

Project Feliann: An AI-Orchestrated, Low-Code Research Support Ecosystem For Student Success and Service Efficiency in a Philippine Higher Education Institution

Tyrone Marcial V. Lopez,¹ and Myrna V. Viernes²

ORCID ID: 0009-0001-3156-8526¹

Aldersgate College Inc., Solano, Nueva Vizcaya, Philippines^{1,2}
rctr@aldersgate.edu.ph

ABSTRACT

Student researchers at small and mid-sized higher education institutions (HEIs) in the Global South wait days, sometimes longer, to secure a research adviser, get compliance documents verified, or receive statistical feedback. This paper documents the design, development, and pilot implementation of Project Feliann, an AI-orchestrated, low-code research-support ecosystem built by the Research, Planning, and Analytics Services (RPAS) office of Aldersgate College Inc., a private HEI in Nueva Vizcaya, Philippines. Following Design Science Research (DSR) methodology, the team combined a low-code orchestration platform (n8n), a large language model (Google Gemini), and existing Google Workspace infrastructure to automate four processes: adviser matching, a secure academic-standing chatbot, compliance triage for data-analysis requests, and preliminary statistical auto-analysis. The pilot began in December 2025 and is approximately 90 percent complete. Two of the four workflows already run on live student data. Internal, non-independently audited estimates point to large reductions in turnaround time and an annual infrastructure cost of roughly ₱4,741, or about US\$84, though formal validation remains pending. The project was separately recognized in the 2026 World's Universities with Real Impact (WURI) rankings, placing 31st in the Student Support and Engagement category. This paper offers one of the few documented accounts of a low-code, AI-orchestrated alternative to enterprise RPA built in-house at a resource-constrained Philippine HEI, and it sets out a concrete agenda for the empirical validation the system has not yet undergone.

Keywords: *artificial intelligence in education, robotic process automation, low-code automation, academic advising, chatbot, design science research, data privacy*

INTRODUCTION

Robotic process automation (RPA), increasingly paired with AI, has streamlined admissions, records management, and student support across universities worldwide, cutting processing time and cost while freeing staff for higher-value work (Bhardwaj & Kumar, 2025). Most documented cases come from large, well-funded institutions running commercial enterprise platforms. The University of Melbourne saved an estimated 10,000 staff hours a year; the University of Auckland saved roughly 23,000. Little is documented about whether smaller institutions without a dedicated automation budget can reach similar gains.

The gap matters because administrative bottlenecks are not unique to large universities. At many small and mid-sized HEIs, a student's path from choosing a thesis topic to receiving usable statistical feedback still runs through manual, person-dependent steps. Coordinators match students with advisers by hand, verify compliance documents themselves, and compute statistics on request. This friction compounds student research anxiety and places uneven demands on faculty (cf. Liu et al., 2025, on the stakes of fair student-supervisor matching).

This was the operating reality of the RPAS office at Aldersgate College Inc. before Project Feliann, a system named in memory of two RPAS leaders, Felicidad O. Furigay and Annabelle Leabres, whose mentorship shaped the office's culture. Three further gaps in the literature motivate the study. First, most educational chatbots handle general inquiries rather than authenticated access to sensitive individual records (Okonkwo & Ade-Ibijola, 2021; Wang et al., 2023). Second, formal fairness-in-matching models such as Liu et al.'s (2025) are rarely reported as embedded in low-cost, non-specialist-maintained production systems. Third, compliance with the Philippine Data Privacy Act of 2012 (RA 10173) is only moderately implemented among sampled HEIs, with breach notification and registration as the weakest areas (Castro, 2021), and no documented Philippine case describes an AI system built from the outset to close that gap rather than widen it.

This paper documents Feliann's design and development using DSR (Hevner et al., 2004; Peffers et al., 2007), describes its pilot implementation and early impact estimates, situates the system within the relevant literature, and proposes an empirical research agenda for independent validation.

METHODOLOGY

Research Design

The study follows Peffers et al.'s (2007) six-activity Design Science Research Methodology: problem identification, solution objectives, design and development, demonstration, evaluation, and communication. This framework fits a project whose central output is a working artifact rather than a tested hypothesis. Before Feliann, the RPAS office's manual

processes produced multi-day turnaround times and no systematic way to balance faculty workload. The loss of two senior staff in 2020 and 2021 sharpened the institution's motivation to preserve its mentorship culture under reduced staffing.

System Architecture

Feliann reuses infrastructure the institution already licensed instead of procuring new enterprise software. n8n, a self-hosted low-code orchestration platform, forms the coordination layer. It receives triggers such as form submissions, spreadsheet updates, and chat messages, then routes them through conditional logic. Google Gemini handles natural-language understanding and generation, including chain-of-thought reasoning before the workflow commits to a decision. Google Sheets serves as the operational database. Gmail dispatches notifications. Google Forms and Drive handle intake and document submission. This low-code layer, built on existing productivity software, is the main architectural difference from enterprise RPA suites such as UiPath and Automation Anywhere used in large-university case studies.

The Four Workflows

Table 1 summarizes the four workflows that make up Project Feliann.

Workflow	Trigger	Core Decision Logic	Output
1. Adviser Matching	New thesis/capstone topic submitted	Keyword extraction matches topic to faculty specialization tags. The least-loaded qualified adviser is selected.	Auto-generated adviser endorsement
2. Class Standing Chatbot	Student initiates chat	Student ID and PIN authenticate the request. A sub-workflow retrieves only that student's grade row. The LLM translates the figures into plain-language feedback.	Conversational grade summary
3. Data-Analysis Request (Compliance Engine)	Student submits statistical support request	Chain-of-thought review checks ethics clearance, format compliance, and adviser endorsement. The analyst with the lowest honorarium count is assigned.	Accept/return decision, analyst assignment
4. Auto-Analysis (Accelerator)	Accepted request with uploaded dataset	The system automatically computes descriptive and inferential statistics, then routes the output to an analyst for mandatory review.	Draft statistical report, analyst-reviewed before release

Ethical and Data-Governance Design

Given the documented compliance gap among Philippine HEIs (Castro, 2021), the chatbot requires two-factor identification, student ID plus a separate PIN, before it releases any grade information. Retrieval is isolated in a sub-workflow, so the conversational agent never sees more than one authenticated student's record. Every AI-generated output that reaches a student passes through human review before release. No formal Privacy Impact Assessment has been conducted, and this paper does not independently confirm the system's National Privacy Commission registration status. Both are flagged as future work in Section 5.

Pilot Scope and Timeline

The pilot is scoped to the RPAS office rather than the institution as a whole. After informal testing, the system entered production use on December 12, 2025. The team's internal assessment currently places the pilot at approximately 90 percent completion.

3. Results

Functional Status

Two workflows, Adviser Matching and the Compliance Engine, run fully on live student data. The Class Standing Chatbot is operational and in active use. Auto-Analysis remains under refinement. Generating statistically reliable analysis automatically is harder than the rule-based logic that drives the other three workflows.

Implementation Iteration

Early piloting exposed the chatbot endpoint through a free local tunneling service, ngrok, which produced intermittent downtime as sessions expired. The team resolved this by migrating to a dedicated, low-cost virtual private server. The episode illustrates the kind of design iteration DSR's demonstration activity anticipates: an early convenience choice proved inadequate under sustained use and was revised.

3.3 Projected Operational Efficiency

Table 2 sets out the team's internal, non-audited projections at piloted scale

Metric (internal estimate)	Manual Process (baseline)	With Feliann (projected)
Estimated annual request volume	~500+ requests/year	~500+ requests/year
Estimated manual labor per request	~2.5 hours	Minutes (automated)
Estimated total labor-hours displaced/year	—	~1,250 person-hours
Estimated monetized value of displaced labor	—	~₱93,750/year
Annual infrastructure (VPS) cost	—	~₱4,741 (~US\$84)
Projected ROI ratio	—	~19:1

Note. Figures are internal, non-independently audited estimates, offered for transparency and to support the future-research agenda, not as validated results.

Qualitative Observations

Adviser-matching requests once depended on a coordinator's personal familiarity with faculty availability. They now run through an auditable least-loaded selection rule. Data-analysis requests are assigned by lowest cumulative honorarium count, replacing an informal, loosely tracked allocation. Both observations echo, in a simpler form, the fairness concerns Liu et al. (2025) raise formally.

External Recognition

Feliann was recognized in the A1, Student Support and Engagement, category of the 2026 WURI rankings, placing 31st globally. WURI's methodology combines invited-evaluator scores, a judging panel, and a generative-AI-assisted evaluation of submitted initiatives, not an independently audited empirical study. Read the recognition as an external, qualitative signal that evaluators found the system's design credible, not as confirmation of the Table 2 figures.

DISCUSSION

Impact in Context

Feliann's projected 1,250 annual person-hours displaced looks small next to the 10,000 and 23,000 hours reported for the University of Melbourne and the University of Auckland (Bhardwaj & Kumar, 2025). The scale difference explains most of that gap. A fairer comparison is proportional: Feliann's projected gains cost roughly ₱4,741 a year in infrastructure, several orders of magnitude below what a commercial RPA license and dedicated automation team typically require. If independent evaluation bears out these figures, it would suggest that automation gains need not stay locked behind enterprise-scale budgets.

A Low-Cost Alternative

This paper's main practical contribution is the low-code layer built on existing infrastructure. It shows that a small RPAS team, without dedicated software-engineering staff, can build and maintain a multi-workflow system of this kind. Whether the approach generalizes to other small HEIs facing similar constraints is a question for the replication study proposed in Section 5.

Fairness Mechanism in Context

Feliann's rules, least-loaded adviser and lowest-honorarium analyst, are deterministic and single-criterion. Liu et al. (2025) model fairness with multiple weighted attributes. Feliann trades that richness for a rule a small team can implement and explain, at the cost of the qualitative fit between a student's interests and an adviser's expertise that a multi-attribute model could capture.

Privacy-by-Design

Feliann's authentication and data-segregation choices respond directly to the compliance gap Castro (2021) documents among Philippine HEIs. Without a formal Privacy Impact Assessment or independently confirmed NPC registration, though, the system's privacy posture is best described as well-intentioned and structurally sound rather than certified compliant. That certification belongs to the institution's data protection officer, not to this manuscript.

Limitations

This paper documents a single workflow ecosystem in a single department of a single private Philippine HEI, so it makes no claim of generalizability without replication. The DSR cycle has reached demonstration, not independently audited evaluation. Table 2 reports projections, not measured outcomes. The corresponding author also served as the

system's principal developer, a conflict of interest managed by explicitly labeling every unverified figure and recommending independent evaluation in Section 5. AI-generated outputs, chatbot grade interpretation, and draft statistical analysis go through mandatory human review, but no systematic audit of their error rate against analyst-verified ground truth has been conducted yet.

CONCLUSIONS AND RECOMMENDATIONS

This paper has documented the DSR-guided development and pilot of a low-code, AI-orchestrated research-support ecosystem built within a small Philippine HEI's existing infrastructure. It stops short of claiming validated empirical impact: the efficiency figures in Section 3 are internal projections, and the privacy safeguards in Section 2.4 have not undergone formal independent assessment. Instead, it offers a transparent account of how a small, non-specialist team can build such a system at low cost, backed by external recognition: the system's 2026 WURI ranking, which evaluators found credible.

The natural next step, and this paper's central recommendation, is for the same team, ideally with independent reviewers, to subject Feliann to the empirical scrutiny its pilot status has not yet allowed:

1. A controlled, timestamped time-and-motion study comparing manual and automated cycle times across at least one full academic term.
2. An independent, finance-office-audited recalculation of the projected annual labor-value savings and return-on-investment ratio.
3. A systematic accuracy audit of AI-generated outputs against analyst-verified ground truth, with error rates reported before further scaling.
4. Structured satisfaction and trust surveys of student and faculty users.
5. A formal Privacy Impact Assessment and independent confirmation of NPC registration status.
6. A multi-site replication study at other small or mid-sized Philippine HEIs.

AI DECLARATION STATEMENT

AI-supported tools helped with grammar and structure in creating this article. No AI tools were used to collect, analyze, or interpret data; all authors are credited for their academic contributions.

REFERENCES

- Bhardwaj, V., & Kumar, M. (2025). Transforming higher education with robotic process automation: Enhancing efficiency, innovation, and student-centered learning. *Discover Sustainability*, 6, Article 356. <https://doi.org/10.1007/s43621-025-01198-6>
- Castro, C. K. (2021). The implementation of the Data Privacy Act among higher educational institutions in the Province of Pangasinan. *Asian Journal of Multidisciplinary Studies*, 4(2), 103–111.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>
- Liu, X., Yue, Q., Hu, B., & Tao, Y. (2025). Graduate students and supervisors matching decision-making considering stability-based fairness based on TOPSIS and grey correlation degrees. *Scientific Reports*, 15, Article 22902. <https://doi.org/10.1038/s41598-025-00388-6>
- Okonkwo, C. W., & Ade-Ibijola, A. (2021). Chatbots applications in education: A systematic review. *Computers and Education: Artificial Intelligence*, 2, Article 100033. <https://doi.org/10.1016/j.caeai.2021.100033>
- Peppers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. <https://doi.org/10.2753/MIS0742-1222240302>
- Segovia-García, N. (2024). Optimizing student support: A review of the use of AI chatbots in higher education. *European Public & Social Innovation Review*, 9, 1–20. <https://doi.org/10.31637/epsir-2024-324>
- Wang, T., Lund, B. D., Marengo, A., Pagano, A., Mannuru, N. R., Teel, Z. A., & Pange, J. (2023). Exploring the potential impact of artificial intelligence (AI) on international students in higher education: Generative AI, chatbots, analytics, and international student success. *Applied Sciences*, 13(11), Article 6716. <https://doi.org/10.3390/app13116716>
- World's Universities with Real Impact. (2026). WURI 2026 rankings. <https://www.wuri.world>