Conversational AI for Bank Client Onboarding

MCP for automating client onboarding with AI-driven data extraction and mapping

Students

Arnel Veladzic

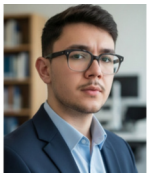

Fabio Gomes Silva

Advisor
Prof. Dr. Mitra Purandare

Subject Area
Artificial Intelligence, Software

Project Partner
Atfinity AG, 8050 Zürich, ZH

Introduction: Atfinity is a Swiss AI-powered no-code platform founded in 2016 that digitalizes and orchestrates client onboarding processes in banking. While these workflows are structured, Relationship Managers (RMs) currently handle the entire case creation manually, which is time-consuming. To improve efficiency, Atfinity plans to extend its platform with AI-assisted capabilities, enabling RMs to create and update onboarding cases automatically and retrieve structured information. This project develops a prototype where an RM can submit client details via a free-text prompt to a Large Language Model (LLM). The LLM extracts and maps the data to Atfinity's internal data model and interacts with the Case Management System (CMS) through the Model Context Protocol (MCP).

Problem: Atfinity's system uses a conditional field architecture, where the visibility of fields depends on prior inputs - for example, fields like Total Wealth or Name appear only after the Account Type is specified. The underlying data model is a JSON structure containing definitions, dependencies and permitted values. Because the platform is highly configurable, banks can define their own fields and rules, resulting in payloads that vary widely and making static mapping infeasible. Additional complexity arises from the lack of established patterns for integrating AI-based extraction with MCP. Limited guidance on architectures, data flow and interaction models necessitates exploratory development.

Result: The project successfully extracted client information from free-text prompts and converted it into structured JSON compatible with Atfinity's CMS. This enables automated onboarding, allowing RMs to submit information conversationally while the system handles extraction, mapping and case population via MCP. Several architectures were explored, of which three were selected and implemented, all sharing the same core architecture but differing in communication flow.

1 – Iterative LLM Mapping (v5)
Adapts to Atfinity's conditional field architecture using iterative LLM calls. It achieved high accuracy in simpler cases (Precision 90.91%, Recall 100%, F1 95.24%) with 23–25s latency. Performance dropped in corporate onboardings (F1 80%, Recall 66.67%).

2 – Iterative Mapping with Preemptive Cache (v6)
Builds a cache before case creation. It achieved 100% Recall but lower Precision (66.67–73.68%) and the slowest speed (32–44s).

3 – Preemptive Cache with Single-Pass Mapping (v7)
Generates a cache once and maps fields without further LLM calls. It achieved the best speed (10–24s), however Recall decreased in complex cases (44.44%).

Summary: Prototype three provides the best speed–quality trade-off, while prototype one is the most reliable in terms of accuracy.

Core Architecture (simplified)
Own presentment

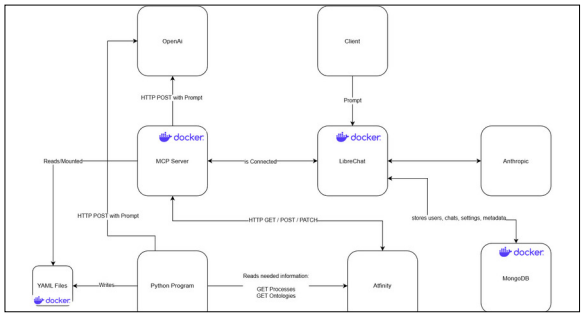

Flowchart (simplified)
Own presentment

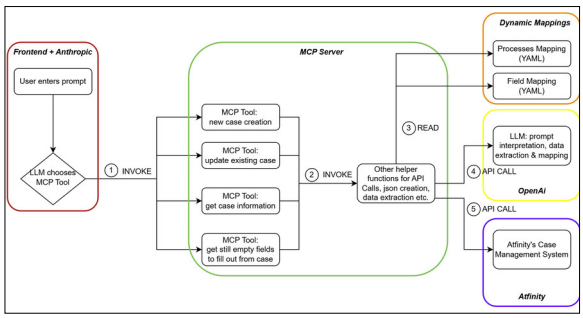

Comparison of Quality Measures
Own presentment

	Architecture v5	Architecture v6	Architecture v7
Test Case 1: Individual Account Onboarding	Precision: 86.67% Recall: 92.86% F1 Score: 89.66% Duration: 25198ms	Precision: 73.68% Recall: 100.00% F1 Score: 84.85% Duration: 36907ms	Precision: 91.67% Recall: 78.57% F1 Score: 84.62% **Duration:** 19813ms
Test Case 2: Simple Onboarding	Precision: 90.91% Recall: 100.00% F1 Score: 95.24% Duration: 23377ms	Precision: 66.67% Recall: 100.00% F1 Score: 80.00% Duration: 34382ms	Precision: 100.00% Recall: 80.00% F1 Score: 88.89% **Duration:** 14991ms
Test Case 3: Corporate Entity	Precision: 100.00% Recall: 66.67% F1 Score: 80.00% Duration: 14606ms	Precision: 90.00% Recall: 100.00% F1 Score: 94.74% Duration: 31937ms	Precision: 100.00% Recall: 44.44% F1 Score: 61.54% **Duration:** 10650ms
Test Case 4: Complex Case with detailed data	Precision: 85.00% Recall: 100.00% F1 Score: 91.89% Duration: 39546ms	Precision: 85.00% Recall: 100.00% F1 Score: 91.89% Duration: 43602ms	Precision: 93.75% Recall: 88.24% F1 Score: 90.91% **Duration:** 23729ms