

IBAC AI Use Case Pilot

Cemiar & La Turquoise “The Perfect Policy”

AI Use Case Study Report





AI Use Case Project: Overview

AI POC Use Case #1: The Perfect Policy

AI Use Case Description

This project aims to address the "coverage gap" challenge, which is the risk exposure resulting from missing or insufficient coverage in a customer's insurance policies when compared to their actual needs. We plan to leverage Artificial Intelligence (AI), specifically Optical Character Recognition (OCR) technology, to automatically process and analyze large volumes of policy documents. The OCR will extract key data points, terms, and conditions, allowing the system to then compare the existing coverage details against a set of predetermined optimal coverage standards for various policy types. This comparison will efficiently identify and flag any instances of missing or inadequate coverage, enabling insurers or brokers to proactively advise clients and close the identified gaps.

Stakeholders

- Broker Name: **Laturquoise**
- Vendor Name: **Cemiar Inc.**

Project Specifics

- Duration: **4 month**
- Cost: **\$30,000 plus volume-based monthly fees.**
- Complexity: This project has a Medium to High level of effort and cost. The effort is driven by the necessity for high-accuracy OCR across diverse policy document types and the subsequent development of complex NLP/AI models to interpret policy language and encode expert coverage gap logic. Successful deployment requires extensive data annotation, testing, and secure integration with existing insurance systems.
- Additional Information: Secure and reliable integration with the Broker Management System (BMS) is critical for proactive proposal generation and seamless workflow execution, demanding significant technical effort and rigor. In addition, OCR and table recognition rely on advanced algorithms to extract and structure data accurately, further contributing to the overall technical complexity.

Project Outcomes

- Measurable results -> Renewal review effort reduced by 80%
- Impact on key business metrics -> Redeployment of staff projected to increase annual NB volume by > 1MM GPW. Tool accuracy ~ 99%. Zero hallucinations
- Insights gained -> Early involvement of end users is crucial to adoption success



AI Use Case Project: Deep Dive

Business Opportunity

- What is the core business challenge or pain point this AI use case aims to address?

The core challenge is the missed opportunity to upsell, the amount of time due to analysis and the risk of Errors & Omissions (E&O) liability due to brokers failing to proactively identify and close existing coverage gaps in client policies.

- Why does this problem exist within the current processes or systems?

This problem exists because manually reviewing complex, diverse policy documents for subtle exclusions and comparing them against optimal benchmarks is a time-consuming, highly error-prone process for human brokers.

- Why is solving this problem critical for the organization's success?

Solving this problem is critical for driving revenue growth through automated, accurate cross-selling and enhancing client trust and retention by proactively mitigating their financial exposure, thereby strengthening the organization's market reputation.

- Are there any missed opportunities or risks associated with not addressing this issue?

Not addressing this risks significant financial losses from potential E&O lawsuits stemming from uninsured client claims, and the loss of revenue from missed opportunities to offer necessary, profitable additional coverage.

Business Impact

Expected Benefits

- What measurable outcomes are projected?

Is projected to reduce analysis time by 80%, enhancing broker productivity and operational capacity.

80% of policies analyzed by the system, significantly reducing the risk of unidentified coverage gaps.

- What operational improvements are expected?

We anticipate a 15% reduction in operational costs, as brokers and broker assistants will be able to process files more rapidly.

Consistency in analysis will be maintained even in the event of unforeseen staff absences, ensuring workflow continuity.

- How will this impact employee productivity or engagement?

- **We also expect to see a positive impact on the sale of supplementary products, although no specific percentage increase has been established at this point.**

Actual Benefits

- What quantifiable results were achieved compared to projections?

The success rate of the comparators is above 95%. This new tool therefore provides sufficient reliability for our brokers to use it confidently.

- Were operational efficiencies realized (e.g., faster processes)?

We are confident that we can save 80% of the analysis time by combining our two AI projects, Comparator and Perfect policy, which significantly reduces the time required to analyze a renewal.

- Did the project deliver intangible benefits?

The analysis process is much faster. In addition, missing coverages are easier to identify, which helps reduce errors and ensures that all protections our clients are entitled to are offered. This also allows us to reinvest that time in business development. Clients also receive better service, as the broker can focus on their advisory role and add coverage.



AI Use Case Project: Deep Dive

Scope of POC Implementation

- Where will the AI solution be applied—enterprise-wide, specific departments, or just a pilot?

The solution will be applied within a specific business area: Commercial Lines. Although it can be standalone, the primary application is integration with the automatic commercial billing system provided by Cemiar.

- Which processes or workflows will this solution impact?

The solution will primarily impact the workflow phase during billing where coverage validation and gap analysis are traditionally performed by brokers

- Is the scope limited to certain business areas (e.g., commercial lines, submission management)?

Yes, the scope is strictly limited to Commercial Lines policies and related client documentation.

- Which users and departments will be involved?

The system will directly involve end-users, primarily the brokers and administrators responsible for client policy management and validation.

- Approximately how many users will be impacted?

The solution is expected to directly impact a core group of approximately 50 to 60 brokers within the organization.

- Are there any dependencies on existing systems or platforms that need to be considered?

The solution's primary dependency is the integration with the Cemiar commercial billing automation, but, a stand-alone drag-and-drop interface is also available.

Technical Solution Summary

- Is the solution SaaS-based, vendor-hosted, client-hosted, or broker-hosted?

The solution is a SaaS (Software-as-a-Service) offering, meaning it is hosted and maintained externally by the vendor, Cemiar.

- Does the solution leverage an AI model, and if so, which one (e.g., off-the-shelf LLM, custom model)?

The primary AI technology leveraged is OCR (Optical Character Recognition), utilizing both third-party tools like Azure Intelligence and Tesseract.

- Does the solution include additional logic or automation beyond the AI model?

Yes, the solution includes proprietary logic: a missing coverage algorithm and a database/module containing coverage definitions for each policy type.

- What systems does the solution connect to, and what data sources does it leverage?

It connects to the Broker Management System (BMS) using the Applied Epic SDK and leverages policy PDF files from the BMS along with OCR providers.

- What are the key architectural components (e.g., chat interface, APIs, database, BMS connectivity)?

Components include the Applied Epic SDK for file fetching, the AI layer (OCR reading), and the Core Logic Engine that uses client-provided templates and algorithms to identify missing coverage.

- Are there any anticipated challenges related to data quality, compliance, or integration?

Key challenges involve achieving high OCR reading accuracy on diverse document formats and ensuring the coverage templates used for validation are precise and up-to-date



AI Use Case Project: Employee Journey Map BEFORE

Renewal reception

Renewal is received in the system and the comparator is generate, available in the client file

Renewal analysis

- The broker or broker assistant analyze the renewal, it typically takes them 5 to 15 minutes to analyze the current policy alongside the renewal.
- He identify manually the differences and the missing or insufficient coverages.

File update

The broker will do an update with the insurer, offer the missing coverage (s) and notify the client of the important differences identified

Price request

The broker or the assistant broker will ask the price of the missing coverage to the insurer and send the update file

AI Use Case Project: Employee Journey Map AFTER

Renewal reception

Renewal is received in the system and the comparator is generate, available in the client file

Renewal analysis

The broker or the broker assistant will take 1 to 3 minutes to analyze the results of the comparison


The AI project practically erase the analysis task and reduce the risk of E&O

File update

The broker will do an update with the client and offer the missing coverage identified or offer to increase the insufficient coverage

Price request

The broker or the assistant broker will ask the price of the missing coverage to the insurer and send the update file



AI Use Case Project: Resume of the project - Phase 1 & Phase 2

The broker or broker's assistant receives the renewal and takes 5 to 15 minutes to compare it with the current policy to identify any differences, they determine any missing coverages, then contact the client to update the file and offer the appropriate protections.

In a previous AI project at La Turquoise, we included a comparative table highlighting the differences in limits, coverages, exclusions, conditions, and other elements between the current policy and the renewal. Based on this table, our new project identifies the missing coverages or limits required to achieve the 'perfect policy', according to each insurer. The entire project will therefore significantly reduce brokers' analysis time, enable them to review more files within the same timeframe, and allow them to quickly request pricing from insurers for the missing coverages.

The complete policy review will require approximately 1 to 3 minutes for the broker, instead of the current 5 to 15 minutes, before proceeding to the next step



AI Use Case Project: Implementation Deep Dive

Step 1: Design

The objective of this project was to build a service that validates insurance coverages within policy PDF documents based on configurable rules defined by clients. Stakeholders included insurance brokers and employees from La Turquoise, who helped align requirements and validate the approach.

During the design phase, multiple policy PDFs and relevant coverage keywords were collected to support the validation logic. Architecture decisions were reviewed with the Cemiar team using diagrams and schemas.

A key challenge was the large variation of keywords in policy documents, requiring the configuration to be improved iteratively. The design phase took a few weeks to complete.

Step 2: Development

During development, a user interface was implemented to let clients configure rules for validating policy coverages. The service was integrated as a plugin within our Commercial Billing system.

Text extraction from policy PDFs used Azure Tesseract (OCR). For data selection, we performed a data filtering process to select high-quality documents relevant to the our current use cases.

Testing progressed from a small dataset for quick feedback, to larger datasets, and finally in an environment replicating production.

A key challenge was OCR AI extraction errors, which were handled with logic to automatically detect and correct common mistakes. The development phase took approximately 2.5 months.

Step 3: Deployment

The service was deployed using a phased rollout strategy within our Commercial Billing system, starting at 10%, then 50%, 80%, and finally 100% to ensure predictable outcomes and avoid impacting other customers.

Development, test, and production environments were used to properly validate features before deployment. Performance checks ensured the validation algorithm ran efficiently, while security was reinforced with a JWT-based system to prevent unauthorized access.

User readiness was confirmed by sending extensive test results, addressing any issues, and ensuring La Turquoise was prepared for deployment.

The key lesson learned was the importance of careful, phased testing to prevent system-wide issues. Deployment took 2-3 weeks.

Step 4: Live & Ongoing Support

Post-launch, a ticket support system was implemented and made accessible to the entire Cemiar team to ensure timely issue tracking and resolution.

Employees were onboarded through short live demonstrations designed to quickly familiarize them with the system and its usage.

No formal enablement resources, such as guides, FAQs, or workshops, were provided at this stage.

Feedback is collected directly from customers and is used to inform ongoing improvements and enhancements to the solution.

Implementation Challenges and Risks

During implementation, the primary challenges stemmed from evolving client needs and requirements, which required iterative optimizations and continuous adjustments to the solution. Data-related challenges were also encountered, as relevant datasets were limited and difficult to source due to the specificity of the use case (e.g., particular insurers and policy types). Resource constraints did not impact progress; however, some delays occurred due to client approval cycles, as updates required validation through test results before proceeding to subsequent phases.

No significant risks were identified regarding security, privacy, or regulatory compliance, as no confidential client data is stored within the system. Dependency on third-party vendors remains low risk, given the use of reliable cloud infrastructure provided by Microsoft Azure.

To mitigate these challenges and risks, an iterative development approach was adopted, enabling rapid adaptation to changing requirements and continuous performance improvements. Close collaboration with clients was maintained through regular feedback loops and validation checkpoints to streamline approvals. Additionally, targeted live demonstrations were used to support user adoption, while the system architecture leveraged stable and scalable cloud services to ensure reliability and minimize operational risk.



AI Use Case Project: Performance Metrics

	High-level summary	How is it collected	Result
1. Validation Execution Rate	Measures the proportion of PDF files for which the validation process completes successfully without technical errors. This metric reflects system reliability rather than validation accuracy.	The metric was derived from a test set of more than 80 randomly selected PDF files, all of which were processed successfully.	Expected: 90% Actual: 100%
2. Coverage Detection Rate	Measures the percentage of cases where the validation process successfully identifies all coverages within the PDF files. This metric reflects the accuracy of the detection process.	The metric was calculated based on validation runs across more than 80 PDF files.	Expected: 90% Actual: 99%
3. Coverage Hallucination Rate	Measures the proportion of content incorrectly generated or “hallucinated” by the AI during PDF extraction.	The metric was calculated based on validation runs across more than 80 PDF files.	Expected: 2% Actual: 0%
4. Broker Analysis Time	Represents the average time required for a broker to manually analyze and validate a policy renewal, including coverage verification, using the Coverage Gap Finder.	We are now in testing phase to collect feedback from the brokers to analyze the time saved, but we are confident of the results.	Expected: 80% time saved Actual:
5. Reinvestment	We projected to reinvest the time saved in business development with the possibility of dedicating the equivalent of two additional full-time brokers to quotations and new business.	When the project will be fully functional in our departments	Expected: + 250 new clients, approximately 1 000 000\$ premiums
6. Client retention	Improving the client experience through more tailored advice will contribute to higher retention of our existing clients.	It is an expectation difficult to quantify at this stage	Expected: N/A Actual: N/A
7. Employee retention	The quality of the tools will undoubtedly enhance the employee experience through increased engagement and retention, while also improving our ability to attract new talent.	It is an expectation difficult to quantify at this stage	Expected: N/A Actual: N/A



AI Use Case Project: Recommendations for Full Implementation

Overall Summary

To scale from the POC to full implementation, the solution should be extended to support auto policies and additional insurers.

This expansion will allow the service to cover a larger range of policies, significantly reducing manual validation work for clients.

Key success factors include satisfaction with the current POC, secured funding, and alignment from stakeholders to support broader deployment.

The suggested approach follows a phased rollout: first gather test data from new insurers, validate results, and then deploy the solution in production. This strategy ensures controlled, predictable expansion while maintaining service reliability.

Technical & Architecture

No infrastructure changes are required for full-scale deployment, as the architecture was designed to support growth from the start. Integrations, APIs, and data pipelines will remain the same, with only new configurations added for additional insurers.

Scalability measures may include increasing server resources to handle higher demand, while security will be strengthened to ensure the system remains robust and protected as usage expands.

Data & Compliance

To ensure data quality and availability at scale, the system will maintain filtered, high-quality PDFs and continuously validate new insurer data before production deployment.

At the same time, privacy and security will be enforced through strong access controls, JWT authentication, and monitoring for unauthorized access, with security measures scaled to match increased usage.

Process & Operations

Workflow will stay the same but run faster with the existing architecture. Change management will involve stakeholder coordination and readiness checks, while operational risks will be mitigated through phased rollout, testing, and monitoring.

People & Enablement

Users will receive a brief training to learn how to set up new configurations in the UI portal. Adoption will be driven organically, with clients showing colleagues how to use the tool, while Cemiar remains available for support.

Long-term success will be ensured through a ticket-based email system that notifies Cemiar of any bugs for swift resolution.

Thank You!

