Atlas for Industries

Innovation in Insurance with Artificial Intelligence

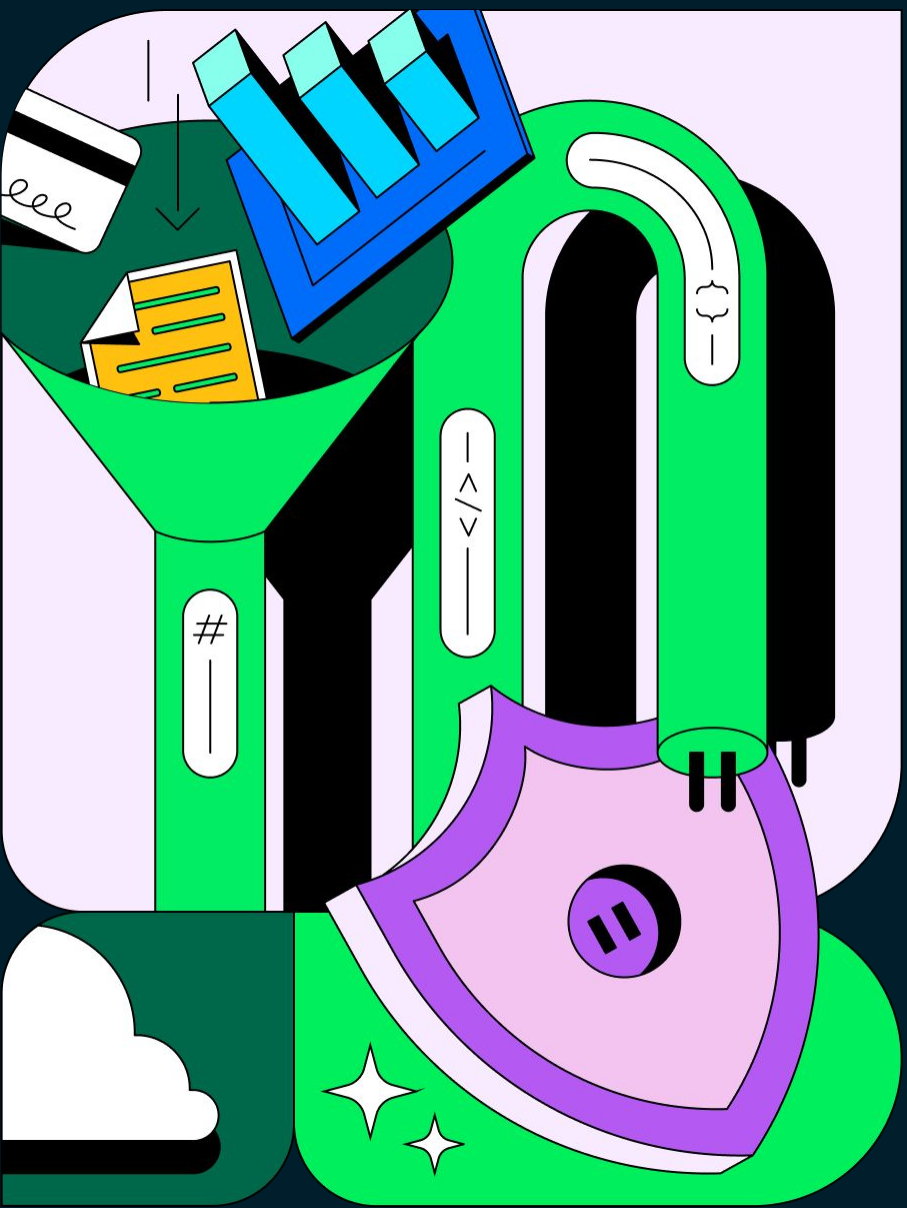

With its ability to streamline processes, enhance decision-making, and improve customer experience with far less time, resources, and staff than traditional IT systems, AI offers insurers great promise.

Architecting with AI for the Insurance Industry

Delivering business value

When architecting for AI-augmented application workflows, it's important to keep your overall business objectives in mind. How will what you do help the insurance business you support? What problem are you solving?

IT portfolio managers who may be driving platform consolidation and modernization agendas inside of your organization may be approaching the space from a technology-first lens, evaluating one platform vs. another in terms of technical capabilities and cost. Data science and analytics folks may be wishing to test or prove data hypotheses. Business product owners may have priorities in terms of raising the bar when it comes to your business applications and their features, hoping to augment their data processing capabilities with the addition of AI.

For you, as the architect or CxO, your charge is to help the organization leverage AI to make meaningful strides in terms of business outcomes, and that means navigating a number of dimensions, not limited to:

- **Business objectives:** A clear understanding of the organization's goals
- **Data and workflows:** Knowledge of business processes and supporting data
- **AI application:** Applying AI to streamline data-intensive tasks

Our end goal is to be able to gather, understand, interact with, and generate data faster by applying the right AI capabilities to the right points in the organization.

Claim as an example

Common organizational goals for an insurer may include:

- **Operational excellence:** Emphasizing efficiency and effectiveness in operations to maximize returns and reduce waste
- **Customer centricity:** Prioritizing investments that improve customer satisfaction and engagement

How well your organization can process and resolve claims, for example, directly impacts the objectives above. Achieving this is directly coupled to how quickly, efficiently, and accurately we can process the data found within claim-handling application workflows.

Claim handling provides us great opportunities to exploit AI in order to accelerate data processing hotspots, so that the organization can return meaningful value from the technology investment.

So what data in a claim-handling workflow is hard to work with? Unstructured data sources such as damage photos, accident forms and reports, claim handler notes, traffic camera videos, and claim-handling guidelines and recommendations. Are any or all of these data sources cumbersome to work with for your employees? Do they need to open and read forms, examine and interpret images, and distill and write case-file notes before your claims can move forward in the handling lifecycle? If a catastrophic event strikes, do you have the staff to handle a massive, sudden influx of claims, all with these types of unstructured data?

Augmenting with AI

Architects are the ones tasked with deciding what type of AI is best suited to solve for insurance data processing needs.

Machine learning models trained on historical data can be used to make predictions and decisions within a business workflow, in lieu of a person.

Generative AI and large language models, however, give us core NLP (natural language processing) capabilities that are particularly well-suited to augmenting data processing abilities. These capabilities, as applied to the claim-handling workflow, include:

Entity extraction to help interrogate and retrieve relevant information from an unstructured source such as a PDF guideline, or a large body of text found within an accident form.

Text and image classification enables claim handlers to determine types of damage or characteristics found on a damage photo.

Text summarization speeds up the synthesis of large bodies of text, or across many sources.

Text generation helps generate case files and give succinct instruction to workers based on larger bodies of guidelines, for example.

Interactive chat makes the previous or additional existing information available more quickly to both your employees and your customers.

Architecting for applications

Your software delivery teams and the applications they support may be divided up by agile delivery domains. It's inside of those domains and applications that you'll need to apply AI in order to effectively move the needle on organizational and process outcomes. In short, your AI belongs in your applications.

The data that supports those applications resides in operational data stores. If we want both our applications and our AI to work with real-time data, it should be accessible within the same underlying data store. What serves our apps should also serve our AI.

Managing operational structured *and* unstructured data:
Your application architecture needs to provide a way to store, serve, and update structured data as part of your workflow and integrate the unstructured data. This may come in the form of both the raw data (PDF's, images, notes), and the vector-encoding representation of it. You'll want the vector-encoding data to be as close as possible to the structured data that already may exist in your application. Several reasons why include:

- **Performance:** Providing application-level SLA (service level agreements) utilizing efficient server compute
- **Security:** Providing consistent app-tier level security controls to determine who can access what data.
- **Ease of application delivery and maintenance:** The fewer components and dependencies, the lower the complexity and cost of building, deploying, and maintaining effective software solutions that leverage AI.

Applying AI to the claim-handling workflow and delivering better business outcomes

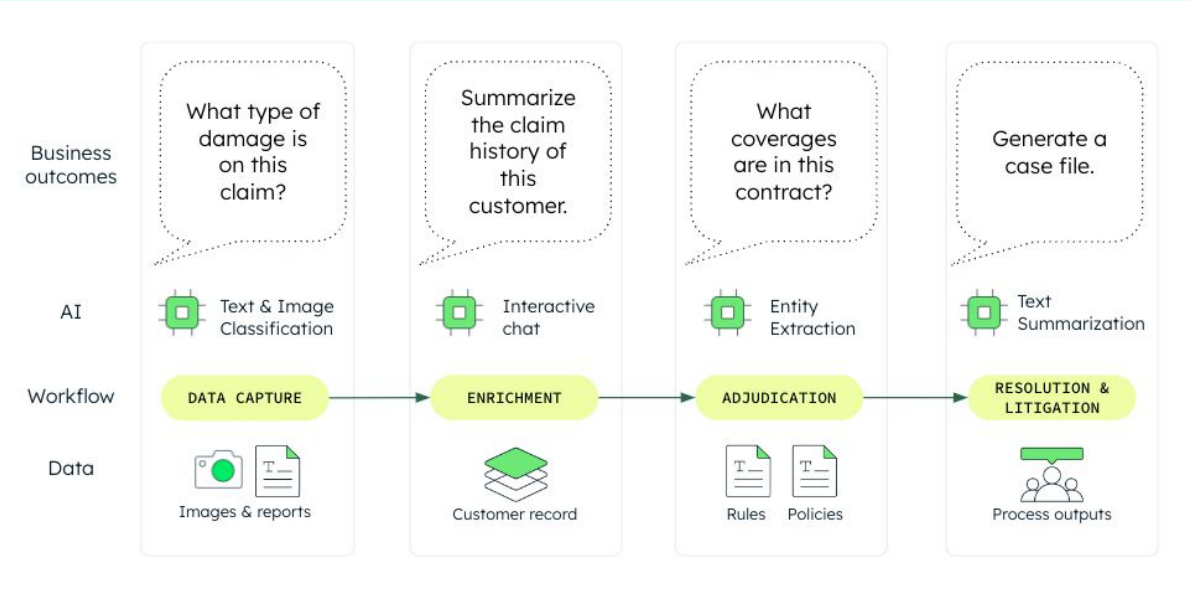

Figure 1: AI use cases aligned to the claim-handling workflow

The above diagram shows practical use of core NLP capabilities that include text and image classification, interactive chat, entity extraction, and text summarization. When applied to a claim-handling workflow, for example, these capabilities can reduce data hotspots, resulting in lower processing times and costs, and improved customer experiences.

Before AI can transform our organizations, we first have to bring it into our **applications**

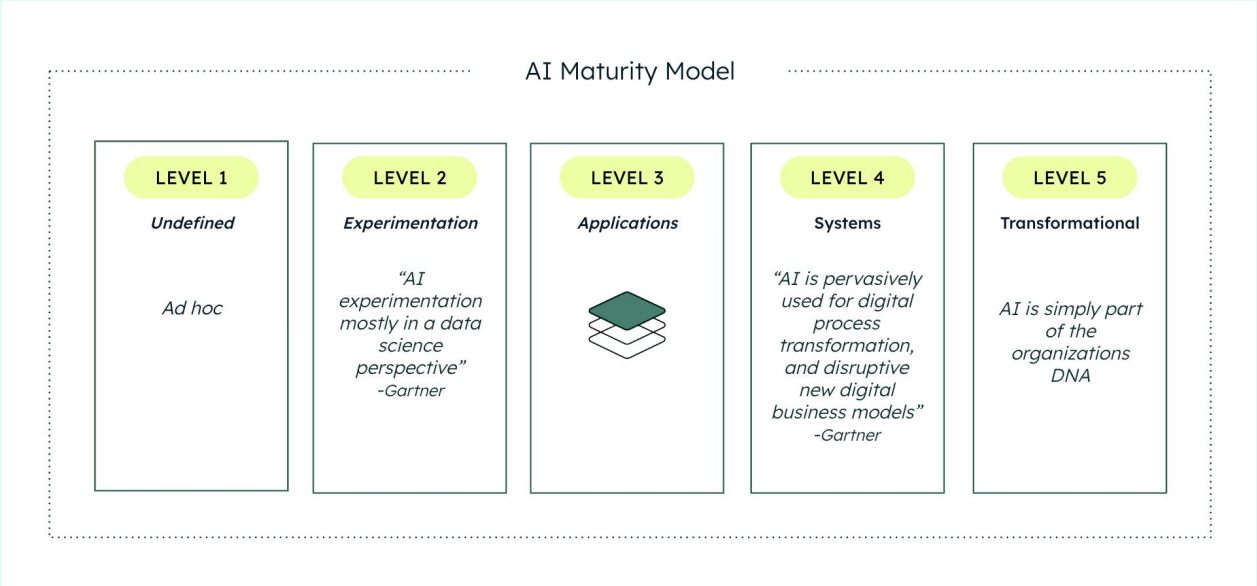

Figure 2: AI Maturity levels from undefined to transformational

The above diagram describes the different stages of AI adoption within enterprises, from early interest to a pervasive and structural integration into processes and decisions. Many organizations struggle with moving out of Level 2, analytics experimentation, to Level 3, deploying AI features within business applications, in order to deliver meaningful business value and outcomes.

The three layers of generative AI

Gen AI applications can be separated into three major layers:

Compute and AI models: Underlying processing power plus foundational and embedding models.

Tooling to fine tune and build applications: Tools that provide foundational models with context by feeding them proprietary data.

AI-powered apps and experiences: The interface and experiences that end-users interact with, and app frameworks that simplify the process of building AI experiences.

Foundational models are extremely powerful, but being trained on public datasets, they lack the domain knowledge and data context needed to adequately support enterprise applications.

This is where the second layer comes into play: the data and tooling that will glue everything together and enable generative AI-enhanced applications to be fully operational, moving your organization from the experimentation level of Level 2 into full-fledged production apps of Level 3 maturity.

MongoDB's operational database stores proprietary structured and vector data, and makes it available to LLMs when an application makes a request, effectively providing the context necessary for a foundational model to answer questions beyond its initial knowledge boundary.

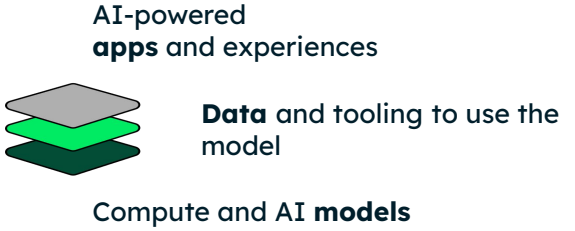

Building applications that leverage AI

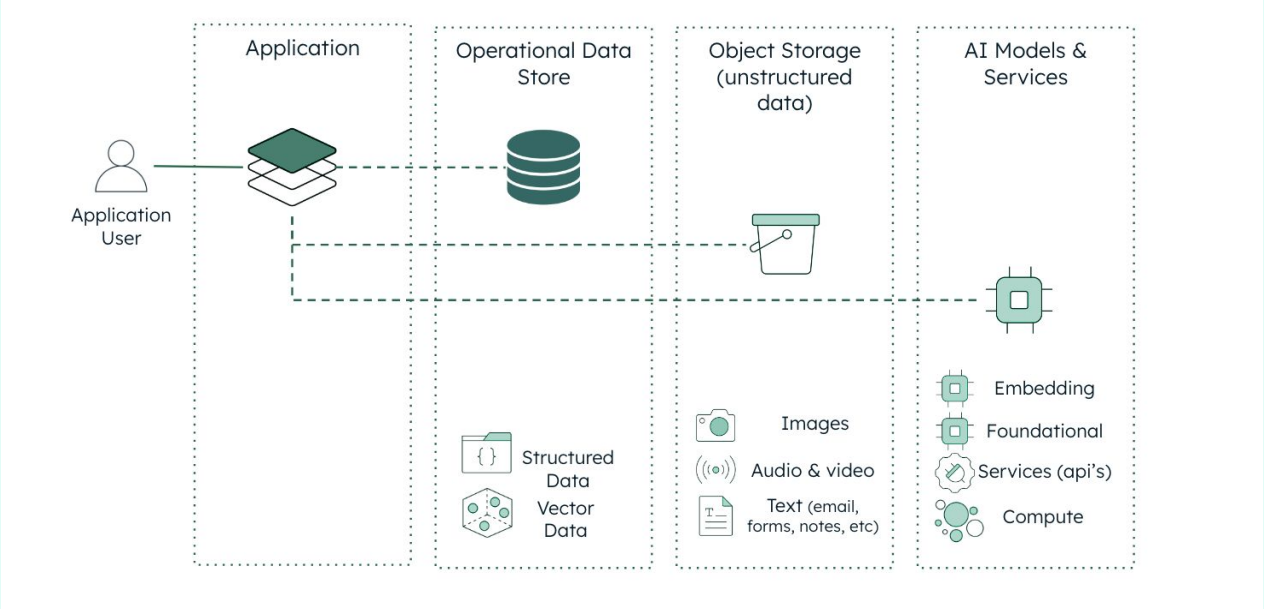

Figure 3: Applications with AI

The above diagram illustrates in greater detail the three layers of gen AI, showing the structured data flowing from the application into the operational database, raw unstructured data managed in object storage, accessible by applications where it can be processed. Part of the processing includes the vectorization and subsequent persistence of those vectors in the operational data store, where it can be readily accessed by the application.

In an inherently information-driven industry, insurance companies ingest, analyze, and process massive amounts of data. Whether it's agents and brokers selling more policies, underwriters adequately pricing, renewing and steering product portfolios, claim handlers adjudicating claims, or service representatives providing assurance and support, data is at the heart of it all.

Given the volumes of data, and the amount of decision-making that needs to occur based on it, insurance companies have a myriad of technologies and IT support staff within their technology investment portfolios. It's no surprise that AI is at the top of the list when it comes to current or prospective IT investments. With its ability to streamline processes, enhance decision-making, and improve customer experiences with far less time, resources, and staff than traditional IT systems, AI offers insurers great promise.

Underwriting & Risk Management

Few roles within insurance are as important as that of the underwriters who strike the right balance between profit and risk, bring real-world variables to the actuarial models at the heart of the insurer, and help steer product portfolios, markets, pricing, and coverages. Achieving equilibrium between exposures and premiums means constantly gathering and analyzing information from a myriad of sources to build a risk profile sufficient and detailed enough to make effective policy decisions.

While many well-established insurers have access to a wealth of their own underwriting and claims data, integrating newer and real-time sources of information, keeping up with regulatory changes, and modeling out what-if risk scenarios still involve significant manual effort.

Perhaps the single greatest advantage of AI will be its ability to quickly analyze more information with fewer people and resources. The long-term impact will likely be profound, and there is tremendous promise within underwriting.

Advanced analytics

Traditional IT systems are slow to respond to changing formats and requirements surrounding data retrieval. The burden falls on the underwriter to summarize data and turn that into information and insight. Large language models are now being leveraged to help speed up the process of wrangling data sources and summarizing the results, helping underwriting teams make quicker decisions from that data.

Workload and triage assistance

AI models are mitigating seasonal demands, market shifts, and even staff availability that impact the workload and productivity of underwriting teams, saving underwriting time for the high-value accounts and customers where their expertise is truly needed. Amid high volumes for new and renewal underwriting, traditional AI models can help classify and triage risk, sending very low-risk policies to touchless automated workflows, low to moderate risk policies to trained service center staff, and high-risk and high value accounts to dedicated underwriters.

Decision-making support

Determining if a quoted rate needs adjustment can take significant time and manual effort. So can preparing and issuing renewals of existing policies, another large portion of the underwriters' day-to-day responsibilities. Automated underwriting workflows leveraging AI are being employed to analyze and classify risk with far less manual effort. This frees up significant time and intellectual capital for the underwriter.

Vast amounts of data analyzed by underwriters are kept on the underwriters desktop rather than IT-managed databases. MongoDB offers an unparalleled ability to store data from a vast amount of sources and formats, and respond quickly to requests to ingest new data. As data and requirements change, the document model allows insurers to simply add more data and fields without the costly change-cycle associated with databases that rely on single, fixed structures.

For every major business entity found within the underwriting process, such as broker, policy, account, and claim, there is a wealth of unstructured data sources, waiting to be leveraged by generative AI. MongoDB offers insurers a platform that consolidates complex data from legacy systems, supports modern applications, and extends those same data assets to AI-augmented workflows. By eliminating the need for niche databases for these AI-specific workloads, MongoDB reduces technology evaluation and on-boarding time, development time, and developer friction.

LEARN MORE

Automating Digital Underwriting with ML*

*mongodb.com/solutions/solutions-library/machine-learning-underwriting-solution

Claim Processing

Efficient claim processing is critical for an insurer. Timely resolution of a claim and good communication and information transparency throughout the process are key to maintaining positive relationships and customer satisfaction. In addition, insurers are on the hook to pay and process claims according to jurisdictional regulations and requirements, which may include penalties for failing to comply with specific timelines and stipulations.

In order to process a claim accurately, a wealth of information is needed. A typical automobile accident may include not only verbal and written descriptions from claimants and damage appraisers but also unstructured content from police reports, traffic and vehicle dashboard cameras, photos, and even vehicle telemetry data. Aligning the right technology and the right amount of your workforce in either single- or multi-claimant scenarios is crucial to meeting the high demands of claim processing.

Taming the flood of data

AI is helping insurers make sense of a trove of data faster and in real time. From natural language processing to image classification and vector embedding, all the pieces of the puzzle are now on the board for insurers to make a generational leap forward when it comes to transforming their IT systems and business workflows for faster information processing.

Claims experience

Generating accurate impact assessments for catastrophic events in a timely fashion in order to inform the market of your exposure can now be done with far less time, and with far more accuracy, by cross-referencing real-time and historical claims experience data, thanks to the power of generative AI and vector-embedding of unstructured data.

Claim expediter

Using vector-embeddings from photo, text, and voice sources, insurers are now able to decorate inbound claims with richer and more insightful metadata so that they can more quickly classify, triage, and route work. In addition, real-time insight into workload and staff skills and availability is allowing insurers to be even more prescriptive when it comes to work assignments, driving towards higher output and higher customer satisfaction.

Litigation assistance

Claim details are not always black and white, parties do not always act in good faith, and insurers expend significant resources in the pursuit of resolving matters. AI is helping insurers drive to resolution faster and even avoid litigation and subrogation altogether, thanks to its ability to help analyze more data more effectively and in less time.

Risk prevention

Many insurers provide risk-assessment services to customers using drones, sensors, or cameras to capture and analyze data. This data offers the promise of preventing losses altogether for customers and lowering exposures, liability, and expenses for the insurer. This is possible thanks to a combination of vector embedding and traditional and generative AI models.

LEARN MORE

AI-Enhanced Claim Adjustment for Auto Insurance

*mongodb.com/solutions/solutions-library/vector-image-search

Customer Experience

Accessing information consistently during a customer service interaction, and expecting the representative to quickly interpret it, are perennial challenges with any customer service desk. Add in the volume, variety, and complexity of information within insurance and it's easy to understand why many insurers are investing heavily in transformation of their customer experience call center systems and processes.

24/7 virtual assistance

As with many AI-based chat agents, the advantage virtual assistants is that they can free up your call center staff to work on more complex and high-touch cases. Handling routine inquiries can now extend to more complex scenarios than ever thanks to the power of vector-embedded content and large language models.

Customer profiles

Every interaction is an opportunity to learn more about your customers. Technologies such as voice-to-text streaming, vector embedding, and generative AI help insurers build out a more robust social profile of their customers in near real-time.

Claims assistance

Generative AI can deliver specific claim-handling guidelines to claim-handling staff in real time, while traditional ML models can interrogate real-time streams of collected information to alert either the customer or the claim-handler to issues with quality, content, or compliance. AI capabilities allow insurers to process more claims faster and significantly reduce errors or incomplete information.

Real-time fraud detection

According to estimates from the Coalition Against Insurance Fraud, the U.S. insurance industry lost over $308 billion to fraud in 2022. With vector-embedding of unstructured data sources, semantic and similarity searches across both vector and structured metadata, and traditional machine learning models, insurers can detect and prevent fraud in ways that were simply not ever before possible.

Domain-Driven AI

Insurance enterprises all over the world have embraced domain-driven design in addition to aligned software delivery teams and the applications they support by core processing domains. With a shift to microservices and event-streaming, AI capabilities can now supercharge this architecture and accelerate the ability to interact with and serve up real-time data. Domain-based operational data stores help segment and enable stronger data lineage, data quality, and data governance for more truthful AI interactions.

Working together: Applications, data & AI

Core insurance business capabilities organized and deployed by domain and domain capabilities augmented by AI:

- Domain- and task-specific AI, enhancing key application workflows and data processing hotspots
- Operational and vector data stored as close to applications as possible for:
  - Real-time context
  - Performance
  - Security & privacy
  - Domain agility
- APIs and events for efficient processing within and across domains.

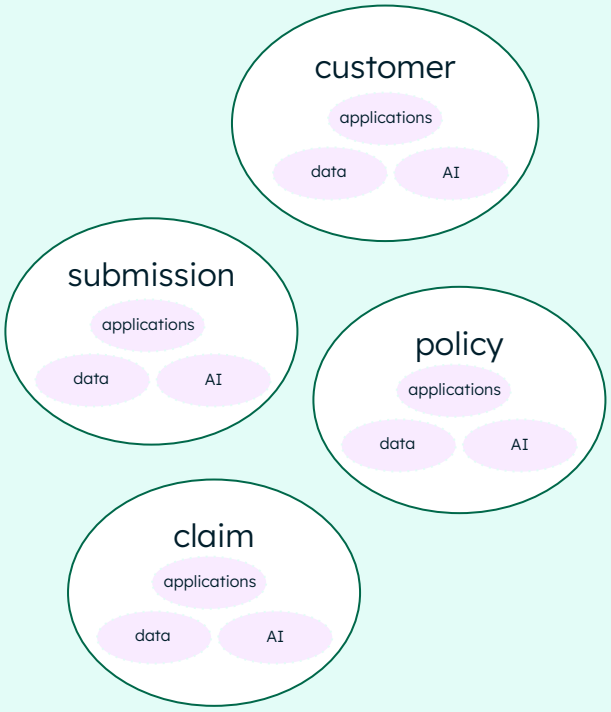

Examples of Domain-Specific AI

ferret

Ferret.ai builds trust with AI & MongoDB Atlas, slashes costs 30%

Ferret.ai helps companies and individuals build trust by providing real-time insights on people and businesses using AI. Powered by MongoDB Atlas, Ferret analyzes vast amounts of public data to identify risks and opportunities, offering a comprehensive view of relationships.

"Through Atlas, it's available as a fully managed service with best practices baked in. That frees developers and data scientists from the responsibilities of running a database so they can focus their efforts on app and AI innovation" Learn more*

Al Basseri

CTO and head of Engineering at Ferret.ai

**mongodb.com/blog/post/building-ai-mongodb-navigating-path-from-predictive-to-generative-ai*

VISO TRUST

VISO TRUST: Accelerating third-party risk assessment with AI

VISO TRUST is an AI-powered third-party cyber risk and trust platform that enables any company to access actionable vendor security information in minutes. By leveraging MongoDB Atlas for efficient data storage and retrieval, VISO TRUST's advanced AI models can process vast amounts of security information, delivering actionable insights in minutes. This streamlined approach significantly reduces assessment time and enhances overall cybersecurity posture. The platform uses sophisticated models and retrieval-augmented generation (RAG) techniques to provide highly accurate and contextually relevant intelligence. This not only accelerates decision-making but also ensures that companies are equipped with the most precise risk evaluations available. Assessments generated from analysis are 80% faster than manual methods with no loss in accuracy

Learn more*

**mongodb.com/blog/post/building-ai-mongodb-cultivating-trust-data*

Solution Library References

The Solutions Library is a curated collection of tailored solutions, offering ready-made templates, best practices, and detailed guides to help developers quickly start projects. With industry-specific examples and sample code, the library enables insurance companies to efficiently implement common patterns, leverage MongoDB's features effectively, and ensure compliance with industry standards. Here are some examples of solutions for the insurance industry.

Claim management using LLMs and vector search for RAG

By converting claim data into vector embeddings, MongoDB's Atlas Vector Search accelerates information retrieval, making it quicker and easier to find relevant details. Large language models (LLMs) then analyze these embeddings to extract valuable insights and context, optimizing claim processing. This combined approach enhances accuracy, efficiency, and overall claims management.

Learn more*

**mongodb.com/solutions/solutions-library/claim-management-llms-vector-search*

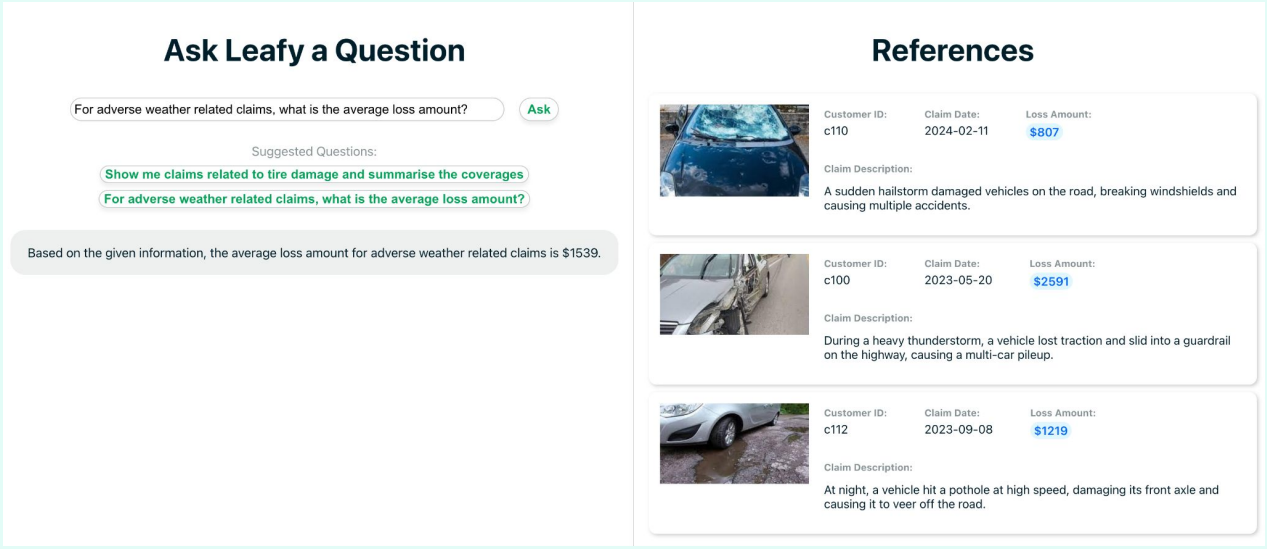

Figure 4: Atlas Vector Search provides answers to user questions about insurance claims, including calculations and detailed claim examples

AI-enhanced claim adjustment for auto insurance

By leveraging AI and vector image search, this solution **automates auto insurance claim adjustments**. Accident photos are compared to a database of past claims, significantly accelerating damage estimates.

Centralized data management in MongoDB Atlas enhances efficiency and consistency throughout the claims process.

Learn more*

*mongodb.com/solutions/solutions-library/vector-image-search

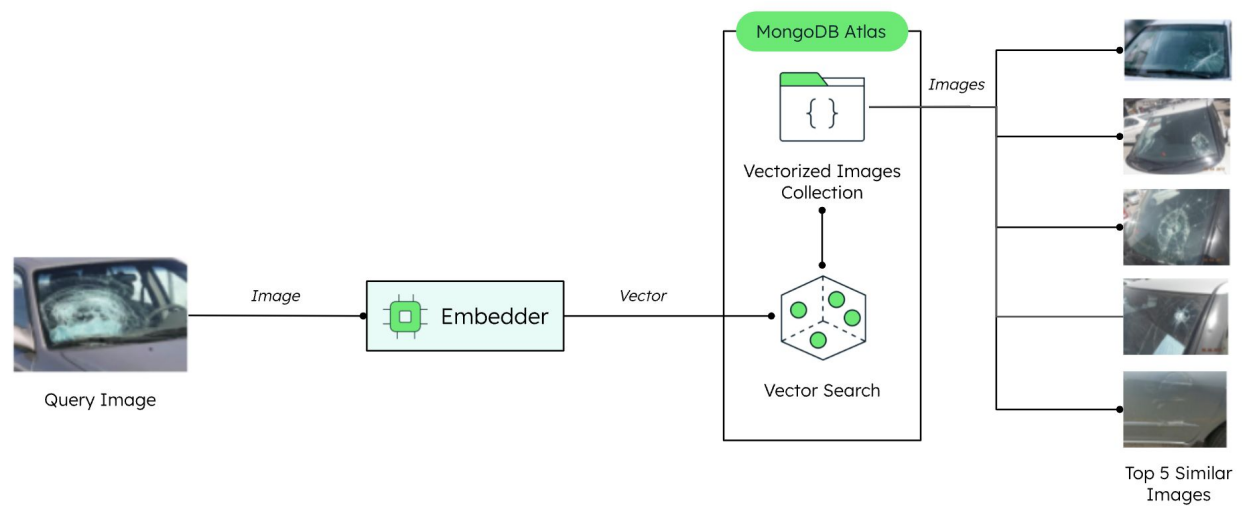

Figure 5: An image similarity query is performed, and the 5 top similar images are returned

Build a PDF search application with vector search and LLMs

PDFs are hard to search, making it tough for insurance workers to find what they need quickly. We're solving this by converting PDFs into a searchable format using tools like Superduper, allowing users to quickly retrieve information and streamline insurance work.

Learn more*

*mongodb.com/solutions/solutions-library/pdf-search-with-vector-search-and-llm

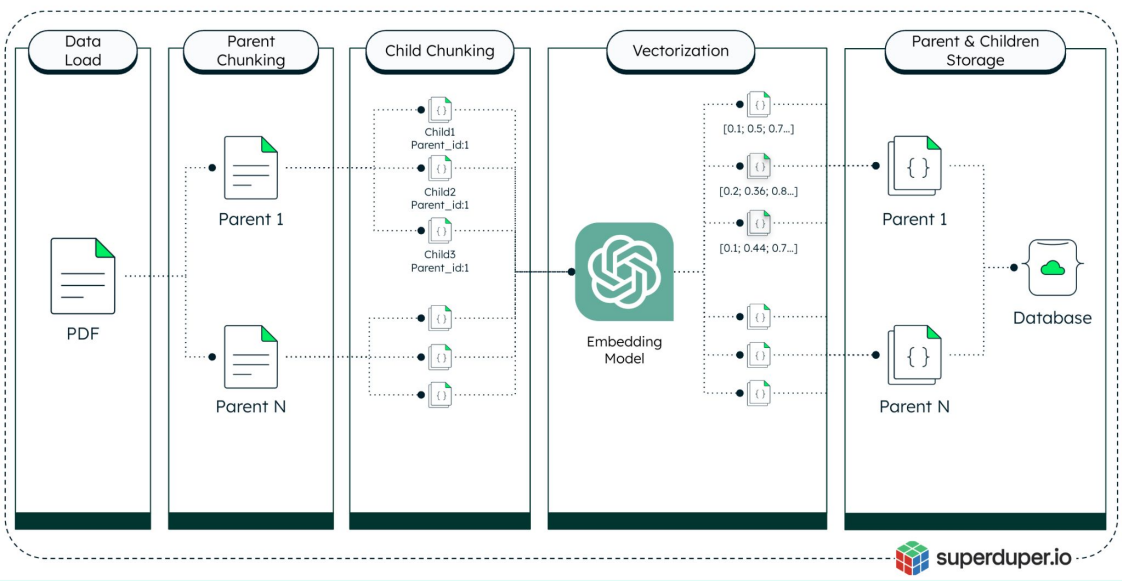

Figure 6: PDF chunking, embedding creation, and storage, orchestrated with Superduper

Other Notable Use Cases

Predictive analytics

AI-powered predictive analytics can anticipate customer needs, preferences, and behaviors based on historical data and trends. By leveraging predictive models, insurers can identify at-risk customers, anticipate churn, and proactively engage with customers to prevent issues and enhance satisfaction.

Predictive maintenance for property insurance

AI-powered predictive maintenance solutions, leveraging IoT sensors installed in buildings and infrastructure, are used in property insurance to prevent losses and minimize damage to insured properties.

Crop insurance and precision farming

AI is being used in agricultural insurance to assess crop health, predict yields, and mitigate risks associated with weather events and crop diseases, which helps insurers offer more accurate and tailored crop insurance products to farmers.

Usage-based insurance (UBI) for commercial fleets

AI-enabled telematics devices installed in commercial vehicles collect data on driving behavior, including speed, acceleration, braking, and location. Machine learning algorithms analyze this data to assess risk and determine insurance premiums for commercial fleets to help promote safer driving practices, reduce accidents, and lower insurance costs for businesses.

Contact Information

Jeff Needham

Insurance Industry
Solutions Principal

jeff.needham@mongodb.com

FOR MORE INFORMATION AND RESOURCES

Visit MongoDB Atlas for Insurance*

**mongodb.com/solutions/industries/insurance*