Business Architecture Copilot: Agentic AI-enabled Business Capability Modelling

Eesha Oaj[1,*], Asif Gill[1,*], Madhushi Bandara[1], Terry Roach[2]

[1] *University of Technology Sydney, 2007, Australia*
[2] *Capsifi Pvt. Ltd, Sydney, 2000, Australia*

Abstract

Business Architecture (BA) plays a significant role in aligning business and IT. Current methods in BA artefact generation and maintenance are largely manual, static, and time-consuming, which limits their effectiveness in dynamic and complex decision-making environments. Large Language Models (LLMs) have showcased their potential in automating the generation of BA artefacts through customised prompting techniques. However, they face a key challenge due to the tendency of LLMs to hallucinate, resulting in outputs that may be inaccurate or inconsistent. Agentic AI offers an alternative approach to LLM-based BA modelling, due to its ability to make autonomous decisions, adapt through learning, and process data in real-time. To investigate their potential, this research project proposes an Agentic AI-enabled business architecture modelling approach (BA Copilot) that leverages AI agents to automate the generation of BA artefacts. Our work employs an agentic system, implemented using the AutoGen 0.2 framework, integrated with an Enterprise Knowledge Graph (EKG) to semantically contextualise organisational knowledge. The proposed system enables AI agents to extract, reason over, and organise business architecture artefacts. The key outcomes of this research project are the Agentic AI-enabled Business Modelling Architecture (ABMA) and an agent-based capability modelling process, designed to enhance decision-making in modern enterprise contexts. By incorporating the knowledge graph and reasoning capabilities through LLM in our design, we aim to facilitate user trust in the AI-generated BA artefacts as well. The current progress of this project involves designing ABMA and implementing a prototype within the scope of business capability modelling (BCM), producing promising preliminary results. We plan to evaluate and extend this work to other BA artefacts and develop a tool that practitioners can use based on our findings.

Keywords

Capability modelling, Enterprise knowledge graphs, Agentic AI, Business Architecture, Decision Making

1. Introduction

Organisations are increasingly interested in the digitalisation of their organisational processes and workflow [1]. This requires organisations to fundamentally change their business models and operating environment by adopting contemporary data and digital technologies, such as AI/ML [2]. Despite the potential influence of data and digital technologies on BA methods, a key challenge for organisations is identifying the capabilities, tasks and decisions that can be augmented or supported via AI. This draws our attention to the BA methods within the enterprise architecture (EA) discipline [3] and how they can support organisations' digital technology-enabled transformations. BA, according to TOGAF [4], is "a description of the structure and interaction between the business strategy, organisation, functions, business processes, and information needs". BA supports various

ER2025: Companion Proceedings of the 44th International Conference on Conceptual Modeling: Industrial Track, ER Forum, 8th SCME, Doctoral Consortium, Tutorials, Project Exhibitions, Posters and Demos, October 20-23, 2025, Poitiers, France

*Corresponding author
eesha.oaj@student.uts.edu.au (E. Oaj); asif.gill@uts.edu.au (A. Gill); madhushi.bandara@uts.edu.au (M. Bandara); troach@capsifi.com (T. Roach)
0000-0003-2461-8369 (E. Oaj); 0000-0001-6239-6280 (A. Gill); 0000-0001-6543-3841 (M. Bandara)

business stakeholders in decision-making by representing and analysing data around entities that govern an enterprise, such as business strategies, goals, capabilities, pain points, business processes, applications, and rules and regulations [4]. Key BA artefacts such as capability maps, value stream maps, and application portfolios provide crucial information for IT investments, performance analysis, and resource utilisation [5]. However, current BA artefacts are static; they become outdated quickly and are unable to keep up with the evolving business decision-making [6]. Static methods refer to traditional, predominantly manual approaches to BA artefact generation, in which the resulting outputs are fixed representations that rapidly become outdated and lack adaptability to dynamic and evolving business environments [7].

Agentic AI, often regarded as the fourth wave of artificial intelligence (AI), evolves from traditional AI agents [8] by integrating deep learning techniques [9], including LLMs [10]. This fusion enables autonomous decision-making, contextual awareness, and the ability to execute complex, multi-step tasks with minimal human input. Agentic AI has the potential not only to replicate but also to exceed human capabilities in managing complex systems [11]. Such systems are capable of autonomously performing reasoning, planning, and executing actions in real time. There is growing research interest in generating knowledge relevant to BA through LLMs by using various prompt-based techniques [12]. However, exploring agentic AI systems in the context of BA remains underexplored in the current literature and practice [13].

This research is motivated by limitations of current BA methods and focuses on business capability modelling (BCM), which is one of the critical parts of BA. BCM is important to analyse and improve the strategic business and IT alignment [14]. BCM involves identifying and visualising the core and enabling capabilities of an organisation or business area, such as customer relationship management capability, claims management, human resource management capability, etc [15, 16]. In particular, this paper aims to address the limitations of existing manual or static BCM by using an agentic AI-enabled approach [17]. This is an ongoing research project aimed at developing a BA Copilot[✗], designed in partnership with Capsifi[*] Pvt. Ltd, a leading platform in BA modelling and operating model transformation.

Preliminary work of this project involved conducting a systematic literature survey on BA methods [7] and developing and evaluating reasoning-based prompting techniques for generating BA artefacts [18]. Through that, we found that by incorporating reasoning-based prompting techniques, particularly Tree-of-Thought (ToT) [19] results in improved BA artefact generation. We are currently extending this work by introducing an agent-based capability modelling process.

It is noteworthy that this project is a work-in-progress where our proposed solution is validated and evaluated iteratively. We present our preliminary findings here, which have been validated by the authors, experts in BA modelling. Our results and insights are intended to inform both research and practice in enterprise modelling.

The key contributions of this research project are as follows:

1. The design of the Agentic AI-enabled Business Modelling Architecture (ABMA), which integrates a multi-agent system with enterprise knowledge graphs for structured BA artefacts.
2. The ABMA process that automates the extraction, reasoning, and organisation of BA artefacts.
3. Proof-of-concept is a tool demonstration of Agentic AI-enabled Business Capability Modelling Architecture (ABCMA), and the modelling process proposed.

The remainder of the paper is organised as follows. Section 2 provides a research method. Section 3 discusses the concept of approach. Section 4 demonstrates a scenario illustration. Section 5 concludes the paper with challenges and future research directions.

[*] https://old.capsifi.com/
[✗] *https://www.digisaslab.org/business-architecture-copilot/*

2. Research Method

This research project adopts the Design Science Research (DSR) method, as BA modelling is naturally complex and organised in layers. DSR works well for this kind of problem, as it helps break down complex issues into smaller, more manageable parts clearly and logically. This project follows the structured DSR hierarchy proposed by Tuunanen et al. [20] as illustrated in Figure 1.

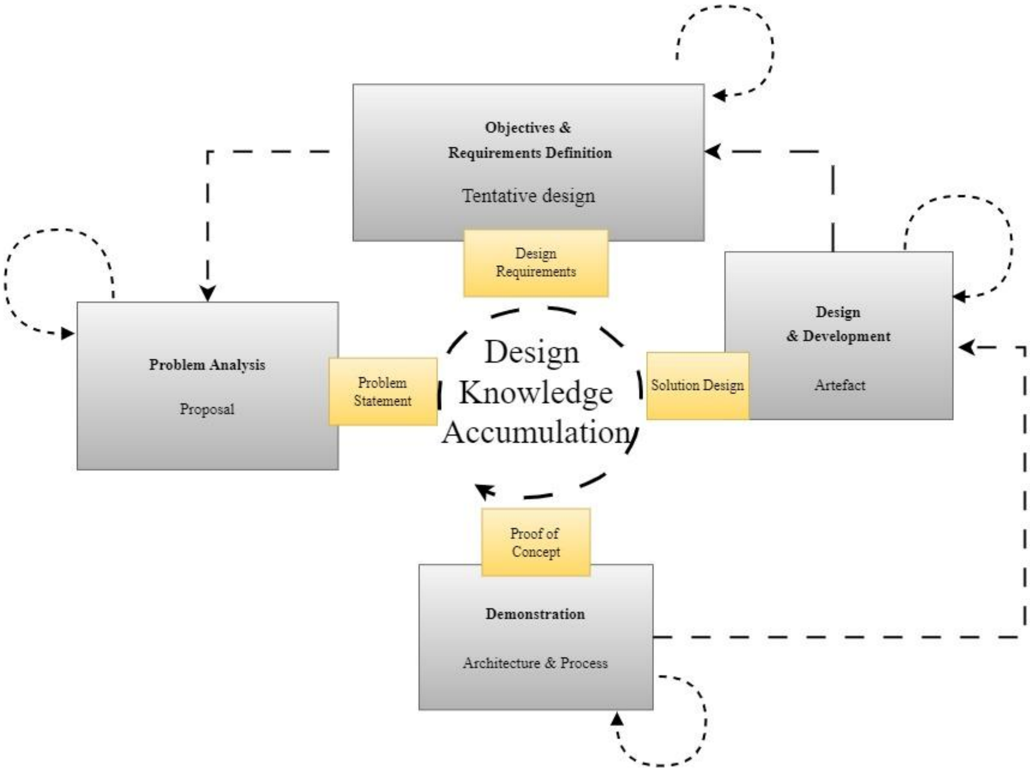

Figure 1. The DSR Method adopted from Tuunanen et al. [20]

The adopted DSR has four stages (Figure 1). Firstly, we identified the research problem, which is the static and manual nature of traditional BA methods, producing static BA artefacts that become outdated quickly and are unable to keep pace with evolving business decision-making. Secondly, our goal was to explore how agentic AI can be leveraged to automate BA artefacts. We applied creative thinking to produce a tentative initial design of ABMA. Comprehensive research was conducted to develop the ABMA design. Thirdly, the design and development are performed iteratively to evolve ABMA. Finally, the results of the scenario illustration will be presented and discussed to conclude the project.

3. Agentic AI-enabled Business Modelling Architecture (ABMA)

3.1. Overview

The ABMA (see Figure 2) presents a modular and AI-driven system designed to automate the generation of BA artefacts through the orchestration of autonomous agents. This system architecture is built upon the principles of agentic AI, natural language interaction, and enterprise knowledge contextualisation to overcome the static and labor-intensive nature of traditional BA methods. The architecture comprises five core layers: user interaction, context extraction, orchestration, processing, and knowledge layer. The underlying cognitive framework of ABMA mirrors the way humans' reason through complex tasks. How each layer interacts and operates is detailed through a demonstration in Section 5.

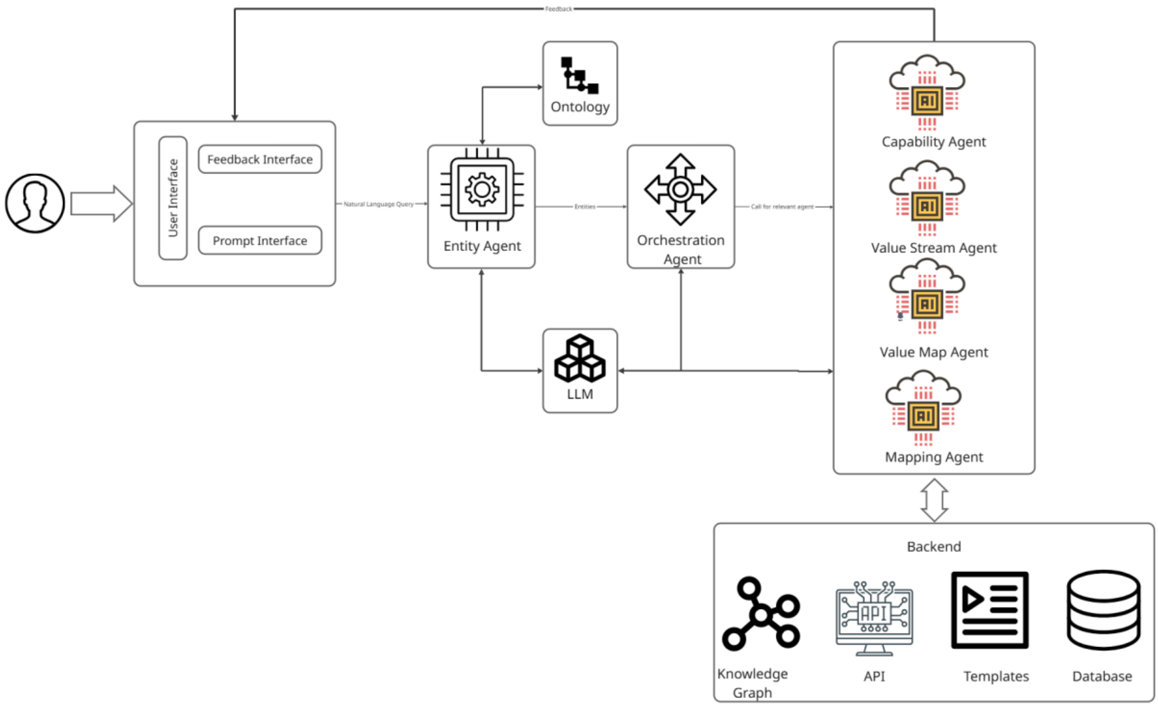

Figure 2. Agentic-AI-enabled BA modelling architecture (Content generation layer)

3.2. ABMA Component Interaction

When architecture is implemented in the real world, the BA artefact generation process is initiated by the user upon submitting a natural language query through the prompt interface within the user interface. There is also a feedback interface for input refinement based on system output. For example, a query can request the generation of capability maps for a specific industry, region, or business function, as detailed in Section 5.

Secondly, the natural language query is passed on to the entity agent, which serves as the first point of analysis. It parses the query to extract contextual information such as industry type, artefact category (e.g. operational capability, strategic capability), and geographic or regulatory scope. The entity agent may delegate tasks to LLM to handle ambiguous or unstructured input. The entity agent is also linked to a reference ontology, providing semantic grounding for the identified entities. This enables the agent to resolve any disambiguation and ensure that the entities identified are aligned with the operational domain. A specific industry or organisation can have its ontology defined and linked for this purpose.

Thirdly, the parsed entities are handed over to the orchestration agent, which is responsible for routing the query to a particular agent based on the identified entities. The agent responsible for the task will generate the artefact in response to the user's query, interacting with the LLM. To ensure semantic accuracy and organisational alignment, the agent will validate the artefacts generated by the LLM against both the domain ontology and the enterprise knowledge graph (EKG). The input knowledge for generating capability lists is derived from user queries, enriched by reference ontologies, and validated against EKG to ensure semantic accuracy, contextual alignment, and organisational relevance. This hybrid approach combines the generative strength of the LLM with the semantic control of structured enterprise knowledge, resulting in accurate and context-aware BA modelling, which leads to increased trust in the generated model. These interactions are illustrated in Figure 2. Finally, once the agent completes its reasoning, the structured output is returned to the user for feedback, refinement, or re-query.

4. Scenario Illustration: Agentic AI-enabled Business Capability Modelling Process:

To demonstrate and validate the practical feasibility of the proposed ABMA, we have developed a proof-of-concept implementation of an agentic AI-based business capability modelling process using AutoGen 0.2, an open-source multi-agent orchestration framework developed by Microsoft. AutoGen 0.2 facilitates the coordination of autonomous agents through structured dialogues and dynamic task delegation, making it well-suited for realising the modularity and autonomy envisioned in the ABMA. For demonstration purposes, we have implemented this through the OpenAI API GPT-4o. Figure 3 illustrates our design in AutoGen 0.2.

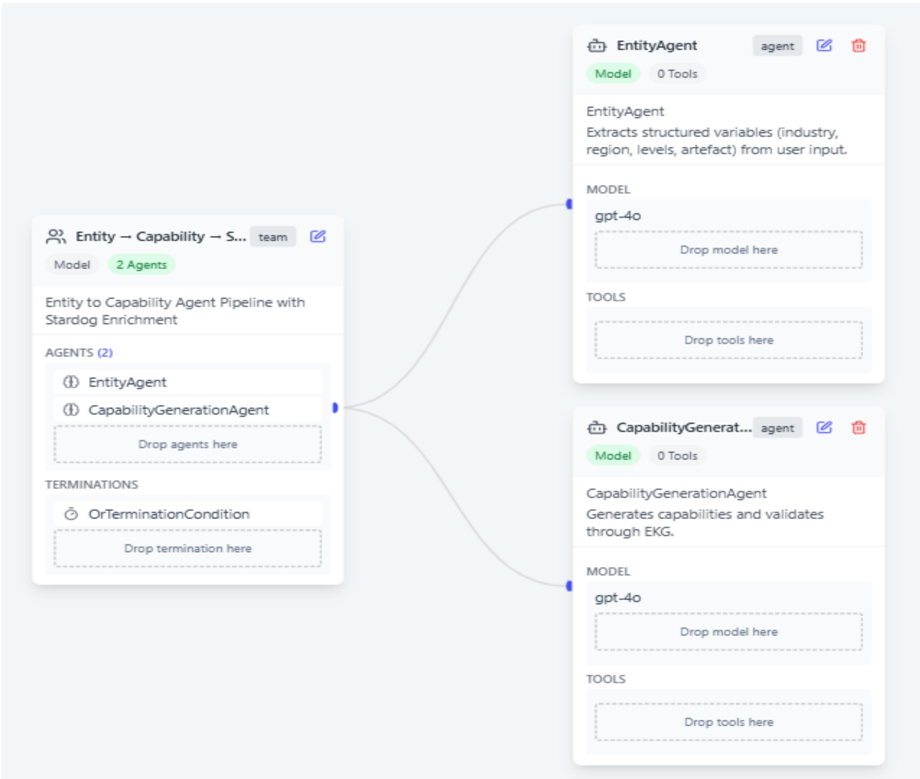

Figure 3. AutoGen 0.2 implementation

To understand the value of ontology and EKG in the proposed architecture, we implement this scenario over two configurations.

- Configuration 1: entity agent and capability agent interact with LLM only
- Configuration 2: entity agent and capability agent access ontology and EKG

In Configuration 2, the Capsicum BA ontology and EKG, provided by the industry partner, are used in our demonstration [21].

4.1. Configuration 1: ToT-Prompted Agent Reasoning

In the initial phase of experimentation (as shown in Figure 4), we created an entity agent and a capability agent. In this configuration, we employed the reasoning-based ToT prompting technique. ToT enhances the reasoning capabilities of language models by allowing them to explore multiple solution paths in parallel and evaluate intermediate thoughts before arriving at a final response. The entity agent was provided with simple instructions to parse user queries into structured entities: BA artefact, industry, region and levels. The capability agent received structured entities as input and was tasked with selecting capability definitions and composing a capability map. At this stage, semantic grounding was limited to rule-based inference and a manually curated ToT prompt technique, without leveraging external knowledge graphs.

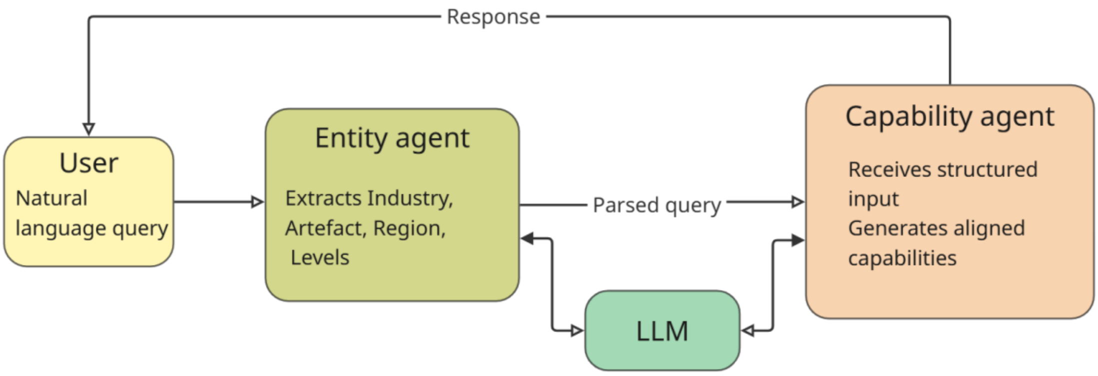

Figure 4. Config 1 prompt-based BCM

The extracted portion (claim management capability) of the response from the complete capability map is shown in Table 1

Query: Generate a 5-level capability map for the insurance industry in Japan.

Table 1. Response to the prompt-based BCM

Level	Capability	Parent	Description
1	Claims Management		Ability to identify, track, and respond to a demand or request to indemnify customers under the terms of an agreement
2	Claims Submission	Claims Management	Facilitating the process of customers submitting claims with the required documentation and information
3	Claims Assistance	Claims Submission	Providing support to customers during the claim submission process to ensure completeness and compliance
3	Claims Intake	Claims Submission	Collecting and logging claim information for initial review and assessment
2	Claims Settlement	Claims Management	Managing the process of evaluating, negotiating, and resolving claims in accordance with policies
3	Claims Evaluation	Claims Settlement	Assessing the validity and coverage of submitted claims against policy terms
3	Claims Payment	Claims Settlement	Coordinate the disbursement of approved claim payments to customers

4.2. Configuration 2: EKG Integration via Stardog

In the second configuration (see Figure 5), semantic validity was introduced through the integration of an EKG, developed by our industry partner, Capsifi. This graph is hosted on Stardog, a graph database platform that supports RDF/SPARQL-based queries and semantic reasoning.

The entity agent was extended to query the Stardog-based EKG, enabling more accurate parsing of user queries by grounding extracted entities in organisational semantics (e.g., industry taxonomies, regional classifications, artefacts). The capability agent used the same EKG to validate whether the requested capability concepts existed within the enterprise ontology. We used the exact basic instructions for the entity agent and the prompt technique for the capability agent, as we have used in Configuration 1. The difference is the addition of an EKG.

In this configuration, we performed a matching test to verify whether the capabilities generated match those in EKG. If a match was found, the agent retrieved the corresponding subgraph and used it to enhance or tailor the generated capability map. If the ontology did not contain a matching concept, the system generated a response indicating the absence of semantic alignment and optionally suggested alternative capabilities or taxonomic nodes.

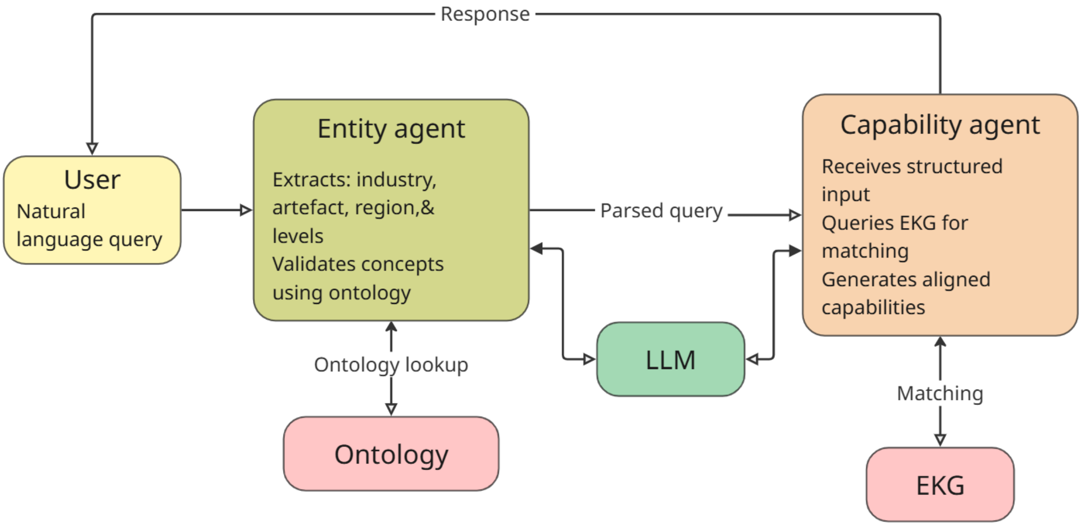

Figure 5. Config 2 EKG-based BCM

The extracted portion (claim management capability) of the response from the complete capability map is shown in Table 2
Query: Generate a 5-level capability map for the insurance industry in Japan

Table 2. Response of EKG-based BCM

Level	Capability	Parent	Description	Match with EKG
1	Claims Management		"Handling the entire lifecycle of insurance claims."	Yes
2	Claims Filing	Claims Management	"Accepting and documenting initial claims from policyholders."	No
3	Claim Information Collection	Claims Filing	"Gathering necessary information and documentation for the claim."	No
3	Claim Form Submission	Claims Filing	"Submitting a structured claim form for processing."	No
3	Initial Claim Verification	Claims Filing	"Conducting preliminary checks to ensure claim validity."	No
2	Claim Assessment	Claims Management	"Evaluating claims to determine payout eligibility and amounts."	No
3	Damage Evaluation	Claim Assessment	"Assessing the extent of the damage or loss claimed."	No
3	Benefit Calculation	Claim Assessment	"Determining benefit payout amounts based on policy terms."	No
3	Claim Investigation	Claim Assessment	"Investigating claims for fraud or misrepresentation."	No
2	Claims Settlement	Claims Management	"Managing the process of finalising claims and disbursing payments."	No
3	Negotiation with Claimant	Claims Settlement	"Engaging with claimants to agree on settlement terms."	No
3	Payment Processing	Claims Settlement	"Processing payments to settle claims."	No
3	Settlement Documentation	Claims Settlement	"Documenting agreements and settlements with claimants."	No

Summary of findings:
Key takeaways from the scenario illustration are:

- Configuration 1 performed better at creative, unconstrained generation of the capability map, but was semantically unverified.
- Configuration 2 offered semantic alignment and potential ontology alignment of capabilities, but it is limited to level 1.

5. Challenges and Future Work

This research project proposes the development of a BA copilot that can automate BA modelling through agentic AI. We present an illustrative scenario of the proposed method, showcasing the potential of Agentic AI in automating business capability modelling through the proposed ABMA. We observe that when EKG is not incorporated into the agent architecture (i.e., Configuration 1 using ToT-prompted agents), entity extraction can be completed using the user query; however, it lacks semantic precision due to the absence of external grounding. Configuration 2 addressed this by integrating an EKG, grounding the entities using organisational ontologies and improving the validity of generated output. These observations suggest that ABMA's modular architecture and agent orchestration proved effective for scalable and context-aware capability mapping. The use of LLMs combined with structured knowledge significantly improved the accuracy and adaptability of generated BCMs.

The next step in our project is to evaluate and expand the illustrative scenario to different industries and other hierarchical levels of capability modelling (Levels 2-3), followed by a systematic evaluation of agent performance. We plan to extend our implementation to additional BA artefact agents, aiming to develop a more intelligent and autonomous BA method. The end goal of this project is to provide architecture that can be adapted by BA experts who are interested in automating their modelling process, while incorporating context-aware reasoning that enables trust in the system.

Acknowledgements

We sincerely thank our industry partner, Capsifi, and the University of Technology Sydney for providing research funding and scholarship to support this research.

References

[1] K. Shoilekova, "Development of a business architecture model and its role in constructing an organisational architecture," in *Software Engineering Application in Informatics: Proceedings of 5th Computational Methods in Systems and Software 2021, Vol. 1*, 2021: Springer, pp. 216-223.

[2] F. W. Mateo and A. Redchuk, "The emergence of new business and operating models under the industrial digital paradigm. Industrial Internet of Things, Platforms, and Artificial Intelligence," *Machine Learning. Journal of Mechanics Engineering and Automation (JMEA)*, vol. 11, no. 2, pp. 54-60, 2021.

[3] A. Nayeem, R. Dilnutt, and S. Kurnia, "Enterprise Architecture Practice and Challenges in Achieving Sustainable Digital Transformation in Developing Countries," 2023.

[4] N. Javanbakht and E. S. Chia, "A Generic Business Architecture Reference Model Using Systems Theory Perspective in Enterprise Architecture," in *2023 IEEE 25th Conference on Business Informatics (CBI)*, 2023: IEEE, pp. 1-10, doi: 10.1109/cbi58679.2023.10187480.

[5] J. L. G. Caceres and A. A. R. Gómez, "Business architecture A differentiating element in the growth of organizations.pdf," 2019.

[6] E. Oaj, A. Gill, and M. Bandara, "Business architecture methods: An initial review and mapping study," 2024.

[7] E. Oaj, A. Gill, M. Bandara, and T. Roach. (2025). Business Architecture Methods: A Systematic Literature Review.

[8] R. Vatankhah Barenji and S. Khoshgoftar, "Agentic AI for autonomous anomaly management in complex systems," *arXiv e-prints,* p. arXiv: 2507.15676, 2025.
[9] A. Mathew, P. Amudha, and S. Sivakumari, "Deep learning techniques: an overview," in *International conference on advanced machine learning technologies and applications*, 2020: Springer, pp. 599-608.
[10] W. X. Zhao *et al.*, "A survey of large language models," *arXiv preprint arXiv:2303.18223,* vol. 1, no. 2, 2023.
[11] A. R. KP and S. Gowrishankar, "Agentic AI for Pathogen-Based Plant Disease Detection," in *2025 6th International Conference on Artificial Intelligence, Robotics and Control (AIRC)*, 2025: IEEE, pp. 323-328.
[12] M. Thanasi-Boçe and J. Hoxha, "From ideas to ventures: Building entrepreneurship knowledge with LLM, prompt engineering, and conversational agents," *Education and Information Technologies,* vol. 29, no. 18, pp. 24309-24365, 2024.
[13] J. Jung and P. Wienke, "Generating Business Capability Maps using GenAI: A Case Study," 2024.
[14] S. Kotusev and A. Alwadain, "Modeling business capabilities in enterprise architecture practice: the case of business capability models," *Information Systems Management,* vol. 41, no. 2, pp. 201-223, 2024.
[15] A. Q. Gill, *Adaptive Enterprise Architecture as Information.* 2022.
[16] G. Koutsopoulos, M. Henkel, and J. Stirna, "An analysis of capability meta-models for expressing dynamic business transformation," *Software and Systems Modeling,* vol. 20, no. 1, pp. 147-174, 2021.
[17] K. Huang, *Agentic AI.* Springer, 2025.
[18] E. Oaj, A. Gill, M. Bandara, R. Maymendpour, and T. Roach. (2025). Prompt Engineering Techniques for GenAI-enabled Business Architecture.
[19] S. Yao, D. Yu, J. Zhao, I. Shafran, T. Griffiths, Y. Cao, and K. Narasimhan, "Tree of thoughts: Deliberate problem solving with large language models," *Advances in neural information processing systems,* vol. 36, pp. 11809-11822, 2023.
[20] T. Tuunanen, R. Winter, and J. vom Brocke, "Dealing with complexity in design science research: A methodology using design echelons," *MIS Quarterly,* vol. 48, no. 2, 2024.
[21] T. Roach, "CAPSICUM - a semantic framework for strategically aligned business architecture," 2011.