



Curtin Singapore

MAY 2026

RESEARCH PULSE



Contents

1 Welcome to Research Pulse

2 Progress Update on the First Call for the Internal Research Grant Scheme

a. A Comprehensive AI-driven Academic Integrity System to Safeguard Authentic Student Learning

b. Student Learning Support Platform - EduAI

c. Attitudes and Perceptions about the Use of Generative Artificial Intelligence (GenAI) in Higher Education in Singapore: A Mixed-Methods Study

d. Generative AI-Based Stock Return Prediction System

e. Strategic Generative AI Automation and Integration to Improve University Research Workflows

3 Student Article - The Application of AI in University Systems: Personalized Empowerment and Institutional Competitiveness

4 Research Funding

5 Publication

1. Welcome to Research Pulse

In this issue, we will share updates on the projects from the first call of the Internal Research Grant Scheme (IRGS) at Curtin Singapore.

The scheme is designed to strategically nurture, inspire, and strengthen a vibrant research culture across the University. It provides targeted seed-level funding that enables researchers to explore bold ideas, initiate pilot studies or proof-of-concept projects, and generate the foundational evidence required to secure larger external grants.

Aligned with Curtin Singapore's research clusters, the scheme aims to enhance institutional research capacity by encouraging meaningful industry collaboration and fostering interdisciplinary innovation. It serves as a catalyst for creativity, excellence, and product development, empowering researchers to experiment confidently, produce impactful publications, and elevate Curtin Singapore's overall research performance.

Thank you for your continued support to Research Pulse.

Best wishes,

Curtin Singapore Research

2. Progress Update on the First Call for the Internal Research Grant Scheme

a) A Comprehensive AI-driven Academic Integrity System to Safeguard Authentic Student Learning

Principal Investigator: *Professor Arun Patil, Curtin Singapore*

Project Timeline: *21 November 2025 – 21 June 2026*

Project Aim

The objective of this research project is to propose a new standard academic integrity system that integrates authorship analysis, behavioural monitoring, provenance tracking, and active AI-driven checks. This system moves beyond “plagiarism detection” to become a proactive safeguard of authentic student learning. This proposed system will adopt the important principle and philosophy of “prevention is better than cure” and provide opportunities to students to minimise the attempts of academic misconduct.

Furthermore, the system will save the time of the “confirmation check” process by minimising the requirements of confirmation checks and validity of the misconduct attempts. It will also save students’ time spent on confirmation checks as well as stress level going through the process.

Progress to Date

So far, significant progress has been achieved, and three milestones (out of four) have been delivered as planned.

A comprehensive Proof of Concept document and SLSP Software platform is developed. In the second phase, more features towards the behaviour of the students were developed and was demonstrated in a video. In the third and recent phase, more improvements towards the AI were developed to make the system more autonomous including some smart checks.

In addition, a very important and user-friendly Leadership/staff dashboard is also developed.

Key Outputs and Deliverables

- A comprehensive Proof of Concept
- The development of SLSP Software platform
- Student behaviour model
- AI cheating model
- Active checks feature
- Smart confirmation checks
- Leadership/staff dashboards

Outcomes and Impact

Project outcomes clearly demonstrate the development of a robust academic integrity tool for Curtin Singapore aiming at minimising academic misconduct attempts and approaches at the university.

There are improvements in reporting efficiency, approval turnaround time, and staff engagement with digital tools. The developed artefacts provide a reusable digital and AI reference architecture for future Curtin Singapore initiatives and generate applied digital transformation research outputs suitable for publication, professional development, and staff capability building.



Next Steps

The final and fourth important milestone of packaging, testing with bug corrections will be delivered in due course. Furthermore, the finalisation of project documentation will be carried out. In the project dissemination, key findings will be submitted to relevant conferences as well as to various internal platforms and avenues. The project remains on track for completion within the approved timeline and budget.

b) Student Learning Support Platform - EduAI

Principal Investigator: *Dr Jaideep Chandran, Curtin Singapore*

Project Timeline: *24 November 2025 – 1 August 2026*

Scope

High-level architecture for Question Generation, Transcript, Summarization, and AI Assistant workflows, including ingestion, retrieval, generation, persistence, and shared service responsibilities.

Executive Overview

This system implements a modular retrieval-augmented generation (RAG) architecture to support three principal capabilities: educator-facing question generation, transcript summarization, and student-facing contextual question answering. The platform follows a common pipeline pattern of **ingest -> retrieve -> generate -> persist**, while allowing each feature to use retrieval and storage components differently based on product requirements.

At the platform level, the solution consists of a browser-based client, an authenticated API gateway, a Python RAG orchestration layer, local persistence in PostgreSQL with vector support, external semantic retrieval via Ragie for course-scale search, and OpenAI services for generation and embeddings.



Functional Pipelines

Question Generator

The Question Generator converts educator-provided materials into structured assessment items. An educator selects a source file from previously ingested materials or uploads a new asset. The request enters the platform through the Express API gateway, which authenticates the session and forwards the workload to the FastAPI RAG service.

For newly uploaded documents, the RAG service extracts text, chunks content, generates embeddings, and persists the output in PostgreSQL. In PDF upload scenarios, Ragie may also be used as an auxiliary indexing layer. At generation time, retrieval is scoped to the relevant user and filenamecontext from PostgreSQL rather than a course-wide search index. Retrieved chunks are incorporated into prompts, and OpenAI generates draft questions. Educators may review the generated output and optionally persist approved items to the question bank.

Transcript Summary

The Transcript Summary workflow ingests transcript-like source material, stores chunked content in PostgreSQL, retrieves those stored chunks, and produces a structured summary through the LLM layer. In the current implementation, this feature does not primarily use Ragie for answering. Instead, it relies on the local database-backed chunk store for retrieval and then calls OpenAI to produce the final summary. Persistence is generally limited to the ingested chunks rather than the generated summary itself.

AI Assistant

The AI Assistant supports student-facing question answering at the course level. A student submits a natural-language question through the frontend, and the request passes through Express to the FastAPI RAG service. Unlike the educator file-scoped retrieval path, this workflow uses Ragie as the principal retrieval layer.

Ragie performs semantic search over previously indexed course materials, and the retrieved context is passed to OpenAI for answer generation. PostgreSQL is not the primary retrieval surface for this chat pipeline in the current implementation, and no persistent server-side conversation store is assumed in the supplied architecture summary.

Comparative Pipeline View

Table 1: The three features share a common control plane but diverge in retrieval and persistence behavior.

Pipeline Stage	Question Generator	Transcript Summary	AI Assistant
Ingest	Upload, chunk, embed, store in DB; optional Ragie indexing for PDF path	Upload, chunk, embed, store in DB	Typically consumes pre-indexed course content
Retrieve	PostgreSQL chunks scoped by user and file	PostgreSQL chunks	Ragie course-level
Generate	OpenAI creates questions	OpenAI creates structured summary	OpenAI generates course-context answer
Persist	Optional question bank write-back	Usually none beyond chunks	UI-only interaction in current implementation

Service Inventory

Shared Platform Services

- **React / Vite frontend:** browser-based client for educator and student workflows, responsible for rendering interfaces and invoking authenticated API requests.
- **Express API Gateway:** single front door for session handling, JWT or cookie-based authentication, route management, and file upload handling.
- **FastAPI RAG service:** Python service that performs ingestion, chunking, embedding orchestration, retrieval selection, prompt assembly, and database persistence.

Persistence and Retrieval Services

- **PostgreSQL + pgvector:** primary storage for document chunks, embeddings, and saved question bank entries. Supports local vector retrieval and structured persistence.
- **Ragie:** external semantic retrieval service used primarily for course-indexed search and, in some upload scenarios, PDF-oriented indexing.

Generation Services

OpenAI: provides question generation, transcript summarization, contextual answer generation, and embeddings where the ingestion pipeline uses OpenAI-backed embedding workflows.

c) Attitudes and Perceptions about the Use of Generative Artificial Intelligence (GenAI) in Higher Education in Singapore: A Mixed-Methods Study

Research Team: *Dr Carolyn Koh, Dr Nik Chong, Mr Jeffrey Woo, and Mr Jack Ng, Curtin Singapore*

Project Timeline: *15 December 2025 – 30 November 2026*

This research examines how students at Curtin Singapore perceive and use Generative Artificial Intelligence (GenAI) in their academic work, with a focus on learning practices, academic integrity, and responsible use. Using a sequential explanatory mixed-methods design, the study combines a quantitative survey with qualitative focus group discussions to provide both breadth and depth of insight.

The project is on schedule, having progressed from the planning stage to active implementation. Ethics approval was successfully obtained from the Curtin University Human Research Ethics Committee, enabling the study to proceed.

Phase 1

The pilot study has been completed. The second stage of actual data collection at both the Singapore and Bentley campuses has been initiated through a survey deployed via Qualtrics, capturing students' attitudes, usage patterns, perceived benefits, and concerns related to GenAI. Phase 1 is currently work in progress and should be completed in May / Jun.

Initial observations from the pilot study indicate that while students recognise the efficiency and support provided by GenAI tools, there remain concerns around over-reliance, academic integrity, and the boundaries of responsible use. These early patterns will inform the design of the subsequent qualitative phase. Once the survey data has been collected. The data will be analysed, the findings will inform the next phase of research.



Phase 2

This phase of the research marks the commencement of the qualitative data collection process. The method for this phase will involve focus group discussions involving students and educators. Before the focus group discussions commence, our industry partner, Adobe, will conduct a workshop for the focus group participants demonstrating their latest AI software. Findings from this phase will be coded and analysed, then synthesised with the quantitative findings.

Findings from phases one and two will be integrated to develop a comprehensive understanding of how GenAI is shaping learning behaviours and decision-making in higher education.

Publications

Concurrent with our data collection, the team has been developing the literature review and methodology sections of the manuscript in preparation for submission to a Q1 journal. We have also sought guidance from Prof Alex regarding suitable journal outlets for submission. This places us ahead of our proposed timeline, which scheduled the literature review and manuscript development for journal publication and conference presentation between 30 October and 30 November 2026.

Further outcomes

Further anticipated outcomes include a second peer-reviewed journal publication, a conference presentation, and an industry report that offers practical recommendations for educators and institutions. The project also aims to support Curtin Singapore's efforts to advance responsible AI adoption in education and to strengthen student AI literacy within an increasingly complex and rapidly evolving digital landscape.



d) Generative AI-Based Stock Return Prediction System

Principal Investigator: *Dr Ronnie Soh, Curtin Singapore*

Co-Investigators: *Mr Leo Kee Chye and Mr Raymond Qiu, Curtin Singapore*

Project Timeline: *5 January 2025 – 31 May 2026*

Project Aim

This project aims to design and develop a Generative AI-based stock return prediction system that integrates financial market data, macroeconomic indicators, technical signals, and large language model extracted sentiment or narrative signals. The project examines whether GenAI and agent-based analytical tools can enhance the interpretation of financial information and support more informed stock return prediction.

The objective is to build a practical and scalable prototype that combines traditional quantitative finance methods with emerging GenAI capabilities. The system is intended to support future research in AI-driven investment analysis, financial sentiment modelling, return forecasting, and finance-related decision-support dashboards.

Progress to Date

Good progress has been made in clarifying the overall project direction and developing an initial working prototype. The team has identified a roadmap to apply LLMs to macroeconomic and stock return prediction using three main categories of inputs: LLM-extracted sentiment or narrative signals, historical price data, and macroeconomic data.

A preliminary dashboard has been completed. The dashboard currently considers several agent-based inputs, including LLM-extracted sentiment signals, simple technical tools, and investment style perspectives commonly associated with fund managers, such as value versus growth orientations. This early-stage output demonstrates the feasibility of combining GenAI-derived qualitative signals with conventional quantitative market indicators.



Key Outputs and Deliverables

The key outputs and deliverables completed or substantially progressed to date include:

- Identification of a project roadmap combining macroeconomic, price-based technical indicators, fund manager investment style perspectives, and sentiment-based data
- Preliminary testing of selected models and APIs from open-source repositories
- Early exploration of Agentic AI to support automation of the research workflow
- Initial conceptual framework for a multimodal stock return prediction pipeline

Outcomes and Impact

The expected outcomes and impact are:

- Demonstrates how financial market data, macroeconomic variables, technical indicators, and LLM-based sentiment signals can be combined within one analytical system.
- Supports finance teaching and research by showing how AI tools can be used to analyse financial markets in a structured and transparent manner.
- Strengthens applied research capability in AI-driven finance, investment analytics, and decision-support systems.
- Offers potential publication opportunities, including a conference paper, applied finance paper, or AI-in-finance teaching case.

Next Steps

The next phase will focus on transforming the preliminary dashboard into a more robust research prototype. A key next step is to strengthen the machine learning component of the project. This will involve preparing suitable training, validation, and testing datasets so that the models can learn from historical market patterns, macroeconomic conditions, technical indicators, and LLM-extracted sentiment signals.

Further development will include improving the sentiment extraction process, testing baseline models, evaluating predictive performance, and documenting the results. The team will also continue exploring Agentic AI for workflow automation, particularly for data retrieval, feature engineering, model execution, and report generation.

The project is expected to be completed around August to September 2026, subject to the availability of suitable data sources, model testing progress, and dashboard refinement.

e) Strategic Generative AI Automation and Integration to Improve University Research Workflows

Principal Investigator: *Dr Kiah Mok Goh, Curtin Singapore*

Co-Investigator: *Ms Jacqueline Ang, Independent Trainer*

Project Timeline: *1 December 2025 – 30 June 2026*

Project Aim

This project aims to design and prototype a compact digital transformation (DX) framework that applies Generative AI (GenAI), business analytics, and automation technologies to enhance university research and academic workflows. The work aligns with Curtin's strategic priorities in digital innovation, emerging technologies, and contemporary pedagogy, while delivering applied research outcomes and institutional capability development.

Progress to Date

Significant progress has been achieved across all planned workstreams. A fully functional Business Intelligence (BI) dashboard has been developed, incorporating curated datasets, interactive key performance indicators (KPIs), drill-down analytics, and row-level security. Customised GenAI tools have also been implemented to support course-level assistance for both students and lecturers.

In addition, a GenAI-enabled training agent and e-learning modules have been configured with structured learning content, knowledge grounding, quiz-based assessments, and interactive simulations to enhance compliance awareness and engagement. An end-to-end travel approval workflow has been successfully automated using Microsoft Forms, SharePoint, Outlook, and Microsoft Teams, enabling real-time notifications and a fully auditable approval trail.

Key Outputs and Deliverables

- Governed BI workspace with analytics dashboards and reporting capabilities
- GenAI agent with training modules, simulations, and assessments
- Automated travel approval workflow with SharePoint-backed audit trail
- Draft conference paper documenting system architecture, innovation, and outcomes

Outcomes and Impact

Early outcomes demonstrate measurable improvements in reporting efficiency, approval turnaround time, and staff engagement with digital tools. The developed artefacts provide a reusable digital and AI reference architecture for future Curtin Singapore initiatives and generate applied digital transformation research outputs suitable for publication, professional development, and staff capability building.

Next Steps

Remaining activities include minor refinements to analytics dashboards and automated workflows, finalisation of project documentation, submission of the conference paper, and dissemination of outcomes through internal knowledge-sharing and capability-building sessions. The project remains on track for completion within the approved timeline and budget.



3. Student Article - The Application of AI in University Systems: Personalized Empowerment and Institutional Competitiveness

Mr Liu Jinwen, COMP2006 student, Curtin Singapore

The tertiary education system, like Curtin University, characterized by independent research and open access to resources, provides a high-quality platform for students with strong self-directed learning habit expectations. For international students with other language backgrounds, different learning methods and preferences, personalized support and guidance could shorten the adaptation process significantly.

The cultivation of students' sustainable learning abilities and independent research literacy, which are the most important criteria in delivering the Unit learning outcome (ULOs), need tremendous work without AI. It includes observation, generating suggestions and continues integration. Beyond those features, application of AI not only provides personalized empowerment for students and faculty but also enhances their learning and research capabilities with analyzable and traceable data.

The core structure of AI appliances in this proposal:

- A. Computing Layer:** University Based Cloud Deployment – Accessibility
- B. Data Layer:** Physically Isolated Dedicated Storage – Confidentiality
- C. Data Collection:** Non-Intrusive Behavioral Tracking – User friendly & Data Genuity
- D. Output Layer:** Human-In-The-Loop Validation – Integrity

Many universities have implemented practical applications: the AI adaptive learning platform developed by Stanford University include Open TutorAI, PAPPL (Personalized AI-Powered Progressive Learning), and simulations like SALT. The AI learning assistant at the MIT realizes full-process personalized reminders for pre-class preview and after-class review, focusing on the long-term cultivation of students' learning habits. It shows the intensive compilation of using AI, the most powerful tools, in advance learning and edging research.

To sum up, AI empowers management and outreach by driving an efficiency revolution rooted in data reusability and research-driven service enhancement. Through the reconstruction of governance logic via a closed-loop data system and the precise engagement of target audiences using intelligent insights, AI ultimately translates into reputational premiums and sustained academic research assets. This process enables a qualitative leap in the institution's overall competitiveness.

4. Research Funding

1. Maritime and Port Authority (MPA) Call for Proposal

<https://www.mpa.gov.sg/maritime-singapore/innovation-and-r-d/call-for-proposals>



2. Singapore–Horizon Europe (SG-HE) Complementary Fund, National Research Foundation

<https://www.nrf.gov.sg/singapore-horizon-europe-complementary-fund/>



3. NRF Fellowship (Class of 2027)

Class of 2027 – Open



5. Publication

Jack Ng-Soon-Chye, Navigating algorithmic markets in Singapore: A practitioner reflection on the backstage realities of algorithmic path dependence and sustainable online retailing, *Journal of Tropical Futures OnlineFirst*, 5 May 2026.

<https://journals.sagepub.com/doi/10.1177/27538931261446664>

