responses to user queries in complex policy environments.

ISSUE/OPPORTUNITY

Because the data relevant to government agencies is stored in different locations and formats, it can be difficult for analysts and policymakers to effectively query datasets and retrieve relevant information in a timely manner. With nomenclature issues,

it can also be challenging to identify associated data topics and types. The result is a diminished ability to digitise policymaking and discussion, as well as increasingly complicated interactions around policy matters.

HOW AI CAN HELP

Al assistant

Al can identify data related to common themes and topics and then summarise that information in response to user queries, helping to identify policy differences, conflicts, and gaps.

Citizen engagement in policymaking

Using AI, governments can create interactive platforms and chatbots that encourage citizens to participate in policymaking discussions. The AI-driven interface can gather public opinions and feedback on policies, making it easier for citizens to voice their views.

























Digitising policymaking

MANAGING RISK AND PROMOTING TRUST



Private

Some of the data relevant to policy issues may be sensitive or restricted, and the AI model may require controls to limit which users can access which datasets.



Fair and impartial

Various stakeholders aim to influence policymaking. Al might be biased in giving higher weight to comments and input coming from some sources over others. This has the potential to produce biased policies that favour certain businesses or segments of society.

POTENTIAL BENEFITS

Data query at scale

By reviewing large volumes of policy documents, AI can help users accelerate information gathering and increase their capacity and efficiency in querying disparate datasets.

Participatory policymaking

Using AI to better identify and incorporate a varied views and stakeholders supports more robust and complete representation in policy matters.













Drafting contracts and SoWs

Procurement

Al can analyse offerings from existing vendors, match them to organisational needs, generate requests for proposals, and analyse the responses.

ISSUE/OPPORTUNITY

Governments procure billions of dollars in goods and services annually. Traditionally, government procurement requires significant volumes of paperwork, which can lead to delays. Many government procurement contracts are highly detailed and often incorporate a range of clauses

and requirements from payment terms to export controls to wage and workforce requirements. Drafting requests for proposals (RFPs) and contracts and then generating statements of work (SoWs) requires significant time and resource investments.

HOW AI CAN HELP

Automated drafting

Al can automate the RFP and SoW writing processes by generating the initial drafts based on templates, historical documents, or specific prompts provided by procurement officials.

Extracting information

Al's advanced Natural Language Processing (NLP) capabilities can help extract relevant clauses and requirements from existing contracts, SoWs, and legal documents. Such information can be used to either create new contracts or assess the risks posed by existing contracts.













Drafting contracts and SoWs

MANAGING RISK AND PROMOTING TRUST



explainable

Transparent and Al may not be able to explain why certain clauses are added to a contract while others are excluded, which is vital information for the human user validating the outputs.



Private

Ingesting existing and historical contract data may pose data privacy and legal hurdles. Model governance is necessary to ensure the Al model, as well as the underlying data, meet privacy rules, regulations, and standards.

POTENTIAL BENEFITS

Time savings

Creating initial document drafts with Al can expedite the writing process and lead to significant time savings, compared to manually creating each RFP or SoW from scratch.

Improved consistency

Al can develop drafts while adhering to predefined guidelines in prompts, which supports a greater level of consistency across report writing.

Multilingual citizen services

Service delivery

Al can help with language translation to support the delivery of services to citizens.

ISSUE/OPPORTUNITY

Many governments around the world serve populations with varying language proficiency and linguistic backgrounds. This challenges agencies to develop multilingual websites, translate official documents, and support frontline workers with translation tools so they can better communicate with all citizens.

HOW AI CAN HELP

Aiding frontline workers

Al can be used to create realtime audio and text messages in different languages as frontline workers interact with residents around a variety of services, such as social care, health care, and emergency response.

Translating official documents

Government agencies often handle the publication of official documents, laws, regulations, and policies. Al can streamline the translation process and help produce accurate and consistent translations.

Announcement and website translation

Government websites and public information (e.g., health and travel advisories) can be translated quickly to make essential information more accessible to the population.













Multilingual











citizen services

MANAGING RISK AND PROMOTING TRUST



Fair and imparti<u>al</u> The data used to train an AI model for use in translation may not be consistently accurate or robust across all languages, which could in turn lead to poorer translations and less access to citizen services for some language speakers than for others.



Private

The translating model may be exposed to sensitive information, necessitating steps to ensure the model does not mishandle or inappropriately divulge protected data and thus violate data privacy regulations.

POTENTIAL BENEFITS

Real-time translation

When audio or text can be translated into a multitude of languages in real time, it enables more seamless and conversational interactions with different language speakers.

Translation at scale

Al can handle large volumes of document translation, giving an agency capacity to ensure government information and services are accessible to a greater audience.













Summarising legislative documents

Legislative administration

Al can help legislative staff more rapidly transcribe and summarise hearings, legislation, documents, and official announcements.

ISSUE/OPPORTUNITY

Legislative offices are expected to hold hearings on important topics, respond to constituents, and make public announcements in the form of press releases. Manually transcribing hearings and meetings is a time-consuming task. Further, developing new legislation (where staff play a pivotal role in research) requires sifting through voluminous policy proposals and research published by experts.

HOW AI CAN HELP

Summarising official documents

Auto-generating transcripts of hours-long committee hearings and summarising important bills and hearings can significantly reduce the administrative burden on staffers.

Process and summarise policy proposals and research

Legislative staff review a large volume of policy proposals and recommendations published by experts. Al can quickly summarise the documents for them, so staffers can spend more time on higher level policy analysis and decision-making.

The Al Dossier | Government & Public Services

Summarising legislative documents

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

Al may perpetuate latent biases based on its training set and generate skewed summaries that are partisan and favour certain ideologies.



Private

Ingesting internal policy proposals can expose sensitive information, requiring organisations to take measures that protect the confidentiality of internal documents.









Reducing burdens

POTENTIAL BENEFITS

Generating summaries of official hearings can reduce administrative burdens on legislative staff so they can focus on more complex tasks.



Saving time

Al can quickly retrieve information and summarise it, saving legislators and staff time when reviewing lengthy, complex, or detailed documents.





Global policy tracking

Automated tracking and analysis of public policy

Organisations in the public and private sector can use AI to monitor, interpret, and analyse public policy developments in real time across hundreds of countries.

ISSUE/OPPORTUNITY

Tracking national policy developments on a global scale is a resource-intensive and highly fragmented process. Policy documents vary in language, structure, formatting, and accessibility, making it difficult for international organisations, governments, and advocacy groups to maintain a coherent and timely view of global policy trends.

HOW AI CAN HELP

Data collection and analysis

Al can automatically gather, structure, and analyse vast volumes of policy documents from government websites and public sources around the world. The technology performs multilingual data extraction, applies natural language processing to categorise and summarise policies, and synthesises insights into structured outputs that can be validated by subjectmatter experts.













B











Global policy tracking

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

The AI model should be designed to avoid reinforcing systemic biases. Human experts from different backgrounds and regions are embedded in the feedback loop to validate model outputs, helping to ensure representation across geographies and policy contexts.



Robust and reliable

Automated systems should undergo rigorous, iterative testing to help ensure the reliability of outputs. Policy insights are continuously benchmarked against human analysis and real-world policy documents to maintain a high level of accuracy and dependability, especially in politically sensitive or underreported regions.

POTENTIAL BENEFITS

Greater accuracy, scalability and knowledge sharing

Al can enable ongoing, real-time monitoring and analysis of thousands of policies across hundreds of countries without an exponential increase in manual effort. It can also enable local entities and other stakeholders to identify global policy trends, compare regional approaches, and uncover best practices.

Improved efficiency

Combining data collection and policy analysis with Al-powered automation can save organisations significant time and resources.

Improved decision-making

Use of AI for policy tracking and analysis can provide organisations with more timely, structured, and reliable data to support effective decision-making and planning.

The Life Sciences & Health Care Al Dossier















The Life Sciences & Health Care Al Dossier

Al is accelerating transformation across life sciences and health care – sectors where precision, timeliness, and trust can mean the difference between life and death. From research labs to hospital wards, AI is enabling faster discovery, more targeted treatments, and more efficient delivery of care. With growing volumes of clinical, genomic, and real-world data, AI is helping organisations extract insights at a speed and scale that traditional methods cannot match.

In life sciences, AI is streamlining drug discovery and development - helping to identify promising new molecules, optimising clinical trial design, and identifying patient cohorts for precision therapies. Al is also being used to automate regulatory compliance and to optimise supply chain performance through improved supply and demand forecasting.

In health care, AI is improving diagnostic accuracy, personalising treatment pathways, and easing administrative burdens. Increasingly, agentic AI is automating complex workflows supporting clinical decision-making, coding and billing, and even orchestrating elements of care delivery with minimal human intervention.

As the technology matures, AI will not just enhance isolated functions – it will likely become integral to advancing scientific discovery, improving population health, and building more resilient, responsive health systems with ongoing human oversight. And Al's capabilities will only improve in the future.

Al is helping

organisations

cannot match.

extract insights at a

speed and scale that

traditional methods

That said, the use of AI in regulated, highstakes environments demands rigorous validation, explainability, and alignment with ethical standards. Data privacy, equity in outcomes, and trust in automation are paramount. To succeed, companies will need to effectively integrate Al within existing care and research frameworks while navigating evolving regulatory requirements and public expectations.

Note: The tags below each use case indicate its primary business function and whether Agentic AI is used

Agentic Al

Primary business function













Multi-modal diagnosis and clinical decision support

Collaborating across data types to improve diagnostic accuracy

Agentic AI systems can act as a clinical co-pilot, with specialised agents analysing different data sources – such as imaging, lab results, and patient records – and then combining insights to support more accurate and timely diagnoses.

ISSUE/OPPORTUNITY

Diagnosing complex conditions often requires synthesising data from multiple modalities: radiology scans, lab tests, patient histories, genetic profiles, and physician notes. This process traditionally depends on coordination across multiple human specialists, each interpreting one slice of the evidence. Limited time, high patient volumes, and fragmented systems make it difficult to integrate all relevant data into a cohesive picture, which can delay treatment or lead to missed diagnoses.

As health systems face mounting pressure from clinician shortages and rising patient loads, the need for decision support that is both comprehensive and efficient is urgent. Multi-agent Al provides an opportunity to automate routine analysis across diverse data sources, highlight patterns that humans might overlook, and generate integrated recommendations that help physicians make faster and more informed decisions.

HOW AI CAN HELP

Specialised data analysis

Specialised AI agents can apply deep expertise in focused areas. For example, imaging agents can flag abnormalities on scans, NLP agents can parse physician notes and pathology reports, and other agents can interpret lab results and genetic markers.

Contextual integration

A planning or coordination agent can link evidence across modalities – for example, connecting a lung nodule on a CT scan with smoking history and lab findings – before proposing next steps such as a biopsy or targeted treatment.

Continuous learning from records

Agents can draw on the electronic health record as common context and shared memory, adapting recommendations over time based on a patient's history and treatment responses (in addition to population-level outcomes).

Workflow automation

By triaging scans, pre-annotating records, and uncovering key insights, agents can reduce the time physicians spend on routine analysis, allowing them to focus on patient interactions and issues requiring complex human judgment.















Multi-modal diagnosis and clinical decision support

MANAGING RISK AND PROMOTING TRUST



Safe and secure Errors in diagnosis can be life-threatening. As such, agents should be deployed with strong safeguards, including rigorous validation of outputs and controlled simulation testing before integration into live workflows.



Transparent and explainable

Because clinical decisions carry major consequences, agents should provide transparent reasoning, including references to underlying data. Validator agents can be used to audit and cross-check other agents' conclusions.



Private

Medical data is highly sensitive, which means that agent systems must follow strict access controls, comply with applicable New Zealand privacy and health regulations and similar regulations, and use privacy-preserving techniques to minimise unnecessary data exposure.



Responsible and accountable

Since liability in clinical settings is complex and risky, agent recommendations should initially be deployed under physician supervision, with final accountability for diagnosis and treatment decisions retained by human clinicians.

POTENTIAL BENEFITS

Faster, more accurate diagnoses

Automated integration of multimodal data helps generate insights earlier and reduces diagnostic errors, enabling quicker, more effective interventions.

Reduced clinician workload

By offloading burdensome, data-heavy tasks such as image triage and report synthesis, agents give physicians more time for patient care and complex decisionmaking.

Improved patient outcomes

Earlier detection and personalised treatment plans can lead to better survival rates, reduced complications, and higher overall quality of care.













Hyper-personalised health care

Delivering tailored, continuous care through multi-agent collaboration

Agentic AI systems can act as a 24/7 virtual care team, with specialised agents monitoring patient data, providing personalised coaching, and coordinating interventions to deliver hyper-personalised health care at scale.

ISSUE/OPPORTUNITY

Managing chronic conditions such as diabetes, hypertension, or heart disease requires continuous monitoring, timely adjustments, and patient engagement outside of clinical visits. Human care managers and clinicians often have limited bandwidth to track individual patients in real time, leaving care gaps where small issues can escalate into serious and costly hospitalisations or complications.

Patients, meanwhile, often struggle

to follow generalised care plans that don't adapt to their lifestyles, personal preferences, or daily health fluctuations. The result is higher readmission rates, preventable emergency visits, and reduced patient satisfaction and health outcomes. Multi-agent Al offers a way to make health care more proactive and personalised, addressing issues early and potentially scaling up to provide individualised support to large populations

HOW AI CAN HELP

Continuous patient monitoring

A monitoring agent can analyse real-time data from wearable devices, glucometers, or other sensors, detecting subtle patterns such as recurring glucose spikes or irregular heart rhythms.

Evidence-based insights

A knowledge agent can stay updated with medical guidelines and population-level data, interpreting broader patient patterns and suggesting personalised, evidence-based adjustments to care.

Personalised coaching

A coaching agent can engage patients directly through apps or chatbots, offering tailored advice, timely reminders, and personal encouragement that reflects their specific health needs and behaviours.

Automated, coordinated care

Agents can schedule follow-ups, suggest medication adjustments (with clinician approval), offer healthy recipes, and even autoorder groceries – all orchestrated through a shared patient profile.













POTENTIAL BENEFITS

Improved patient outcomes at lower cost

Timely, personalised interventions can help improve patient adherence to treatments, reduce complications, prevent costly emergency visits, and enable earlier corrections when managing chronic diseases. More broadly, agentic Al systems can boost operational efficiency and clinician productivity.

Scalable personalised care

Al agents can enable clinicians to provide individualised care to thousands of patients simultaneously, increasing reach and helping to address workforce shortages.

Higher patient engagement and satisfaction

Al can help patients feel supported between clinical visits, improving adherence to treatment plans and fostering trust in their care providers.

Hyper-personalised health care

MANAGING RISK AND PROMOTING TRUST



Biased training data could lead to less effective guidance for certain demographic groups. To minimise problems and avoid reinforcing existing disparities, Al agents should be trained and validated on diverse patient datasets.



Private

Because agentic AI systems continuously collect and analyse sensitive patient information, they must provide strict privacy protections and fully comply with data regulations such as applicable New Zealand privacy and health regulations. This includes obtaining clear patient consent.



Transparent and explainable

Patients and clinicians could initially be sceptical of following Al-driven guidance. To help build trust, Al agents should provide understandable explanations for their recommendations, including specific references to the data or guidelines that informed them.



Responsible and accountable

Clinical decisions can have major safety implications. As such, Al agent outputs should be reviewed by care teams where appropriate, with final accountability for treatments and adjustments retained by licensed clinicians.

Automated customer service

Improving health plan member interactions through multi-agent systems

Agentic AI systems can handle a wide variety of customer service tasks – such as benefit inquiries, claims status updates, and enrolment support – seamlessly escalating complex cases to human representatives.

ISSUE/OPPORTUNITY

Health plans face large volumes of inquiries from members and prospective members across multiple channels, ranging from routine coverage questions to complex claims issues. Traditional call centres and support teams often struggle with long wait times, inconsistent service quality, and high staffing costs, all of which can frustrate members and erode customer satisfaction.

In today's competitive marketplace, a poor service experience can increase churn, while inefficient processes can drive up operating costs through the roof. Health plans need a way to deliver fast, accurate, and personalised support without continuously scaling human staffing levels, particularly during high-volume periods such as open enrolment.

HOW AI CAN HELP

Specialised AI agents

Individual agents can handle specific tasks such as answering benefit questions, processing claims inquiries, and assisting with new member onboarding, providing responses that are timely, accurate, and consistent.

Seamless collaboration

Al agents can share context across tasks, improving their service quality and allowing for smooth escalation to human representatives when needed, without requiring members to repeat their information.

Dynamic prioritisation

Requests can be triaged based on urgency, member status, or complexity, ensuring that highpriority cases receive timely attention.

Proactive engagement

Al agents can initiate reminders or follow-ups, such as prompting prospective members with enrolment information or reminding existing members of required documentation.

























Automated customer service

MANAGING RISK AND PROMOTING TRUST



Private

Health plan support requires handling sensitive personal and medical information. Agentic Al systems must be designed with strict data protection protocols and full compliance with applicable New Zealand privacy and health regulations



Robust and reliable

Since members rely on accurate information to make critical financial and health decisions, Al-generated responses should be validated against authoritative data sources and continuously monitored for accuracy.



Responsible and accountable

Health-related interactions often require a strong dose of empathy and humanity. Al agents should be implemented as a complement to human staff, with clear escalation paths to ensure that members always have access to human support when needed.

POTENTIAL BENEFITS

Improved member experience

Faster, more accurate, and more consistent responses across channels help improve the overall customer experience, build member trust and satisfaction – and reduce churn.

Increased operational efficiency

Automating routine tasks frees up human staff to focus on complex or high-priority cases, boosting effectiveness while at the same time improving productivity and efficiency.

Enhanced scalability

Agentic AI can help health plans maintain consistent levels of service and responsiveness in times of fluctuating inquiry volumes (e.g., during open enrolment) without proportionally increasing staffing.

Smarter clinical trials

Improving trial design, recruitment, and monitoring with multiagent systems

Agentic Al systems can optimise clinical trials by designing protocols, selecting sites, identifying participants, and monitoring progress in real time, helping trials run faster, safer, and more efficiently.

ISSUE/OPPORTUNITY

Clinical trials are one of the most expensive and resource-intensive stages of drug and therapy development, often slowed by challenges such as inefficient site selection, poor patient recruitment from narrow patient pools that don't fully reflect the broader population, and frequent protocol amendments. These inefficiencies can lead to delays, higher costs, and a higher risk of trial failure – ultimately impacting the speed-to-market and efficacy of new treatments.

In the face of growing therapeutic pipelines and increasing pressure to reduce time-to-market, life sciences organisations need tools that can improve trial design and execution while maintaining compliance with strict regulatory standards.

HOW AI CAN HELP

Optimised site and protocol design

Al agents can analyse hospital capabilities, historical trial data, and patient access to recommend optimal trial sites, appropriate sample sizes, and robust protocols that balance scientific rigor with operational feasibility.

Data integration and harmonisation

By combining data from electronic health records, registries, and wearable devices, agents can create a holistic, real-time view of trial operations, ensuring decisions are made with the most current and complete information.

Adaptive trial management

Agents can continuously monitor trial progress and outcomes, dynamically adjusting resource allocation or protocol details to keep trials on track and compliant.

Patient stratification and recruitment

Specialised agents can identify and stratify eligible patient populations, improving recruitment efficiency while reducing delays caused by insufficient enrolment.













Smarter clinical trials

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

Al models should be trained and validated using demographically broad datasets to avoid systemic bias in recruitment strategies. Representation audits can help ensure no population segments are inadvertently favoured or excluded.



Private

Since clinical trials depend on sensitive patient data, agents must be designed with strong privacy protections, data governance controls, and full compliance with applicable New Zealand privacy and health regulations.



Robust and reliable

Trial outcomes directly impact patient safety and therapeutic approval. Al-generated recommendations need to be validated against clinical expertise and continuously refined through real-world feedback.



Responsible and accountable

Over-reliance on automation can lead humans to become complacent and miss important clinical nuances. Al agents should be treated as decision-support tools with physicians and trial managers maintaining final accountability.

POTENTIAL BENEFITS



Improved trial design and adaptive management reduce delays and inefficiencies, driving down the high costs associated with clinical development. Also, real-time recruitment monitoring increases the likelihood of trial completion.



By including a broader and more demographically representative participant pool, clinical trials are more likely to produce results that reflect the real-world effectiveness of a drug across different populations.

Accelerated time-to-market

Smarter site selection, faster recruitment, and fewer protocol amendments shorten trial timelines, enabling faster delivery of new therapies to the marketplace.

Better commercial outcomes

Drugs that are tested and proven effective across a wide range of demographic groups are more likely to gain broad clinical adoption, which increases their commercial potential.



























Dynamic inventory management

Using coordinated AI agents to optimise supply chains and inventory

Agentic AI systems can transform supply chain and inventory management through specialised agents that help optimise inventory levels and supply chain performance while reducing costs.

ISSUE/OPPORTUNITY

Managing inventory in life sciences and pharmaceuticals is particularly challenging due to complex supply chains, strict regulatory requirements, and unpredictable demand for critical products. Traditional inventory management often relies on fixed rules and periodic actions, which can lead to costly overstocks, stockouts that disrupt patient care, and inefficiencies in production and procurement.

In addition to regulatory and operational complexity, companies must also navigate factors such as geographic variations in demand, perishability of products, and external disruptions like transport delays or weather events. These challenges can result in shortages with serious health implications, or excess inventory that drives up costs and waste. To remain competitive and compliant, organisations need inventory management systems that can respond dynamically to real-time changes in demand, supplier performance, and production capacity – minimising waste and ensuring that essential therapies and products are always available.

HOW AI CAN HELP

Demand forecasting agents

Al agents can analyse historical sales data, market trends, and external signals to help predict demand fluctuations and proactively adjust inventory levels and plans. Al can also integrate local geographical characteristics, disease prevalence, and socioeconomic factors to generate highly accurate, micro-market-specific demand forecasts.

Supplier and procurement agents

Agents can monitor supplier performance, lead times, and procurement activities, ensuring purchasing decisions align with both short-term needs and longer-term forecasts.

Production and scheduling agents

Al can optimise production schedules, identifying bottlenecks and aligning output with available resources and expected demand. By leveraging cross-silo data and advanced analytics, Al can identify patterns and trends that traditional methods might miss, helping to mitigate product shortages.

Distribution and logistics agents

Agents can monitor distribution flows, detect inefficiencies, and propose adjustments in real time to help ensure inventory is positioned where it is most needed.

Coordinated orchestration

An orchestration layer confirms that actions across procurement, production, and logistics remain aligned, enabling holistic, end-to-end inventory optimisation.













Dynamic inventory management •

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Poor forecasts or system errors can trigger significant supply chain disruptions. Al agents should be validated against human planners and real-world performance, with continuous monitoring to ensure accuracy and stability.



explainable

Transparent and Inventory decisions can have a major impact on costs, compliance, and patient access. As such, agent-driven decisions should be supported by clear rationales that supply chain professionals can understand and verify.



accountable

Responsible and Automated systems can overlook subtle operational factors that a human expert might consider crucial. Conversely, over-reliance on automation can make human experts complacent. Human oversight of Al is essential, with supply chain professionals empowered to override or adjust agent recommendations as needed. 🗖

POTENTIAL BENEFITS

Reduced costs and improved financial performance

Dynamic, data-driven inventory planning powered by agentic AI can lower inventory holding costs, reduce excess stock, and minimise spoilage of perishable products.

Improved service levels

Real-time monitoring and coordinated decision-making reduces the risk of stockouts, ensuring patients and providers receive products when needed.

Greater supply chain agility

Organisations can respond more quickly to disruptions or demand spikes, maintaining resilience in the face of changing market or regulatory conditions.

B











End-to-end autonomous drug discovery and development

Accelerating drug discovery through collaborative AI agents

Al-enabled workflows can accelerate aspects of drug discovery under expert oversight – from molecular design to clinical trials – using specialised agents that work together to accelerate innovation and reduce costs.

ISSUE/OPPORTUNITY

Drug discovery and development is one of the most expensive and time-consuming challenges in life sciences, often costing billions of dollars and spanning many years or even decades. Researchers must manually generate candidate molecules, predict their properties and efficacy, assess toxicity, and evaluate manufacturability (often across disconnected systems and siloed teams). This fragmented approach slows innovation, drives up costs, and increases the risk of late-stage failures.

With growing demand for faster therapeutic innovation and mounting R&D pressures, life sciences organisations need ways to both increase the throughput of candidate discovery and improve confidence in early-stage selections. Multi-agent AI offers an opportunity to streamline molecule design, integrate diverse data sources, and automate iterative refinement of promising compounds.

HOW AI CAN HELP

Agent-driven design of novel compounds

Specialised AI agents can generate novel chemical structures, simulate compound interactions, and predict pharmacological properties, providing a broad and high-quality candidate pool. AI can also propose modifications to known molecules in early-stage drug development, rapidly evaluating their therapeutic potential and feasibility.

Multi-modal integration

Al systems can incorporate data from different domains – including chemical, biological, and safety data – creating a holistic view of each compound's potential and reducing the risk of pursuing weak candidates.

Collaborative refinement

Toxicity, manufacturability, and regulatory-focused agents can evaluate compounds in parallel, sharing data with molecule generation and simulation agents to iteratively improve designs. An orchestration agent can oversee and manage the specialised agents, integrate outputs, and prioritise workflows, facilitating seamless progression from initial ideation to candidate selection.

Clinical trials and regulatory submissions

Agents can design preclinical studies and trial protocols based on predicted outcomes. Also, agents can prepare submission-ready documentation, track compliance requirements across regions, and adapt protocols to evolving regulatory standards.













End-to-end autonomous drug discovery and development

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Early design errors can lead to costly downstream failures. Al outputs should be validated against human benchmarks and continuously monitored.



explainable

Transparent and Regulatory approval requires meeting strict standards for safety, efficacy, and traceability. Agent reasoning and decision paths should support explainable outputs, including transparency into the data and assumptions used.



Responsible and accountable

Drug discovery is a complex challenge with life-or-death implications. Al agents should be deployed as co-pilots to human scientists, with humans retaining accountability for final go/no-go decisions.

POTENTIAL BENEFITS

Accelerated time-to-market

Automated generation and evaluation of compounds can significantly shorten earlystage development timelines – from years to months - enabling rapid responses to emerging diseases and other unmet needs.

Lower R&D costs

By rapidly identifying and prioritising candidates that meet key property constraints, AI can reduce development costs and improve the quality of compounds advancing to later stages.

Higher likelihood of success

Comprehensive, multi-modal analysis and real-time adaptation improve the chances that selected candidates will advance successfully through clinical trials. The results could include reduce late-stage attrition, improved ROI, and timely delivery of safer, more effective therapies.

A co-writer for appeals

Denial appeal letters

Al can be used to draft denial appeal letters, drawing from patient records and medical policies and guidelines in a faster, more cost-effective way.

ISSUE/OPPORTUNITY

When a medical insurance claim is denied, hospital billing staff face a costly and lengthy process of reviewing patient records and medical policies to create an appeal letter. Part of the challenge is the amount of time required for staff to compile an appeal.

HOW AI CAN HELP

Retrieving policies and guidelines

An Al-based retrieval model can reach across large volumes of medical policies and member plans to identify the necessary information for a claims appeal.

Extracting patient data

Using extractive algorithms, the organisation can rapidly consult unstructured medical notes, medications, lab results, and other electronic health records.

Writing the appeal

With the necessary information gathered by Al, an LLM can be used to generate an appeal letter.

























A co-writer for appeals

MANAGING RISK AND PROMOTING TRUST



accountable

Responsible and When consulting highly detailed guidelines, policies, and records to appeal a claim denied for vague reasons, Al models may misinterpret the denial or the records, leading to an unsuccessful appeal. Ultimately, a human needs to be accountable for validating appeal letters.



Private

By drawing from electronic health records, the model is consuming health information whose protection is subject to laws and regulations. Organisations must ensure that the data ingested and information outputted aligns with data protection and patient privacy expectations.

POTENTIAL BENEFITS

Reclaim revenue

Automating the denial appeal process can supplement hospital billing resources, leading to more denial appeals filed and potentially more revenue recovered.

Efficiency improvement

Compared to traditional manual methods, Al technologies can enhance the speed and efficiency of appeals activities, such as substantiating claims and drafting appeals. Their potential to streamline processes for both simple and complex cases can make the legal workflow more timely and cost-effective.

Faster admin for payers, providers, and patients

Accelerated prior authorisation

Using AI to analyse medical policies, guidelines, and provider-submitted information about underlying issues, patient needs, and medical history, an organisation can automate a Prior Authorisation submission (Provider) or generate a Prior Authorisation approval or denial (Payer).

ISSUE/OPPORTUNITY

The Prior Authorisation process is manual and labour-intensive for both health care payers and providers. The process requires the input of coders who understand the intent of a payer's Prior Authorisation policies, as well as the need for medically necessary care management plans. The

time required to analyse medical records and policies to make determinations on Prior Authorisation submission, approval, or denial can lead to a long administrative process between the payer and provider, which can negatively impact patient satisfaction and the customer experience.

HOW AI CAN HELP

Supporting the provider

For providers, AI can help prepare a Prior Authorisation submission by analysing submission requirements and guidelines and cross-referencing with a patient's medical records to ensure necessary requirements are met. AI can then aid in submission to the payer and continually learn which best practices tend to lead to Prior Authorisation approvals.

Supporting the payer

For payers, AI can help reduce the time required to make a Prior Authorisation decision, impacting the patient experience. It also helps mitigate fraud by determining if there are anomalies in a provider's coding practices and supports compliance by analysing submitted Prior Authorisation requests and records against the payer's policies and procedures.

More efficient operations

For both payers and providers, using AI can reduce work burdens and streamlines the ability to handle Prior Authorisations, which can reduce costs while improving patient experiences.













Faster admin for payers, providers, and patients

MANAGING RISK AND PROMOTING TRUST



Safe and secure

Prior Authorisation requires the provider and payer to communicate sensitive patient data, such as protected health information (PHI) and personally identifiable information (PII), etc., which means this data is exposed to the model. Risks include unauthorised thirdparty access, as well as AI systems inadvertently revealing sensitive information during the generation process, thus compromising patient data confidentiality.



Fair and impartial

The process for submitting and responding to a Prior Authorisation involves a standard set of rules and the patient's medical history, which introduces the potential for bias in Al models. This bias might arise from the historical data used to train the model (e.g., disparities in health care treatment or outcomes), which could lead the AI model to inadvertently perpetuate and even amplify such biases by making its own biased decisions or recommendations. The use of standardised authorisation rules and patient-specific medical history, alongside continuous monitoring and careful evaluation, helps mitigate this risk and promotes fairer outcomes.



Robust and reliable

While the process for submitting and responding to a Prior Authorisation revolves around a standard set of rules and the patient's medical history, there is a risk the model will misinterpret nuanced medical conditions of underrepresented populations that were not in the training dataset, and thus falsely deny the need.

POTENTIAL BENEFITS



With AI, providers and payers may require less time to understand policies, research patient medical records for compliance, and generate, approve, or deny a Prior Authorisation request.



An AI feedback loop refers to the cyclical process whereby the AI model's output is presented to users or evaluators for feedback, which is then used to iteratively update and refine the model. This enhances the consistency and quality of outputs, enables providers to gain a deeper understanding of payer policies, streamlines decision-making processes, and helps payers optimise their procedures.

Improved patient experience

As the Prior Authorisation process becomes more efficient, patients can receive the care management they need without needless waiting for administrative processes to conclude. This can increase patient satisfaction by improving the administrative and patient experience.















Simplifying claims submission

Medical coding

Al can be used to create codes for a claims department to categorise incoming claims and billing for medical services and procedures, which can improve the accuracy, efficiency, and speed of the claims process.

ISSUE/OPPORTUNITY

The claims submission process in the medical industry can be laborious and error-prone, involving the manual categorisation of a large volume of incoming claims with complex medical codes. This time-consuming task leads to backlogs, delays, and potential payment issues for health care providers.

HOW AI CAN HELP

Transformed claims processing

Using AI to help categorise incoming claims and analyse and assign accurate codes can improve the overall accuracy, efficiency, and speed of claims processing. This can result in faster reimbursements for providers and a streamlined experience for both the claims department and patients.

Reduced labour burden

By leveraging AI, human effort in the claims submission process can be redirected to higher value tasks, which could result in administrative cost savings for the payer.

























Simplifying claims submission

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

A Large Language Model (LLM) used in medical billing may be susceptible to bias from skewed training data, incorrect labels, and under-represented cases, potentially leading to incorrect claim categorisation. To mitigate these issues, careful data collection, diverse model testing, and continuous monitoring and adjustment are vital for ensuring fair and accurate performance.



Private

To assess coding accuracy, the LLM compares the billed codes with the patient medical history, which exposes the patient's data to the underlying model and creates potential privacy risks that need to be mitigated.



Robust and reliable

Medical coding is highly regulated with strict penalties for over/under coding. The accuracy and reliability of LLM outputs in this regard is essential, as mistakes could have serious consequences. Reliability may be challenging in part because patient medical history may contain multiple modalities (e.g., text, images, and video).

POTENTIAL BENEFITS

Accuracy to limit revenue loss

Leveraging an LLM can help reduce the risk of coding errors. This can increase billing accuracy and decrease revenue loss due to errors.

Time efficiency

Automating the review of medical records can save valuable time for health care practitioners, enabling them to focus on more meaningful work.

A physician's message manager

Provider in-basket management

An LLM can be used to process messages in a health care provider's in-basket, accelerating responses while enabling physicians to focus on patient-facing care.

ISSUE/OPPORTUNITY

The amount of time required for primary care providers (PCPs) to accomplish both administrative and patient care responsibilities can exceed what is possible in a day. In some cases, upwards of two-thirds of their time is spent on administrative, non-patient facing work.

The 21st Century Cures Act encourages electronic medical records (EMR) in-basket usage, which led to a dramatic increase in in-basket messages during the COVID-19 pandemic. The result is a significant burden on PCPs, which is contributing to physician burnout.

HOW AI CAN HELP

Triaging the in-basket

Al can be used to review routine messages (e.g., Rx refills, scheduling) and delegate simpler tasks to automation.

Message assistant

PCPs can leverage AI to summarise complex clinical messages for review and use the model to draft replies for provider input and response. AI models consult prior in-basket replies and EHR data when creating drafts.

Insights at scale

By using Al-enabled in-basket message systems at scale, organisations can identify issues related to patient negativity in their messages. The insights into complaints, expressions of dissatisfaction, frustration, confusion, or concern about care can inform interventions that may improve the patient experience.













A physician's message manager

MANAGING RISK AND PROMOTING TRUST



Safe and secure

Use of AI for in-basket systems involves collecting, processing, and storing large amounts of sensitive patient data, such as medical history, diagnoses, and treatment plans. This data is subject to strict privacy laws, and any unauthorised third-party access could result in legal and financial consequences for health care providers.



Responsible and accountable

If messages are composed or summarised with inaccurate information, it could lead the PCP to erroneous decision-making or poor patient engagement, which can have significant consequences for patient health, trust in the health care provider, and the reputation of the organisation.

















POTENTIAL BENEFITS

Physician support

By using an Al-enabled in-basket system, the PCP's time-consuming administrative tasks are reduced, permitting more patient-facing work and mitigating one cause of physician burnout.

Timely responses

A more efficient process for working through in-basket messages can lead to faster responses to patient needs, contributing not only to a better patient experience but potentially also better health outcomes.

Patient sentiment

By identifying and tracking signals of negativity at scale, health care providers can gain insights into common pain points in the patient experience. This could help them proactively address these issues, whether by adjusting their practices, improving communication, or implementing other interventions to enhance patient satisfaction.













Democratising AI model creation

Knowledge domain model development

Generative AI can be used to improve existing AI models by removing user interface (UI) hurdles through reinforcement learning (RL) without the need for technical staff.

ISSUE/OPPORTUNITY

Developing novel AI models for life sciences and health care organisations continues to demand a high degree of technical proficiency to perform data exploration, feature engineering, model training, and evaluation. Frequently, the steps involved in model training lack a user-friendly interface, posing accessibility challenges for health care professionals and domain experts who may not

possess extensive technical backgrounds. Simultaneously, the quality and relevance of model outputs hinges significantly on domain expertise and practical experience. Overcoming this divide between technical acumen and domain knowledge remains an obstacle in harnessing the complete capabilities of AI within the field of life sciences and health care.

HOW AI CAN HELP

Empowering professionals

With its capacity for learning from and adapting to iterative feedback, generative AI can act as an enabler for professionals across various sectors. It offers the opportunity to continually refine existing domain-specific AI models by adding new training data. This iterative enhancement increases the model's accuracy, utility, and relevance to the user's specific professional needs. In this way, generative Al can empower professionals by providing them with tailored, precision AI tools that evolve with their work.

Streamlining health care model development

Generative AI can help simplify model development in the complex and highly regulated health care industry. By focusing on intuitive user interface designs and automated processes, generative AI minimises UI obstacles, making it more

accessible for professionals to refine and improve their existing models.

Improving alignment

Generative Al leverages reinforcement learning (RL) techniques, a type of machine learning where an AI system learns to make decisions by trial and error, to validate and improve its own outputs. This process assists in mitigating prevalent AI challenges, including hallucinations or confabulations, ambiguity, and colloquialism misuse. As a result, it bolsters Al's reliability and furnishes professionals with more precise models and predictions, thus aligning AI capabilities more closely with user requirements.













Democratising Al model creation

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Hallucinations or confabulation could lead to the execution of incorrect procedures or use of suboptimal reagents and equipment, causing inaccurate experiments and inefficient use of resources. Particularly in medical or pharmaceutical labs, inaccurate information could even lead to compliance or regulatory issues.



Transparent and explainable

The generative AI system incorporates tools that offer transparency into data engineering pipelines, including data preparation stages. This inherent transparency facilitates an understanding of the AI model's functioning within the organisation, fostering trust in the accuracy and reliability of the AI system's outputs. It is a crucial component of the AI use case, demonstrating the system's accountability and promoting its acceptance across the organisation.

POTENTIAL BENEFITS

Enhance institutional knowledge access

Al can help reduce institutional knowledge loss due to employee exits and enable on-demand access to domain-specific knowledge across the organisation.

Increase development throughput

Domain area experts can drive more self-sufficient model experimentation and development by utilising natural language model outputs and synthesising insights about optimal procedures, reagents, equipment, and techniques into a comprehensive and accessible format.

Cost management

This democratic approach to AI model development empowers employees to take part in model experimentation, reducing costs associated with machine learning operations and technical specialists.













Optimising lab procedures

Experimental design

Al can be used to create procedural templates and recommendations on best practices (e.g., reagents, equipment, techniques).

ISSUE/OPPORTUNITY

Laboratory personnel, including researchers, technicians, and managers, often face challenges in maintaining up-to-date procedural templates and maintaining the consistent application of best practices, especially as scientific knowledge evolves rapidly. These challenges can lead to inefficiencies,

errors, and inconsistency in experiments or analyses. Additionally, without a central source of curated recommendations, time and resources may be wasted sourcing and comparing various reagents, equipment, and techniques. These pain points present an opportunity for AI to streamline and enhance laboratory processes.

HOW AI CAN HELP

Generation of novel processes

Leveraging historical data and scientific principles, an Al model could suggest novel experimental designs, more efficient processes, or alternate uses of reagents and equipment, stimulating innovation in laboratory procedures.

Data analysis and interpretation

Al uses a large language model (LLM) to analyse data from lab protocols, equipment specifications, previous experimental designs, reagent usage, and techniques, providing a holistic understanding of laboratory procedures and principles.













Optimising lab procedures

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

The integration of multimodal text and images of complex structures and processes in experimental design presents complexity. This can heighten the risk of unworkable, unfeasible, or inefficient designs, as interpreting and accurately representing this diverse and intricate data can be challenging. These challenges could potentially lead to errors in the design and execution of experiments, resulting in failed or less reliable outcomes and unnecessary time and resource expenditure.



Responsible and accountable

In the event of erroneous design recommendations, accountability may be an issue. Determining who bears the responsibility for incorrect designs and their potential consequences is important. The roles of human oversight and system validation need to be clearly defined.



Transparent and explainable

With the application of AI in experimental design, there may be challenges related to explainability. If scientific or academic papers are to be published based on the results, authors need to be able to adequately explain the methodology behind the AI-recommended designs, which can be inherently complex due to the black-box nature of some AI models.

POTENTIAL BENEFITS

Efficiency

LLMs can reduce the time and effort needed for experimental design by streamlining and accelerating data analysis and procedure consolidation, and by providing best practice recommendations.

Lower cost

With less time required for experimental design, organisations can reduce the overall operational costs of experiments while also increasing throughput.

Revealing the rules

Automated regulatory compliance

Al can be used to support compliance by processing large amounts of regulatory documents from multiple geographies.

ISSUE/OPPORTUNITY

Compliance with ever-changing regulations in every geography is a costly, time-consuming process for pharmaceutical companies. Even with significant investment in legal help, regulatory compliance can be hard to achieve.

HOW AI CAN HELP

Text processing

Al can be used to extract regulations for a specific purpose from thousands of pages of regulatory texts, enabling compliance.

Transforming the legal support ecosystem

As AI handles the laborious, detail-oriented process of regulatory text processing, it can also lead to a commensurate decrease in the need for third-party legal and compliance support.

Mitigating financial risk

By employing Al in regulatory compliance, the potential financial risk associated with noncompliance can be reduced.

























Revealing the rules

MANAGING RISK AND PROMOTING TRUST



explainable

Transparent and Al models may produce outputs that are hard to interpret, making it difficult to validate them and explain the reasoning to regulatory authorities.



Private

While regulatory authorities may vary, data privacy around personal health information remains a priority, and data that is not anonymised first may leak and become inappropriately disclosed.



Robust and reliable

An Al model trained to extract compliance factors from regulatory documents may be susceptible to outputting information that looks accurate but is a hallucination, making human validation an important element for mitigating risks around reliability.

POTENTIAL BENEFITS

Cost reduction

Using AI to process regulatory documents reduces the need for humans to perform time-consuming tasks, potentially lowering the cost of compliance.

Fuel for growth

When regulatory compliance becomes tractable across geographies because of AI processing capabilities, it helps the organisation confidently expand business operations globally.

20/20 impurity detection

Al-driven visual inspection for particulate matter in IV fluids

Computer vision powered by AI can be used to detect particulate contamination in IV bags, reducing product waste and improving patient safety in life sciences manufacturing.

ISSUE/OPPORTUNITY

Pharmaceutical manufacturers, particularly those producing IV fluids and life-saving therapies, face a persistent and costly challenge: detecting particulate matter in sterile products. Despite sterile manufacturing environments, small particles – such as plastic, dust, or other foreign materials – can still enter IV bags, posing serious health risks to patients.

This issue is not new – dating back to the 1940s – and despite ongoing improvements, a scalable, reliable, and cost-efficient solution has remained elusive. Historically, detection has relied on manual inspection, often using contingent labour, leading to inconsistent results, high labour costs, and significant product waste. Also, every incident of contamination risks brand reputation, regulatory scrutiny, and potential product recalls.

HOW AI CAN HELP

Automated, real-time inspection

By combining Al vision capabilities with generative Al models trained on synthetic and real-world data, manufacturers can automate the inspection process at scale. High-resolution imaging and computer vision detect anomalies in fluid packaging with greater precision and consistency than the human eye.

Continuous learning

Al enhances the system by learning from historical defect data, adapting to new defect types, and identifying potential causes through pattern recognition across vast datasets. The AI system not only flags potential contamination in real time but also enables traceability - helping identify root causes by analysing patterns across manufacturing lines, geographies, or specific production lots. This insight enables proactive correction and long-term process improvements.

























20/20 impurity detection

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Al models should undergo rigorous testing across multiple manufacturing lines and environments to help ensure high accuracy and minimal false positives/negatives. Redundancy checks, human-in-the-loop validation, and performance monitoring help ensure reliable operation even under variable lighting or packaging conditions.



Transparent and explainable

Al-based contamination detection can provide clear, trackable results, allowing manufacturers to understand why a product passes or fails inspection. Detailed imaging and reporting help ensure accountability and regulatory compliance while enhancing confidence in quality control.

POTENTIAL BENEFITS

Improved patient safety and confidence

Consistent detection of contaminants before they enter the supply chain increases patient safety, and fewer quality incidents enhance brand trust with hospitals, regulators, and patients.

Operational efficiency and scalability

Replacing manual inspection with AI reduces reliance on contingent labour and speeds up quality control processes. Also, AI-based solutions can be deployed at scale across multiple products, manufacturing lines, and facilities worldwide.

Improved ROI and reduced costs/waste

Al helps drive measurable ROI through improved efficiency, reduced labour costs, and avoidance of costly recalls. Also, fewer discarded bags due to false positives or late-stage detection leads to significant material savings.

Quality and traceability

Root cause analysis helps address upstream issues in the manufacturing line, improving overall process quality.

The Technology, Media & Telecommunications Al Dossier

















The Technology, Media & Telecommunications Al Dossier

Al is reshaping the technology, media, and telecommunications sectors from the inside out. For businesses providing the digital foundation of the modern economy, AI is driving both product innovation and operational reinvention. From code generation and content creation to network optimisation and audience analytics, AI is accelerating time to market, enhancing user engagement, and unlocking new revenue models.

In the technology sector, AI is transforming how software is designed, built, and maintained. Developers are using Alpowered tools to write, test, document and debug code faster, while IT operations teams are deploying intelligent agents to monitor systems, predict outages, and automate resolution. In hardware, Al is enabling smarter design of chips and devices, shortening development cycles and improving performance.

Media organisations are leveraging generative AI to create hyper-personalised content, automate editorial workflows, and synthesise massive volumes of user

data into actionable insights. Increasingly, agentic Al systems are moving beyond task automation – autonomously managing workflows, supporting technical sales, and orchestrating elements of content management across the industry landscape. Meanwhile, synthetic voice, image, and video capabilities are opening new formats for storytelling and user interaction.

In telecommunications, AI is enhancing technician support, streamlining network maintenance, and optimising operations - helping providers deliver more reliable service and respond faster to customer needs. As 5G expands, AI will be critical to orchestrating complex, distributed infrastructure while enabling real-time services – from autonomous vehicles to immersive media.

Yet as Al becomes embedded in core platforms and content, the stakes rise around governance, transparency, and responsible use. Issues such as IP ownership, misinformation, algorithmic bias, and deep fakes present real risks. To stay competitive and credible, companies must balance speed of innovation with safeguards that build trust.

To stay competitive

companies must

balance speed of

innovation with

safeguards that

build trust.

and credible,

Note: The tags below each use case indicate its primary business function and whether Agentic AI is used.

Primary business function

Agentic Al















Al-powered technical sales

Streamlining the end-to-end sales process with role-based AI agents

Agentic AI systems can augment sales teams by using specialised agents that mirror traditional sales roles, automate research and CRM tasks, and provide real-time support to boost sales performance and efficiency.

ISSUE/OPPORTUNITY

Technology sales are increasingly complex, requiring teams to manage long sales cycles, multiple stakeholders, and extensive product knowledge.

Sales representatives often spend large amounts of time on administrative work – researching prospects, updating CRM

systems, and preparing materials – leaving less time for relationship building and deal closure. These inefficiencies can reduce responsiveness to customers, lead to inconsistent engagement, and limit organisations' ability to scale sales operations effectively.

HOW AI CAN HELP

Prospecting and research support

An Al-based business development representative (BDR) agent can research prospects, prepare initial pitch materials, and draft outreach communications, helping to provide consistent and timely engagement.

Lead qualification and CRM automation

Specialised Al agents can automatically update CRM systems, validate customer data, and handle routine recordkeeping, reducing errors and freeing sales staff for higher-value activities.

Account executive (AE) enablement

An account executive AI agent can support the next stage of the sales process by preparing detailed materials, assisting with live customer inquiries, and coordinating documents to support the close of the deal.

Workflow integration

The Al agents can connect with sector-specific sales tools and platforms, align processes with industry best practices, and collaborate to provide seamless workflows across the sales organisation.













Al-powered technical sales

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Incorrect data entry or poor workflow execution could result in errors that undermine credibility and damage customer relationships. To minimise such problems, Al agents should be continuously monitored, validated, and tested against live sales processes.



Private

Because sales processes involve handling sensitive customer information, agents should comply with all relevant data protection regulations (such as GDPR and CCPA), minimise unnecessary data exposure, and use encryption to secure communications and CRM updates.



Responsible and accountable

While Al can support sales teams with research and recommendations, ultimate responsibility for customer relationships, pricing decisions, and contractual commitments should remain with human sales professionals and business leaders.



Safe and secure

Sales platforms and CRM systems are prime targets for cyberattack. Al agents should be integrated with enterprisegrade cyber security safeguards to prevent breaches, account takeovers, or data leaks that could damage trust with customers.

POTENTIAL BENEFITS

Increased sales productivity and lower costs

Automating research, CRM updates, and other sales support tasks can reduce costs while enabling human sales teams to spend more time engaging customers and closing deals.

Improved customer experience

Faster responses and more consistent communications improve customer satisfaction and help strengthen relationships and trust throughout the buying journey.

Scalable sales operations

Al agents can increase sales capacity without proportional increases in headcount, enabling organisations to expand their business in new and existing markets.

HOW AI CAN HELP



Customer success planning

An overarching customer success agent reviews CRM and sales data to draft initial success plans and then continuously updates them with new inputs about usage, deployments, and engagement.

Ongoing monitoring and insights

Specialised agents analyse customer health scores, product usage, and support interactions, identifying risks and opportunities for proactive intervention by CSMs.

Strategic review preparation

Another Al agent prepares first drafts of materials for quarterly or strategic business reviews, providing customer-facing teams with timely, relevant insights.

Renewal and upsell support

A renewal agent drafts a renewal playbook, updates the CRM system, prepares quotes, and recommends product or pricing mixes, then cycles information back to the customer success agent for onboarding or cross-sell support.

Conversational support

Agentic Al-powered virtual assistants orchestrate advanced language models to deliver natural, adaptive conversations - handling a wide range of customer queries in real time and multiple languages. By combining orchestration with generative Al capabilities, these agents provide fast, accurate, and consistent responses, while automatically documenting interactions for future reference. This enables scalable, personalised support at lower cost.













Al agents for customer success

Using agentic AI to improve post-sale support and customer success

Agentic AI systems can improve customer success and maximise long-term customer value by using AI to handle numerous post-sale support activities, allowing humans to focus on relationship building.

ISSUE/OPPORTUNITY

For technology companies, post-sale customer success is increasingly central to revenue growth, with a customer's longterm value to the business hinging on renewals, upsells, and cross-sells. However, customer success managers (CSMs) are often stretched thin, spending significant time on administrative work such as preparing success plans, drafting strategic review materials, and managing renewals in CRM systems. This leaves less capacity for relationship building and strategic conversations that secure customer loyalty.

As organisations scale and expand their account portfolios, they have a critical need to automate routine post-sales activities while delivering high-quality, personalised support. Agentic Al offers a way to maintain high-touch engagement - even for down-market or high-volume accounts – without proportional increases in staffing.

Customer Service

Agentic Al

The Al Dossier | Technology, Media & Telecommunications

Al agents for customer success

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Since inaccurate data or flawed insights could undermine customer trust and loyalty, Al agents should be validated against actual customer usage patterns and regularly updated with new business rules.



explainable

Transparent and For customer-facing teams to trust and act on Al-driven recommendations, agents need to provide clear rationales for their outputs and highlight supporting data.



Responsible and accountable

Strong customer relationships are built on personal empathy and trust. As such, humans will always be a critical part of the sales and support process (pre- and post-sale) and have ultimate responsibility for building and maintaining relationships that last.







POTENTIAL BENEFITS

Higher customer retention

By proactively monitoring customer health and addressing problems early, Al agents can help organisations maintain strong relationships and reduce churn.



Agents can identify upsell and cross-sell opportunities that drive account expansion and long-term revenue growth.

Lower costs and improved productivity

Automation of routine planning, reporting, and CRM updates reduces costs and frees CSMs to focus on relationship building and other high-impact activities that boost customer satisfaction and value.









HOW AI CAN HELP



Specialised coding agents

Al agents can generate code from natural language requirements, translating business or product needs directly into executable code.

Automated testing agents

Testing agents can validate functionality, identify defects, and ensure that new code integrates seamlessly with existing systems.

Debugging and troubleshooting agents

When errors are found, agents can automatically propose fixes or apply patches, quickly resolving problems and accelerating development timelines.

Deployment and operations agents

Specialised agents can handle integration, deployment, and monitoring, facilitating smoother releases and more resilient production environments.

Code summarisation and documentation

Agentic AI can automatically generate clear, human-readable documentation alongside code, improving maintainability and accelerating future development - without burdening engineers with manual write-ups.













Al agents for software engineering

Automating the software lifecycle with multi-agent collaboration

Agentic AI systems can accelerate and improve software development by using specialised agents that generate, test, debug, and deploy code, enabling faster delivery of higher-quality applications at lower cost.

ISSUE/OPPORTUNITY

Software engineering is an increasingly complex challenge, with organisations under pressure to quickly deliver new features and products while maintaining high levels of reliability, security, and compliance. Traditional development approaches require large teams of software engineers to execute repetitive

tasks such as writing boilerplate code, testing, debugging, and managing deployments. These activities are timeconsuming and costly. Also, they can impair innovation and development by diverting skilled engineers away from strategic design and problem-solving.

R&D/Product Development

Agentic AI

MANAGING RISK AND PROMOTING TRUST

engineering

Al agents for software





Faster development cycles

debugging shortens release timelines, enabling quicker delivery of features and products.

refinement can reduce defects and lead to more reliable, secure software.

By offloading repetitive software development tasks, engineers can focus on higher-value activities such as strategic design, innovation, and problem-solving.

POTENTIAL BENEFITS

Automation of coding, testing, and

Improved code quality

Continuous Al-driven testing and

Greater productivity and efficiency















Robust and

explainable

reliable

Agents must avoid exposing sensitive business logic, proprietary algorithms, or customer data during development and deployment. Data privacy protections and secure handling of intellectual property are essential, particularly when using shared or cloud-based training data.

In software development, code quality and integrity are

production to help ensure reliability at scale.

Transparent and Since developers must understand, maintain, and extend

or fixes were chosen.

paramount. Al agents should be validated against test suites, subjected to rigorous regression testing, and monitored in

Al-written code, agents should provide clear documentation

includes explaining why specific code structures, libraries,

and traceability of their logic and coding decisions. This





Although Al agents can perform coding tasks well, ultimate responsibility for design decisions, compliance, and final code deployment should rest with human engineers and engineering managers. Clear governance frameworks should ensure that accountability is not abdicated to autonomous systems.

B











Al agents for service lifecycle management

Automating technology and telecom service activities

Agentic AI can manage key service lifecycle functions – including customer support, network operations, and billing – operating with minimal human intervention to improve efficiency and scalability.

ISSUE/OPPORTUNITY

Telecom and technology providers face rising demand for responsive, always-on service delivery. Traditional service lifecycle processes – such as handling customer inquiries, managing network performance, and processing billing – often

rely on large teams performing repetitive tasks across multiple systems. These manual approaches can create delays, inconsistencies, and higher operating costs, while limiting the ability to scale operations effectively.

HOW AI CAN HELP

Specialised functional agents

Al agents can be tailored for specific functions such as customer service, network monitoring, billing management, and field support, each handling tasks autonomously with high speed and consistency.

Customer support automation

In support environments, Al agents can triage tickets, resolve common issues, and escalate complex cases as needed, whether through voice or text-based channels.

Operational integration

Integrated into existing enterprise platforms, agents can update records, monitor systems, and provide real-time recommendations, seamlessly aligning with current workflows.

Scalable orchestration

Multi-agent systems can coordinate activities across service lifecycle stages, enabling efficient end-to-end management with minimal human oversight.

On-site troubleshooting support

Al-powered tools equip field technicians with access to troubleshooting guides and generate clear, step-by-step resolution plans, improving on-site problem solving and network performance.





POTENTIAL BENEFITS

Greater efficiency and lower costs

Automating routine service tasks can improve the productivity of support and operations teams, reducing both the time and cost required to address customer inquiries and technical issues.

Improved customer satisfaction

Al-driven service that is faster, more consistent, and more reliable enhances customer satisfaction and loyalty.

Scalable operations

Organisations can expand support capacity and manage higher service volumes without proportional increases in headcount. Also, existing service staff can shift their focus to activities that are more complex and strategic.













Al agents for service lifecycle management

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Maintaining data integrity and ensuring accurate ticket resolution are critical, especially when agents operate autonomously. Al agents should be continuously validated for accuracy and stability.



Safe and secure

Given the risks of cyberattacks targeting critical telecom and technology infrastructure, agent systems should be hardened against intrusion and designed with robust cyber security measures to protect networks from sophisticated attacks.



Private

Since agents often handle sensitive customer and operational data, strong security protocols and strict compliance with privacy regulations are essential.



Responsible and accountable

Agentic AI can make mistakes or overlook nuances in complex service cases. Human supervisors need to closely monitor Al agents and retain final responsibility for actions and decisions, intervening whenever necessary – especially in situations that involve sensitive customer information or regulatory compliance.

HOW AI CAN HELP



Financial forecasting

and profitability.

analysis

Marketing effectiveness

A marketing effectiveness agent

recurring insights on ROI, ROAS, and customer lifetime value.

performance data to provide

can analyse historical and current

A scenario planning agent can run simulations of different budget strategies to model impacts on subscription starts, retention, and overall profitability.

Scenario planning

Marketers need tools that can continuously analyse can rule performance, forecast outcomes, and support decision-making on where best to invest their resources. A financial forecasting agent can analyse marketing budget inputs and business assumptions to generate full-year projections for revenue, spend,

Coordinated decision support By sharing insights across forecasting, effectiveness,

forecasting, effectiveness, and scenario planning, Al agents can collaborate to provide a comprehensive view of marketing's financial contribution.











Efficient marketing spend

Using AI agents to improve the efficiency of marketing spend

Agentic AI systems can help media and entertainment companies get more value from their marketing spend by forecasting financial outcomes, evaluating ROI, and running budget scenarios to optimise growth, retention, and profitability.

ISSUE/OPPORTUNITY

Media and entertainment companies invest heavily in marketing to drive subscriptions, viewership, and engagement; however, measuring and optimising financial impact remains a complex challenge. Marketers must balance budgets across multiple channels, predict outcomes such as customer acquisition and retention, and spend their money as efficiently and effectively as possible – particularly during highprofile events, which have outsized financial implications.

Traditional forecasting and ROI analysis often rely on manual processes and disconnected tools, making it difficult to produce timely, accurate insights. Without more adaptive, data-driven approaches, marketers risk overspending, missing revenue opportunities, or failing to adjust strategies quickly enough to reflect changing consumer behaviour and market conditions.

Marketing Agentic AI

164













Efficient marketing spend

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Inaccurate forecasts and simulations can lead to poor budget decisions. To improve reliability, AI agents should be continuously tested and refined to reflect real-world market conditions.



Private

Because customer-level data is used to measure marketing effectiveness, strong data governance and compliance with privacy regulations are essential.



Responsible and accountable

Al agents are fallible and cannot perfectly anticipate all external factors or sudden market shifts. Also, Al hallucinations remain a concern. Ultimately, humans need to be responsible and accountable for all decisions and actions.

POTENTIAL BENEFITS

Higher revenue

By automating forecasting and scenario planning, agentic AI helps drive higher revenue through more accurate bundle sign-up predictions and targeted marketing efforts.

Improved efficiency and conversion rates

Greater efficiency in owned and paid channels – including both traditional linear TV and connected TV (CTV) – can lead to better conversion rates for subscriptions.

Faster, more agile decisions

Agentic AI enables a business to respond more quickly to market changes and optimise marketing spend during high-impact events.

Al-supported budget allocation

Using AI agents to optimise how marketing budgets are allocated

Agentic AI systems can inform and optimise global, regional, and national marketing budget allocations across cost centres, brands, content types, and performance channels.

ISSUE/OPPORTUNITY

Global media and entertainment companies operate across multiple regions, products, and channels, with significant marketing investments needed to drive subscriptions and engagement. Traditional budgeting processes often rely on static planning cycles and siloed data, making it difficult to dynamically reallocate spend across cost centres, brands, content types, and performance channels.

Without better budgeting tools, organisations risk under-investing in high-performing areas, overspending in less effective channels, and creating global and regional strategies that are out of line with local market realities.

HOW AI CAN HELP

Cost centre analysis

Al agents analyse marketing spend across cost centres and recommend reallocations that optimise key performance indicators (KPIs) such as sign-ups, churn, and viewership.

Product, channel, and vendor optimisation

Other agents assess product, channel, and vendor mix within each market, using predictive models to maximise subscription starts and profitability.

Regional and country-level scenarios

Scenario planning agents run simulations of budget allocations at regional and country levels, evaluating the impact of incremental or cost-saving changes on subscriptions, retention, and profitability.

Coordinated recommendations

Agents collaborate to provide a unified view of marketing performance and deliver actionable recommendations for smarter, more flexible budget allocation.



























Al-supported budget allocation

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Bad marketing budget decisions can have serious consequences for the broader business. Agents should be validated regularly against historical budgets and real-world outcomes, using continuous performance tracking to help ensure accuracy and minimise the risk of misallocation.



Transparent and explainable

Agents should be able to provide clear explanations of their reasoning – e.g., how reallocating \$X from channel A to channel B affects ROI – so marketing leaders can understand and trust their recommendations.



Responsible and accountable

Al agents should be used as decision-support tools, with final responsibility reserved for marketing and finance leaders who understand strategic business contexts.



Safe and secure

Given the sensitivity of financial plans, systems must be built with enterprise-grade security to help ensure budget forecasts and proprietary data are protected from breaches or misuse.

POTENTIAL BENEFITS

Improved financial outcomes

Optimising the product, channel, and vendor mix within each country can maximise subscription starts, customer retention, and overall profitability.

Greater spend efficiency

Efficiency improvements in both CTV and owned marketing channels can drive higher conversion rates and reduce wasted marketing spend.

More flexible decision-making

Real-time scenario planning enables marketing teams to adjust budgets dynamically, aligning strategies with rapidly changing market conditions.

Al for gamers

Game content development

Developers can leverage AI – including generative models – to maintain and update their game with new assets and content in line with user community requests and interests.

ISSUE/OPPORTUNITY

Game development requires a massive up-front investment in time, resources, and capital. AAA games can cost tens of millions of dollars to develop and take years to complete. These costs will only rise as players increasingly demand more complex games, more post-release

support, and more frequent content updates. Al provides the gaming industry with an opportunity to bend the cost curve through enhanced development efficiency, while also simultaneously meeting player demands.

HOW AI CAN HELP

Ongoing content development

Post release, developers can rapidly generate and deploy new gaming assets as expansions or microtransactions, such as seasonal or downloadable content (e.g., new characters and skins). Developers can use text prompts to generate new content in line with the existing game and requests from the community, then upload those assets to the game.

























Al for gamers

MANAGING RISK AND PROMOTING TRUST



Responsible and accountable

Generated content resulting from a model trained with proprietary third-party data may lead to copyright claims if it is deemed to be too similar (without substantial variation).



Safe and secure

Players' personally identifiable information (PII) could be fed into the models as they interact within the game, which raises risks around cyber security and regulatory compliance. The collection of PII, even inadvertently, places an obligation on the organisation to secure the data as it is accessed, transferred, and stored.



Fair and impartial

Generated assets may over-index on player segments providing feedback or residing in specific regions. This uneven sampling of the input data could lead to bias in what assets are generated, and it may lead to missed opportunity and revenue as some customers are ignored.

POTENTIAL BENEFITS

Greater efficiency for greater creativity

By automating the process of creating game content, developers have more capacity to work on creative game designs and explore new, innovative ideas.

Cater to gamers

More immersive, controllable, responsive, engaging, and unique experiences for gamers (based on community requests and existing popular assets) has a direct impact on player lifetime value.

Drive new revenue

When add-on content can be generated with minimal human involvement, it creates new revenue streams with minimal investment.

Content creation with AI

Al-enabled creative tools

Content creation can be facilitated and enhanced with AI tools that minimise the need for manual editing and time-consuming content management.

ISSUE/OPPORTUNITY

Content creators and managers are faced with large volumes of data that require considerable time to generate, edit, and oversee. Creators also face tight deadlines that require high levels of efficiency for content management and editing.

Significant time and resource investments are needed for video and image editing, and the volume of content creates challenges around data management and finding the right content at the right time.

HOW AI CAN HELP

Creative assistant tool

Al can be used to create imagery and apply edits using descriptive commands. Features like conversational editing, text-to-template, and text-to-image allow users to expedite the editing phase of the content creation process.

Picture editorial

Producers can automate footage management with video-to-text Al to evaluate and create tags for scenes and content. Text-to-video commands (e.g., "add more rain to this scene") can be used to enhance and accelerate the editing process.

Al "reshoots"

Content creators can use scripts and 3D scans of actors to generate new content, alter footage to create more realistic special effects, and allow studios to make edits without the need for reshoots.

























Content creation with AI

MANAGING RISK AND PROMOTING TRUST



accountable

Responsible and Al tools may be trained with large databases of media and content, some of which may be copyright protected. As a result, the model outputs may include aspects of a creator's or studio's work or style that are not attributed to them, which raises legal risks.



Robust and reliable

Noticeable changes in style and brand quality due to Al content creation and editing may erode consumer trust in the brand and content.



Private

If bad actors access the underlying models or applications, it could contribute to the spread of fake content on behalf of the organisation, leading to misinformation. Model owners should ensure strong privacy and access controls to mitigate this risk.

POTENTIAL BENEFITS

Greater efficiency

Content management stakeholders can gain efficiencies by leveraging creative tools to facilitate work and even create net-new content across the production lifecycle.

Improved content quality

Generating novel content can supplement the human creative process and potentially lead to a higher-quality product.

Content tailored to the audience

With AI, creators can hyper-personalise content with prompts driven by consumer trends and interests.



Marketing content multiplier

On-brand publishing

Using AI, marketing content generation can be quicker and more cost effective, while still preserving the company's brand identity.

ISSUE/OPPORTUNITY

When multiple authors are contributing to a piece of marketing or business content, there can be quality and consistency issues with tone and brand values. Authors are challenged to consistently balance product promotion with thought leadership and insight. As such, on-brand publishing is a significant time and cost investment

that requires a long-term commitment to generating content that establishes the organisation's (or its leaders') subject matter authority. However, the return on investment for on-brand publishing can be difficult to measure because the impact itself is complex and challenging to quantify.

HOW AI CAN HELP

Cohesive content generation

Al systems can be trained with on-brand content to mimic the style of company marketing materials and generate new, high-quality content rapidly and on demand.

Ideation with generation

Marketing departments can leverage AI to quickly create multiple versions of content in various styles to identify the most compelling and persuasive option.

Tailored, personalised messaging

With AI, organisations can easily create multiple versions of the same on-brand marketing tailored to different customers and audiences.

















POTENTIAL BENEFITS

Instant marketing

Companies can create content better tailored to their brand and customers, iterating through multiple drafts as needed.

Time savings

As AI systems instantly generate content, human staff can shift to an editorial role.

Broader marketing range

With the ability to easily create content across various formats, styles, and topics, companies enjoy greater flexibility in how they reach their customers. This also allows companies to more rapidly adapt to marketing trends.













Marketing content multiplier

MANAGING RISK AND PROMOTING TRUST



Transparent and explainable

Personalised advertisements may be customised based on data collected or purchased from individuals. This may be off-putting to consumers who realise the organisation has such broad access to their data, potentially harming the enterprise's brand reputation and undermining consumer trust. One way to mitigate this outcome is to ensure data collection and usage policies are transparent and communicated meaningfully to the consumer.



accountable

Responsible and Content produced by Al systems may not be subject to the same protections as human-generated content. Companies need to be wary of infringing on copyrighted material used to train Al systems.



Safe and secure

When brand data is used to train Al, there is a risk of data leaks that could result in sensitive information or IP being divulged to competitors. Companies need to ensure that their proprietary information is safely stored, transferred, and used, as well as monitor model outputs to validate that protected information is not being revealed.



Language translation at scale

Content localisation

Al can be used to quickly and easily scale content across regions by translating and converting text and audio into regional languages.

ISSUE/OPPORTUNITY

The ability to create and translate content at scale can be a competitive differentiator for multinational enterprises, but it can also command significant time and resources, and rapid, on-demand translation may be difficult to achieve.

HOW AI CAN HELP

Tools for custom localisation and quality assurance

Al can be used to help organise and manage complex file type, analyse content before translation to optimise localisation, and integrate glossaries, term bases, and language tools into workflows.

Content personalisation across industries

Al-powered content personalisation can supercharge localisation efforts by helping to improve engagement, build brand loyalty, and increase conversions.

Speech recognition during translation

Al can be leveraged to enable voice user interfaces (VUI), transcribe video and audio content into text, and simultaneously translate spoken content into the target language.

























Language translation at scale

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

Bias in the data used for content personalisation could lead to unequal and unfair recommendations for certain groups of customers. In addition, Al applications are often trained on datasets from major languages, which means LLMs may have lower accuracy rates for less common languages and alternative dialects.



Transparent and explainable

Messaging and tone may change with language translation, which may negatively impact the text or audio being generated and the overall quality of the content. Localisation should be audited to make sure that the messaging remains consistent with the original intent.

POTENTIAL BENEFITS

Enhancing translation

Translation processes using AI can lead to improved speed, accuracy, and scalability.

Improving the customer experience

A wider availability of language resources with the quality and speed enabled by Al promotes a higher-quality user experience.

Ensuring quality

Organisations can leverage Al to automate quality assurance for the localisation of digital assets by providing more accurate natural language processing.

Enhancing chip innovation

Semiconductor chip design & manufacturing

Al can be used to iterate chip designs by having designs "compete" across a set of performance dimensions.

ISSUE/OPPORTUNITY

With demand for ever more powerful semiconductor chips, design complexity is rising. As semiconductor sizes continue to shrink, density scaling becomes a challenge, since upgraded features are required to fit on perpetually smaller chips.

HOW AI CAN HELP

Iterative chip design

Al can generate and iterate chip designs and improve the outputs by having chip designs "compete" across a set of performance dimensions. At each new iteration, chip parameters are tweaked based on learnings from the best-performing designs in past iterations. These models are trained on existing layouts to learn patterns and constraints and generate new layouts that meet specific design requirements.





















Cost and time

By shortening the development lifecycle, the enterprise can reduce total development costs.

Create new ideas

Al can help improve designs or discover entirely novel designs that optimise performance based on specific criteria, such as power consumption, performance, location, and manufacturability.













Enhancing chip innovation

MANAGING RISK AND PROMOTING TRUST



Safe and secure

With each new generation of novel designs, there is a risk of IP leakage and data breaches for proprietary chip designs and technical specifications generated by the LLM that could severely damage the organisation's competitive advantage. There should be rigorous security protocols in place to protect against this.



Responsible and accountable

When using AI for design, the organisation needs to consider how to secure copyrights or patents and protect the IP of chip designs that are moved into production.



Transparent and explainable

For complex simulation processes, the organisation needs the capacity to understand how and why the model determined a scenario or design to be optimal. Design validation requires users and stakeholders to be able to understand the reason for the outputs.

Tech specs on demand

Technical sales, operations, and field staff knowledge management

Al can help sales, operations, and field staff quickly find and translate technical specifications to enable faster knowledge retrieval.

ISSUE/OPPORTUNITY

Technology offerings require deep technical understanding and the ability to find the right technical specifications in a timely manner. When it comes to translating technical specs and responding to customer technical questions, sales, operations, and field staff can be challenged to translate the information and effectively communicate it to the customer.

A big part of the problem is the timeconsuming process of scouring through vast amounts of unstructured information and knowledge documents that contain the specifications and answers customers are seeking.

HOW AI CAN HELP

Spec summarisation and search

Al can be used to create summaries of technical specifications based on targeted text-based queries to help understand which products meet customer requirements. It can suggest features and integrations that align with the customer's existing technology stack and vendors, as well as provide links to articles or an internal knowledge base for future reference.

Knowledge management update

Sales case histories and other current documents can be used to update knowledge management so similar technical inquiries in the future can be rapidly addressed with previous resolution steps and summarisations.

Automated technical demos

Al can be used to automate the creation of software demonstrations tailored to specific clients and use cases. This is achieved by training on demo scripts and sample interactions to generate demonstrations showcasing a solution's key features and benefits.

























Tech specs on demand

MANAGING RISK AND PROMOTING TRUST



Private

Customer data (e.g., sales case history, customer tech stack/vendors) must be processed by the model, making it necessary to continuously monitor model outputs and safeguard customer data to mitigate privacy risks.



Robust and reliable

Al models are susceptible to hallucinations, or factual inaccuracies, making human validation essential to establishing trust in the outputs and the decisions they inform. What is needed is a verification process to ensure the accuracy and reliability of information derived from the model (e.g., spec summarisation, demos). This has a direct impact on answering customer questions, and by extension, customer satisfaction and sales.

POTENTIAL BENEFITS

Efficiency with automation

Al's ability to quickly consult and summarise technical specifications greatly reduces manual effort for sales, operations, and frontline staff when responding to technical sales inquiries.

Tailored to the customer

Greater personalisation in responses and demonstrations improves the customer sales experience and increases the chances for conversion.

Supporting sales staff and other stakeholders

With AI, staff can rapidly create content to support the sales and marketing process, and to address specific questions from customers and partners.



B











Al-powered RFP and knowledge assistant

Automated proposal generation and sales knowledge management

Al can produce RFP responses automatically and help sales teams prepare for pitches by providing easy access to internal knowledge resources through smart chatbots.

ISSUE/OPPORTUNITY

Sales processes are often constrained by how quickly teams can access institutional knowledge and respond to Requests for Proposals (RFPs). Many sales teams have only days to coordinate across multiple departments and deliver detailed technical and commercial responses. Their ability to respond can be slowed by manual processes, fragmented internal documentation (e.g., playbooks and product briefs), inconsistent proposal quality and knowledge reuse across teams, and limited tools to extract and synthesise key information.

Al-powered tools can accelerate sales professionals' ability to retrieve, understand, and reframe information for client needs – without requiring technical expertise or deep coordination across departments.

HOW AI CAN HELP

Providing easy access to internal knowledge through chatbots

Salespeople can converse with Alpowered chatbots to quickly and easily retrieve sales playbooks, technical specs, competitive positioning, and customer references directly from internal documentation repositories.

Automatically drafting RFPs

Al models can produce highquality, tailored RFP responses by finding and summarising relevant content from existing sales documents, aligning answers with internal knowledge bases, and incorporating reusable proposal components.

Providing individualised sales support with little or no coding

Non-technical users, including sales reps and subject matter experts, can generate summaries, extract insights, and draft proposals through a simple user interface – no prompts or coding required.

Enabling customised sales processes and tools

Technical users can integrate Al tools directly into other internal systems, workflows, or dashboards to build more personalised applications.

Al-powered RFP and knowledge assistant

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Al-assisted workflows can deliver higher-quality outputs – in less time – than traditional approaches. Users can flag incorrect responses or incomplete information; these are logged and reviewed in recurring QA cycles. Also, fallback mechanisms should exist to ensure consistent availability if problems arise with the AI models.



Transparent and explainable

Documentation should be provided for both business users and developers to explain how the system processes inputs and generates outputs. The chatbot interface includes citation tracing, where users can see which source documents were utilised to generate responses. Proposal-generation tools can allow users to edit and review outputs before submission, promoting human-in-the-loop oversight and transparency.



Safe and secure

All data and model interactions should occur within a secured internal environment, with no calls to third-party APIs unless vetted and approved. Systems should support audit logging for all user interactions to help ensure traceability and compliance. Role-based access controls can ensure only authorised personnel are able to view or generate sensitive proposal content.



Private

The system should not log personally identifiable information (PII) unless required by specific business rules and protected under internal data governance protocols. Feedback mechanisms should be anonymised where appropriate, helping to ensure user privacy while supporting continuous improvement. RFPs and customer documents processed in the system should be stored temporarily and purged according to data retention policies.

POTENTIAL BENEFITS

Faster deal cycles

Sales teams can respond to RFPs and prepare sales pitches/collateral much more quickly than before, accelerating the sales cycle.



With centralised, Al-assisted knowledge access, sales teams can produce responses that are more consistent and comprehensive - reducing errors and potentially improving win rates (especially for opportunities with time-sensitive budget windows).

Increased sales rep productivity

Salespeople can search for materials or draft proposals more quickly, freeing them to focus on sales strategy, client relationships, and personal follow-ups.

Path to commercialisation

Once validated internally, Al-powered sales tools have the potential to be offered to external customers, turning ar internal efficiency driver into a revenuegenerating product.















HOW AI CAN HELP

code generation.

Automating test creation

Al systems can also help identify gaps in testing coverage and can prioritise high-risk areas based on historical failure data, although integration with structured data and internal governance systems remains an ongoing challenge.

Al tools can be used to create new test cases from product requirement documents, bug histories, and structured datasets. These tools can assist engineers by proposing a wider set of test scenarios including ones not previously considered and by automating portions of test implementation through













Automated test case generation

Al-powered test case generation and automation in chip development

As chip designs become more complex and product cycles accelerate, engineering teams are leveraging AI to automate test case generation and validation.

ISSUE/OPPORTUNITY

Chip development demands exhaustive testing and validation due to increasing functional complexity and the high cost of post-release defects. Human testers may struggle to keep pace with the volume and sophistication of required test cases, leading to potential quality

issues, slower development cycles, and growing verification costs. Yet, security vulnerabilities or missed bugs can result in major product delays, public backlash, and brand damage, prompting chip manufacturers to add even more layers of testing.

Automated test case generation

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Generated test cases can be validated against known test results and manually vetted to help ensure they hold up under real-world complexity. Also, systems can be stress-tested with increasingly complex product requirement documents to assess scalability and robustness across chip generations.



Transparent and explainable

Al-generated test cases can be accompanied by natural language summaries or rationales explaining why certain logic or edge conditions were selected. Engineers can trace outputs back to source inputs (e.g., PRD sections, bug databases), enabling better understanding and debugging of the Al system itself.



Safe and secure

The development and inference processes can occur in sandboxed environments with strict access controls to prevent accidental leakage of proprietary information. Integration with external AI services should be carefully managed to ensure no sensitive IP or design data is exposed to third-party systems.

POTENTIAL BENEFITS



Increased test coverage and enhanced product quality



Al can enable the generation of more comprehensive test cases than previously possible with human effort alone, allowing for earlier defect detection. Also, by identifying edge cases and potential failure modes, Al can reduce the risk of catastrophic bugs slipping into production.



Faster time-to-market



Automation accelerates the validation process, allowing development teams to keep up with faster chip release timelines and feature rollouts.



Operational efficiency and cost control



Al helps teams do more with less, reducing reliance on manual testers and mitigating the need to grow headcount to handle increasing workloads.

Improved development process

As the test tools mature, there is potential for deeper integration with the design and verification phases, improving end-to-end development flow across decentralised teams.



B











Al-powered source separation for music remastering

Separating mixed audio tracks into their component parts using AI

Al can separate vocals or instruments from mixed audio tracks even when the original files are not available, opening up new opportunities for licensing, remixing, archival restoration, and monetisation.

ISSUE/OPPORTUNITY

Many recordings in music labels' back catalogues were produced at a time when multitrack preservation practices were inconsistent, and, in many cases, the original recordings have been lost, damaged, or never existed in isolated formats. This limits the ability to fulfil requests for custom edits – such as instrumentals, a cappella songs, or remixes – thereby stalling lucrative licensing deals,

particularly for synchronisation (music in film, television, and advertising) and derivative content creation. Manual audio reconstruction is costly, time-consuming, and often technically infeasible at scale. Yet demand for high-quality, tailored audio continues to grow, especially with the global expansion of streaming and sync opportunities.

HOW AI CAN HELP

Separating music into its component parts

AI, particularly deep learningbased source separation models, can analyse a fully mixed audio file and isolate its constituent elements - vocals, guitar, bass, drums, ambient noise, etc. into discrete audio tracks with high fidelity. These models have matured significantly in recent years and can now perform at a level sufficient for commercial use in many scenarios. Rather than depending on traditional DSP (digital signal processing) or manual studio methods, the Al learns from large datasets of music to "de-mix" the sound using learned patterns of frequency and structure...

Leveraging Software-as-a-Service

Most deployments today use Al-powered SaaS platforms that allow internal teams to process catalogue tracks quickly and securely. Internal quality control – along with artist or management approval – is then layered on to ensure that the extracted stems meet the creative and technical expectations of the project.

AI-powered source separation for music remastering.

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

All outputs from Al models are subject to expert human review. Because source separation can introduce artefacts, tracks should be assessed case-by-case to determine if the fidelity is suitable for commercial or creative use. Teams should be trained to identify when alternative methods or manual interventions may be more appropriate.



Transparent and explainable

Processes for using Al in audio separation should be clearly defined internally and communicated externally as needed. Stakeholders – including sync partners, artists, and producers – should be informed when Al-generated stems are used, and how those stems were derived from the source material.



Responsible and accountable

All source separation use should be logged, and responsibility for approving commercial use should rest with both label- and artist-facing teams. If stems are to be reused, remixed, or publicly released, the appropriate clearance workflows - including licensing and revenue-sharing – must be followed.

POTENTIAL BENEFITS



Commercial monetisation of back catalogues



Al-powered source separation can make more recordings available for synchronisation deals, remixing projects, or global reissues in alternate formats.



Accelerated time-to-licence



The speed and efficiency of AI can minimise delays associated with locating or recreating stems, enabling a faster turnaround for timesensitive media productions...



Cost-efficient alternative to studio sessions



Al offers a high-quality yet faster and less expensive alternative to manual isolation or re-recording, which are both time-consuming and expensive.

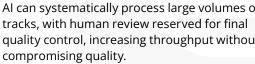
Artist-led remix and reimagination projects



Using AI to extract source elements, artists can revisit and reinterpret their own work or collaborate across genres. Even in less creative scenarios, artists and labels can maintain full control over what gets extracted and used, ensuring all usage aligns with legal, creative, and ethical standards.

Operational scalability

Al can systematically process large volumes of tracks, with human review reserved for final quality control, increasing throughput without















AI-powered archive access and extraction

Transforming historical news content into a valuable asset

All enables news organisations to recover legacy content lost to system or format issues – turning dormant information into a usable, searchable, and monetisable asset.

ISSUE/OPPORTUNITY

News archives hold cultural, journalistic, and commercial potential. But over time, many of the most significant stories – especially interactive long-form journalism, investigative pieces, and special coverage – have become inaccessible due to technological evolution, changes in content management systems (CMS), format obsolescence, and a lack of centralised archives.

Reporters and editors often cannot locate stories they know exist, especially from the early digital era (late 1990s to early 2010s). Multimedia components such as photos, graphics, and maps have not always been retained or migrated, rendering even recovered content incomplete.

HOW AI CAN HELP

Document extraction and digitisation

Al models can process and extract structured information from legacy formats such as PDFs, microfilm scans, and outdated HTML, even when metadata is missing or inconsistent.

Content reconstruction

Al tools can intelligently identify article structure (headlines, subheads, body text, captions, bylines), reconstruct layout context, and reassemble fragmented articles into coherent, readable documents.

Semantic indexing and search

Large Language Models (LLMs) enable content to be semantically tagged and categorised, improving discoverability across themes, time periods, people, and places – even when specific keywords are not used.

Metadata enrichment and linking of multimodal assets

Al can supplement missing or corrupted metadata (e.g., publication date, author, topic) by analysing linguistic and contextual clues. Also, the technology can cross-reference and re-link associated images, graphics, or videos from various archives where files may have been separated during prior migrations.

Improved access

Al can provide improved interfaces – such as chat-style queries or timeline exploration – to help users engage intuitively with the archive.

AI-powered archive access and extraction

MANAGING RISK AND PROMOTING TRUST



Fair and impartial Systems are designed to ensure access to historical content across different eras and communities. Bias mitigation strategies are incorporated into model training and metadata tagging to avoid skewed representation of topics, regions, or individuals.



Robust and reliable

Extraction and structuring workflows are tested across various content types and legacy formats to help ensure consistent quality. Human oversight is embedded throughout the process to validate the accuracy and fidelity of reconstructed articles.



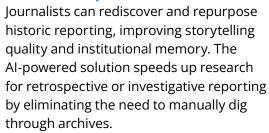
Transparent and explainable

A clear audit trail should be maintained for all Al-processed content, including logs of when and how specific items were extracted, tagged, and categorised. Explanatory overlays and metadata annotations help end-users understand the origin and limitations of Al-reconstructed documents.

POTENTIAL BENEFITS



Editorial improvements





Monetisation

Al can enable news organisations to expand their relationships with libraries, educational institutions, and content platforms while providing the foundation for new archive-based products, such as nostalgia-based newsletters and historical collections. What's more, it positions news organisations to negotiate more effectively with AI companies looking for premium training data by presenting them with a



curated, high-quality proprietary dataset.



Improved operational efficiency The solution can reduce ad-hoc archive retrieval and reduces the need for internal technical support to help recover content. Also, it strengthens the organisation's institutional capabilities for structured knowledge management.



Contact us

Deloitte AI contacts



Mariette Van Niekerk Trustworthy Al Lead mvanniekerk@deloitte.co.nz



Aravind Subramanian
Partner, AI & Data
aravsubramanian@deloitte.co.nz



Silvio Bruinsma
Partner, Assurance & Advisory
sibruinsma@deloitte.co.nz



Daniel Lawrence Director, Transactionsdanlawrence@deloitte.co.nz



Phillip Claridge
Director, Tax
pclaridge@deloitte.co.nz



Dan Hellyer Partner, Business Advisorydhellyer@deloitte.co.nz



Dr. Amanda Williamson
Director, Deloitte New Zealand Al Institute
amawilliamson@deloitte.co.nz

www.deloitte.com/nz/AlInstitute

Deloitte.

This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited, its member firms, or their related entities (collectively, the "Deloitte Network") is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. No entity in the Deloitte Network shall be responsible for any loss whatsoever sustained by any person who relies on this communication.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.

Copyright © 2025 Deloitte Development LLC. All rights reserved.