

ORACLE



AI in Action:

10 Cutting-Edge Innovations to Explore Now

—

These real-world—often multimodal—use cases
plus your data can deliver competitive advantages






Table of contents

AI in action 3

Image analysis extends vision 4

Send in the advanced agents 5

Automation, accelerated 7

AI sentiment analysis is here to help 8

An AI assist to developers 10

Upgrade your upsell 11

HR’s little helpers 12

Cybercops on the beat 14

Fix before break 15

Ambient AI keeps watch 16

How Oracle can help 17



AI in action

By Joseph Tsidulko
Senior Writer

Rapid advances in generative AI and complementary machine-learning-powered technologies, such as object detection and trend forecasting, have moved even the most aspirational use case wish list items from “wouldn’t it be nice” to “let’s do it.” After all, the point of investing time and money in training intelligent algorithms is so they can help people work better, faster, and safer.

We’re talking about emerging applications that incorporate AI-powered tools like computer vision, speech recognition and synthesis, natural language processing, and forecasting engines. Many cutting-edge products now employ multimodal AI, which strings together AI models, both traditional and generative, and their associated data to perform complex tasks. Consider a self-driving car: The passenger speaks an address to a chatbot, the navigation system plots an optimal route, real-time image analysis interprets video coming from cameras and sensors, and another AI hits the gas and steers.

The same concept works to automate complex business processes, often using higher functioning AI tools called agents. You’ll hear a lot about how they’re enabling creative approaches to automating digital and physical processes.

Let’s look at 10 use cases that may spark your imagination.

Organizations train intelligent algorithms so they can help people work better, faster, and safer.



Image analysis extends vision

It wasn't so long ago that researchers struggled to build AI that could identify simple objects that humans recognize instantly. In 2012, when a neural network powered by 16,000 processors proved adept at finding cats in YouTube videos, the achievement was hailed as a milestone, and for good reason—different breeds, colors, and sizes of cats appeared in the videos at different angles, under different lighting. It was a breakthrough. It took only a few more years, however, before image analysis models trained by a subset of AI called computer vision had advanced to being able to quickly classify thousands of distinct breeds of cats and animals of other species.

By making it possible for AI developers to harness large arrays of processors, especially GPUs, for training sophisticated models on bigger, better data sets, the public cloud has been a major driver of rapid advances in object detection. Innovative companies are now developing tools that can spot subtle and meaningful details faster and more accurately than humans ever could. Some companies have trained deep neural networks to scan X-rays and images of biopsies to help radiologists better detect cancers¹. Other systems pore through thousands of drone images covering hundreds of acres of fields to pinpoint pests, diseases, and crop defects before they ruin a harvest.

Almost every industry is exploring potential applications of this technology—object detection is finding use cases as diverse as personalized marketing, construction safety, anomaly detection, self-driving vehicles, and digital media organization. As one example, the US Transportation Security Administration [plans to use AI](#) to help inspect X-rays screening carry-on luggage, easing an important yet fatiguing task for its human agents.²

This technology is also being coupled with LLMs to deliver AI that can not only detect an anomaly, be it a manufacturing defect or a cancer, but also suggest a remedial action.

¹ [PMC PubMed Central](#)

² [Nextgov/ECW](#)

Send in the advanced agents

GenAI-powered chatbots, a familiar type of agent that connects the human and machine worlds, are a qualitative leap from the previous generation of rules-based conversational assistants, which operated primarily based on predefined rules or logic. But there are many more [AI agents](#) discreetly plying their trades, querying relevant data sources and applying generative models to autonomously execute complex actions—think anonymous workers in a tucked away cubicle adeptly shuffling along reports, processing payments, responding to complaints, and more.

When these agents come together, working in unison to fulfill multiple discrete functions in a comprehensive workflow, you get agentic AI. And that technology is on the verge of dramatically changing how companies go about engaging customers, managing employees, conducting financial operations, and administering transactions.

Consider onboarding new employees, a routine yet cumbersome business process. An AI agent can contact new hires with relevant notifications and materials, help them understand their benefits, assist them in completing tax documents and other paperwork, schedule trainings, and grant access to internal resources, such as IT systems and facilities required by the new hire's role.

Multimodal agents get to work

Multimodal AI refers to models that can take in a diverse range of data beyond just text. Think images, audio, video, sensor data, location, and more. By combining and analyzing these different forms of input, a multimodal agent can operate with a more comprehensive understanding of the world.

Examples include a chatbot that uses understanding of both text and images, such as in a schematic diagram, to answer questions about a product, or a recommendation engine that analyzes a customer's browsing and buying history then taps product images and reviews to suggest what to purchase next. Applications abound in healthcare, warehouses, and more.

By combining and analyzing these different forms of input a multimodal agent can operate with a more comprehensive understanding of the world.

And agents are poised for far wider adoption now that a major stumbling block for enterprises has largely been overcome. Agents employ models trained on universal sets of data; they don't know your business specifics, and you wouldn't want to give their developers the confidential data that would let a model incorporate your private data. Retrieval-augmented generation, or RAG, solves that problem by directing agents to proprietary information that tailors outputs, be they text or operations. Essentially, RAG offers a context for the foundation model by incorporating what's in your company's restricted databases and applications without exposing any of that knowledge to the model itself.

RAG empowers agents to take actions useful to a business as authorized—search inventory, email prospects, enter patient symptoms into medical records, and so forth. By specifying context, RAG helps keep the output of the foundation model up to date while reducing instances of the hallucinations that erode trust in GenAI.



RAG empowers agents to take actions useful to a business as authorized—search inventory, email prospects, enter patient symptoms into medical records.

Automation, accelerated

Business process automation (BPA) may not be as exciting as humanoid robots assembling machines or the self-driving cars increasingly cruising our neighborhoods. But like those flashier AI use cases, BPA has seen significant advances in recent years and is poised to deliver far greater efficiency and reliability to enterprises in almost every industry. In fact, today's systems can run entirely autonomously while keeping human approvers in the loop. This helps build trust in the AI.

Many firms have already put [robotic process automation](#) (RPA) to work automating repetitive tasks. RPA deploys software bots to interact with user interfaces of enterprise applications and websites—mimicking the way humans do those same rote tasks. That approach, dictated by straightforward rules, is effective at entering structured data into standard forms or extracting it from them and shuffling along files in predefined workflows. But these bots need help when the next course of action isn't obvious.

That assistance need no longer come solely from humans. RPA bots trained using machine learning can handle more sophisticated workflows than those obeying predefined rules, and they can improve their performance over time. AI-powered services, such as computer vision and natural language processing, further empower RPA to operate autonomously by imparting deeper understanding of varied images and documents used within business processes.

Intelligent agents can add another layer of sophistication to RPA's recognition and decision-making capabilities. By employing LLMs to understand user prompts and textual context and adeptly bring humans into workflows when they're needed for guidance or approvals, GenAI empowers RPA to direct workflows based on evolving factors and conditions. That lets even more business processes be automated.

Shored up by multimodal AI, BPA systems will increasingly incorporate unstructured data, optimize workflows without human input, adeptly interact with humans when needed, compare results to forecasts, and seamlessly operate across diverse IT systems.



AI sentiment analysis is here to help

Every business wants to know how its customers feel about its products and services. And increasingly, companies are relying on machines to tell them.

AI is becoming more adept at interpreting human sentiment. This type of analysis, sometimes called “emotion AI,” has applications beneficial to all kinds of businesses, especially in the areas of customer service, sales, marketing, and product development.

Already, AI routinely parses emails, social media posts, chatbot dialogues, surveys, and other textual interactions to determine customer sentiment. Based on its analysis, the system can make a good guess as to whether a customer is angry or satisfied, if a service ticket needs to be elevated, or if a troubling trend on social media needs to be addressed.

Speech-to-text engines have expanded this use case, enabling businesses to apply sentiment analysis to phone calls and other voice interactions—often in real time—by converting spoken words to written text. The combination of speech to text and AI sentiment analysis can be applied to customer service calls, then offer timely suggestions on how to better engage with customers, helping resolve issues or highlight solutions that might benefit them. The technology is also being used to review recorded sales calls and provide agents with insights that improve future customer interactions.

It's important to keep in mind that while AI can make a pretty good guess as to whether someone is angry, frustrated, or just tired, interpreting sentiment is a highly imprecise science—and not just for machines. Even with our finely honed intuitions as to speech patterns and physical and vocal cues, studies show that different humans agree on sentiment only about 80% of the time. And we've all learned the perils of interpreting too much from a text message. But AI is on the cusp of becoming much more proficient in this subjective task. A qualitative leap will come when the sentiment engine takes in multimodal audio and visual cues, such as vocal inflection and facial expressions, that convey feelings beyond the plain text of the words spoken. That will empower AI to gauge the intangibles that make humans pretty good at reading other people: a subconscious drop in vocal tone, changes in pitch and intonation that correlate with stress or anger, or microexpressions and body language that betray frustration, disappointment, or pleasure.

AI is becoming more adept at interpreting human sentiment. This type of analysis, sometimes called “emotion AI,” has applications beneficial to all kinds of businesses, especially in the areas of customer service, sales, marketing, and product development.





An AI assist to developers

Developers can sometimes struggle to deliver performant, secure, and scalable applications on tight deadlines. They're also typically eager to take advantage of the latest technologies, which explains why software development teams are rapidly embracing AI as a powerful assistant for generating, reviewing, and documenting secure and consistent code.

GenAI-assisted code development tools are helping build feature-rich, mission-critical apps that consistently implement current design principles and security best practices. One analyst survey shows that about half of developers either already use or expect to use these tools—and that figure may be even higher among younger developers.

While this technology has seen significant advances in recent years, humans aren't being taken out of the equation, at least not anytime soon. At some point, it will likely become possible to describe the feel and functionality of an application in natural language and let AI generate several prototypes to choose from. But for now, it still takes skilled developers to design the underlying logic of an enterprise application, even if they ask the AI to then write the code. Still, GenAI-assisted code development along with low- and no-code tools have made it easier and less expensive for large and small businesses alike to rapidly build and deploy engaging websites and user interfaces to meet specific business needs.

But it's within a managed data environment that generative coding can really be unleashed.

When GenAI code development capabilities integrate with enterprise-grade data management systems, the AI gains understanding of the data it has available to work with and secure access to that data within an authenticated permission structure, freeing developers to spend more time imagining possibilities.

Software development teams are rapidly embracing AI as a powerful assistant for generating, reviewing, and documenting secure and consistent code.



Upgrade your upsell

Your AI may know your customers better than you do. And that means intelligent agents can enhance your upsell game in ways that are win-win.

AI models have been applied to upselling and attribute-based selling, through which consumers select specific features to customize their choices, for quite some time. We've all likely directly experienced these revenue-boosting tactics when booking hotels, buying event tickets, shopping on ecommerce platforms, or even watching movies on our favorite streaming service.

These algorithms anticipate customer interest by analyzing a trove of relevant data: buying patterns and past engagement, purchases by similar customers, market trends, social media output, and call center inquiries. But they've always had their limitations when it came to identifying products that reflect an individual's unique tastes and preferences.

Fashion is an excellent example of a product category that has traditionally been difficult for machines to effectively upsell. When you enter a store, an experienced retail assistant gets a sense of your personal style and suggests clothing or other items you might like. AI is now making headway in assessing these types of subjective preferences. As such, emerging AI-powered upsell platforms are beginning to empower businesses to sell higher-value goods in ways that were uniquely difficult for past iterations of the technology—driving sales without annoying potential customers with a barrage of unwanted offers.

HR's little helpers

Under the supervision of HR professionals, AI is helping organizations staff themselves with better-trained and more fulfilled employees.

This starts with targeted recruitment. GenAI can create custom career sites accessible across devices, help generate engaging job postings, even reply to candidates with personalized messages. AI tools are being used to place job notices where they're more likely to be seen by pools of applicants possessing the employer's desired experience and skills.

Say a new graduate, Maria, is looking to apply for a role at a large company. An AI agent can provide step-by-step instructions on how to complete the application, highlighting required documents and deadlines, help schedule an interview and send reminders, then generate regular application status updates so Maria knows where she is in the process. The agent saves time for recruiters, too, by answering questions and providing scheduling assistance

But hiring is only one step in successful staffing. AI agents embedded in human capital management software suites may assist in onboarding new employees by, for example, helping them submit required tax documents, understand and enroll in benefits plans, and access corporate resources and support. The technology can also be used to optimize mundane workforce management tasks, such as scheduling shifts.



GenAI can create custom career sites accessible across devices, help generate engaging job postings, even reply to candidates with personalized messages.

Employees are increasingly using AI to access career-development resources that lead them to opportunities for ongoing learning, upskilling, and effective networking. This empowers them to rise in the ranks, which can decrease churn, enhance employee satisfaction, and feed a pipeline of talent for senior positions.

Many organizations, especially those in industries that mandate learning curriculums to stay current with evolving practices, regulations, and policies, are leaning into AI to help train employees. They're using LLMs to write scripts for microcontent courses geared to specific audiences and text-to-speech engines to convert those scripts into voiceovers on onboarding or training videos. GenAI can also retrieve or create relevant images to accompany the training material, develop tests to validate knowledge, and certify completion.

Under the supervision of HR professionals, AI is helping organizations staff themselves with better-trained and more fulfilled employees.



Cybercops on the beat

AI is patrolling corporate networks day and night, holidays and weekends. Intelligent, autonomous agents constantly scan systems to detect anomalous behavior, suspicious logins, data exfiltration attempts, malware, and network traffic that hints at an intruder on the prowl. Unfortunately, AI is being used on both sides of the cybersecurity wars—attackers are increasingly relying on purpose-trained AI to probe for vulnerabilities and spread malicious software.

If GenAI has become a crucial weapon in the ever-escalating arms race between attackers and defenders of enterprise systems, why not look to a new crop of AI-powered security tools might tilt the balance of power to the good guys?

[Zero Trust Packet Routing](#), or ZPR, is one potentially game-changing AI-powered solution for protecting networks. With ZPR, security administrators can write natural language policies that limit traffic based on the resources and data services that employees need to access. ZPR then deploys AI to allow traffic flows among approved systems for approved uses.

Learn all about zero trust security, including how to get started with ready-made models and overcome common hurdles.



[Download the eBook](#)

Biometrics is another way that AI is enhancing the security of corporate IT systems. AI-powered voice, facial, and fingerprint recognition can eliminate the need for passwords, those 17-digit nuisances that are hard to remember and susceptible to attack.

AI-powered incident response tools are also becoming more sophisticated and essential in mounting a layered defense. When a cyberthreat is detected, every millisecond counts. AI agents can quickly coordinate and execute automated responses that contain the scope of the intrusion, limit the amount of data compromised, take measures to eradicate the threat, document the incident for further analysis, and hopefully minimize the overall damage to critical IT systems.

Fix before break

Predictive maintenance has long been preferable to the alternative of preventive maintenance, which relies on scheduled services and upkeep. And it's certainly better than reactive maintenance—fixing things after they break. But facilities and equipment managers have enjoyed only limited success in applying rules-based evaluation methods to predict when complex machinery was on the verge of failure.

Now, with AI monitoring your machinery, maintenance can finally become a preemptive affair. By ingesting large amounts of real-time data from sensors and smart devices, AI can spot trends that fall outside defined rules, delivering cost and time savings to businesses that rely on production machinery, office equipment, assembly and storage facilities, and fleets of vehicles to maintain operations.

With AI-powered predictive maintenance, machine learning trains models on the ideal functional state of your machines by monitoring them in real time and/or assessing historical data that can include performance characteristics, operating temperatures, acoustics, and infrared scans. The AI then watches for subtle deviations from ideal operations that otherwise would go undetected and, importantly, predicts the consequences of those deviations.

By alerting staff to conduct maintenance, this technology can greatly reduce unplanned downtime and avoid expensive repairs, as well as the reputational cost of missed deadlines. The lifespans of equipment, and the operations that rely on that equipment, can become far more reliable.





Ambient AI keeps watch

Imagine a world in which AI always surrounds you. You don't see the constant churning of algorithms, but they're there in the background, tracking your travels, monitoring your purchases, and more—surfacing only to offer help when they think it's needed.

This, no doubt, sounds like the premise of a dystopian sci-fi movie.

But AI that we want to run passively in the background, waiting for environmental cues to spring into action, will inevitably become a greater part of our personal and professional lives. And despite the discomfiting overtones of what's being called ambient AI—and legitimate concerns, especially around privacy—the potential benefits across a wide variety of use cases are very real.

The healthcare industry is becoming an early adopter of ambient AI with the aim of freeing time for doctors and nurses while helping record important medical details. During a clinical consultation, the ambient AI listens in on the interaction between doctor and patient: complaints of symptoms, results of examinations, prescriptions, and recommendations for follow-up treatment. After they finish, the AI creates a report extracting the important details, sends it to the doctor for review, and files it in the patient's electronic health record. It might even suggest a diagnosis for the doctor to review. Meanwhile, patients leave confident that the doctor focused on their needs and their medical issues were properly documented.

This basic use case can be applied to a variety of businesses. Ambient AI can be an executive assistant that's always there. During product demonstrations, the technology could transcribe speech to text in real time; find and bring up on a screen relevant charts, documents, and images; and then create and circulate a detailed report summarizing key points and conclusions. Or on customer service calls, ambient AI running in the background, without responding to any specific query, can suggest useful information to an agent, present relevant reference materials or documentation, suggest follow-up questions, and formally document the resolution of the interaction.

How to get started

If you're an Oracle customer, a range of AI capabilities are available to you today. Oracle not only has a long history of working with artificial intelligence technologies and incorporating AI into its products, it's also at the forefront of development.

GenAI is a fully managed Oracle Cloud Infrastructure (OCI) service that provides a set of state-of-the-art, customizable LLMs that cover a wide range of use cases, including chat, text generation, summarization, and creating text embeddings. Dozens of AI agents embedded in Oracle Fusion Cloud Applications—with more being added regularly—enable enterprises to apply GenAI where they need it, without leaving the ERP, sales, marketing, and other applications they use every day.

You can try OCI for free to test out ready-to-use pretrained models or create and host your own fine-tuned custom models based on your own data on dedicated AI clusters.

[🔗 Oracle Cloud Infrastructure](#)



AI agents

With Oracle Cloud Infrastructure (OCI) Generative AI Agents, you can leverage a conversational interface to query enterprise data stores in natural language. This combines the power of large LLMs and RAG with your own data, providing real-time information with the ability to act directly on it.

[🔗 Learn more](#)



Generative Development

Introduced in Oracle Database 23ai, generative development for the enterprise is a groundbreaking AI-driven application development environment. The innovative AI-backed technologies help developers rapidly create sophisticated applications with enterprise-grade reliability, scalability, and security.

[🔗 Learn more](#)



OCI Zero Trust Packet Routing

OCI Zero Trust Packet Routing helps prevent unauthorized access to data by managing network security policy separately from the underlying network architecture.

Using an easily understood and intent-based policy language, security administrators can define specific access pathways for data.

[🔗 Learn more](#)

How Oracle can help

Organizations large and small have made tremendous inroads using GenAI, but this technology can do so much more. As you set your strategy, we invite you to explore Oracle's comprehensive AI services and state-of-the-art innovations on our data platform and in our cloud applications—all on a best-in-class AI infrastructure.

Register for AI workshops

Explore our AI solutions

Connect with us

Call +1.800.ORACLE1 or visit oracle.com

Outside North America, find your local office at oracle.com/contact

Copyright © 2025 Oracle, Java, MySQL and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

