International Multidisciplinary Research Journal
ISSN: 2424-7073
Available Online at https://imrjournal.info/
Volume V; November 2026; Page No.94-105
DOI: https://doi.org/10.47722/imrj.2001.95

IMR JOURNAL CANADA

Research Article

DESIGNING FOR WORKFORCE DEMAND: AN AGILE, AI-ASSISTED PROCESS FOR DEVELOPING AN APPLIED AI BACHELOR'S DEGREE

Zahra Sadiq[1]; Sukhman Kaur[2] J. McCuaig[3]

[1 & 2]*Research Assistants, School of Computer Science, University of Guelph, Guelph, Canada*

[3]*Associate Professor Emeritus, School of Computer Science, University of Guelph, Guelph, Canada*

ARTICLE INFO

Article History:

Received 15.07.2026
Accepted 15.10.2026
Published 30.11.2026

Keywords:

AI literacy,
curriculum development,
applied AI,
agile methodology,
workforce readiness

ABSTRACT

67.7% of employers rated AI literacy as essential or very important for new hires, according to a cross-sector employer survey (n = 171; 77.8% completion rate). Employers primarily emphasized the responsible use and critical evaluation of AI outputs rather than AI system development. Yet a concurrent environmental scan of 157 Canadian AI-related programs found only 16 with an applied focus. A new Bachelor of Artificial Intelligence degree with a major in Human-Centred AI at the University of Guelph is a concrete response to that gap. This paper presents a replicable methodological blueprint for rapid, AI-assisted development, demonstrated through a team of three producing a complete degree proposal, ready for external review, in approximately three months. The development process comprised five components: a systematic scan of Canadian and international AI programs; a cross-sector employer survey with quantitative and NLP-assisted qualitative analysis; supplementary stakeholder consultations with faculty, current students, and parents; a multi-agent AI consultation process for developing and refining program learning outcomes; and AI-assisted program sequencing to resolve prerequisite structures and scheduling constraints. Artificial intelligence tools were used substantively at each stage (document synthesis, thematic analysis, learning outcome generation, and logistics modelling), and this use is documented as a methodological contribution in its own right. The resulting degree proposal is organized around employer-validated competencies: responsible AI use, critical evaluation of outputs, data governance literacy, and domain-integrated applied learning. We discuss the design decisions behind the curriculum, the role AI played at each stage, and the limitations encountered in practice. The methodical blueprint offers a replicable approach for institutions seeking to develop AI literacy programs rapidly and responsibly in response to documented workforce demand.

Copyright© The Author(s) 2025. Open Access under the Creative Commons Attribution 4.0 International License (CC BY 4.0). This license permits use, sharing, adaptation, distribution, and reproduction in any medium or format, giving proper credit to the original author, including a link to the license, and indicating if any changes were made. To view a copy of the licence, please visit https://creativecommons.org/licenses/by/4.0/

1. Introduction

A cross-sector employer survey conducted as part of this project found that 67.7% of respondents across diverse sectors rated AI literacy as essential or very important for new hires.

The dominant workforce expectation was not that graduates build AI systems, but that they use AI tools responsibly, evaluate outputs critically, and understand the governance and ethical implications of AI deployment.

Yet a parallel scan of 157 Canadian AI-related programs found only 16 with an applied focus. The majority of existing offerings are concentrated in technical computing and machine learning disciplines, with limited attention to applied, cross-disciplinary AI literacy. Conventional program development cycles are too slow to respond to the rapidly evolving field of Artificial Intelligence (AI). This paper presents a program development process that actively engaged AI agents in the creation of a university degree in AI. Scholarly works concerning the integration of AI tools into curriculum and program design are sparse. Existing research has largely focused on AI's role in delivering curriculum, such as personalized learning systems, adaptive assessments, and tutoring. The use of generative AI to support the full arc of degree-program development, from environmental scanning through learning outcome generation and program sequencing, has not been documented as an established practice (Karataş et al., 2025).

This paper documents that practice. A team of three employed AI agents as collaborators at each stage of developing a new Bachelor of Artificial Intelligence (B.A.I) degree. The development team's direct experience with AI, including the impacts of its affordances and limitations, is detailed in the following sections.

2.Background

2.1 AI Literacy

AI literacy, a subset of digital literacy, has become increasingly relevant in literature. While AI literacy can be interpreted in different ways, some constants remained unchanged across literature: the ability to use and understand AI (Liu et al., 2025; Long & Magerko, 2020), efficient and adaptive application (Chee et al., 2025), critical evaluation of AI (Chiu, 2025), and the ethical implications associated with the use of AI (Chee et al., 2025).

Zhao et al. (2022) introduced a concise framework which encompasses effective curriculum development. Their four-dimensional model encompasses knowledge, application, evaluation, and ethics (Chee et al., 2025), which serves as the working definition of AI literacy for curriculum development. By adopting this structure, the curriculum supports responsible AI use while enabling an interdisciplinary approach, ensuring the program is not limited to students with computer science backgrounds but open to learners from diverse fields. Given the growing importance of AI literacy across disciplines, institutions must be able to develop relevant curricula efficiently, despite the lengthy processes typically involved in academic program development.

2.2 Curriculum Development Timelines in Higher Education

Developing and implementing a new degree program for higher education is time-intensive and can take 2-3 years from conception to launch (Ontario Universities Council on Quality Assurance, n.d.; Trent University, n.d.). At the University of Guelph (UoG), units proposing new programs must first submit a Notice of Intent, which highlights the program's rationale, anticipated demand, enrollment projections, and resource implications, which must be reviewed and approved by Institutional Research and Planning. A New Program Proposal,

an external evaluation, a response to the assessment report, and subsequent revisions, follow before institutional approval and external ministerial consent. (University of Guelph, n.d.).

In parallel, a development team must conduct environmental scans, stakeholder consultations, and iterative curriculum design in advance of proposal submission. This pre-submission work often constitutes 4-12 months of coordinated labour (Duklas, 2015). The Human-Centred AI major for the B.A.I degree was developed to the point of external review readiness in approximately one-fourth of the usual time. AI tools were used at each stage to reduce the time required for synthesis, stakeholder alignment, and draft generation without displacing the interpretive and evaluative judgment of the human team. The result was a compressed, thoroughly completed, and carefully reviewed process outlined in the following sections.

3.An AI-Assisted Program Development Process

The proposal for the BAI: Human-Centred AI major was developed between January and March of 2026. The AI-assisted degree proposal process comprised six key tasks.

1. An agent-assisted systematic environmental scan of Canadian and international AI programs
2. A cross-sector employer survey with agent-assisted quantitative analysis
3. Stakeholder consultations with faculty, current students, and parents
4. AI-assisted financial modelling
5. A multi-agent AI consultation process for developing and refining program learning outcomes
6. Agent-assisted course selection and program sequencing.

The development team comprised three people: two undergraduate co-op students who led the environmental scan, employer survey administration, and supplementary data collection; and one contracted senior academic who directed the project, reviewed and finalized the survey instruments, led the learning outcomes and program sequencing work, and finalized the proposal. The environmental scan and survey instrument were developed in parallel during the first month; survey deployment overlapped with the learning outcomes work and initial course selection in the second; and program sequencing, synthesis, and proposal writing occupied the third. AI assistance greatly reduced the time required to integrate and analyze. AI agents were invaluable for managing the growing corpus of working documents by extracting and summarizing relevant content, thereby eliminating the need for team members to re-read full documents to locate specific findings. They also supported the synthesis of large document sets, primarily during the environmental scan. Other agentic collaboration included thematic analysis of open-text survey responses; structured multi-agent consultation for generating and refining program learning outcomes; and logistics modelling for resolving prerequisite structures and scheduling constraints in the program sequence. All AI-generated outputs were reviewed and adjudicated by the human team.

3.1 External Needs Analysis

In the first phase of the environmental scan, a comprehensive enumeration of Canadian post-secondary institutions was carried out to identify all AI-related credentials at any level,

resulting in a catalogue of 157 programs. Of these, 68 were bachelor-level degrees. 16 of those were applied programs rather than computational. In the second phase, a qualitative analysis was conducted on 24 domestic and international comparators. Programs were selected based on structural similarity to the proposed degree: applied orientation, cross-disciplinary framing, and explicit AI literacy mandate.

Perplexity Pro was used for web-sourced research and program discovery as it retrieved verifiable information about Canadian and international AI credentials. ChatGPT was used to support general synthesis and question-and-answer tasks as the scan progressed, helping organize findings across programs. The division of tools by task emerged through practice as the team identified which tools performed most reliably for each type of work. A representative prompt used during the synthesis phase took the following form: "Here is the program page for [Program Name] at [Institution]. Please summarize the program's structure, stated learning outcomes, and overall orientation, specifically whether the program appears applied, technical, or both. Note any interdisciplinary elements, experiential learning components, and any explicit mention of AI literacy or responsible AI use." All AI responses were reviewed by the human team against primary institutional sources before being incorporated into the catalogue.

Two types of data quality issues were encountered. Some institutions provided incomplete or outdated information on their public-facing websites, requiring the team to draw on multiple sources or record gaps explicitly in the notes field. Second, the applied versus technical distinction proved difficult to identify consistently, as many programs combined both orientations without clearly prioritizing one. In cases of ambiguity, the team manually assigned a program to the most appropriate categories. These classification decisions were made by the human team rather than delegated to AI tools, as the distinction required interpretive judgment about program intent that AI agents could not reliably produce independently.

Employers were selected as a primary external stakeholder for this program proposal. A key program outcome is to produce graduates who apply AI responsibly in professional contexts, and employers are the most direct judges of whether that preparation is adequate. The University of Guelph Co-operative Education office facilitated distribution across employment sectors. The survey was administered through Qualtrics and was open from February to March 2026. It measured four structured constructs: AI exposure timeline, AI literacy importance, expected graduate competency levels, and perceived credential advantage. Participation was voluntary.

A total of 171 employers responded, of whom 133 (77.8%) completed all questions. Respondents represented diverse sectors; multiple sector selections were permitted. The five largest were Engineering or Technology (35.1%), Business or Finance (28.1%), Science or Research (15.8%), Sustainability or Environment (13.5%), and Social Sciences or Public Policy (6.4%).

An AI agent assisted with the design of the R script to analyze the results by suggesting metrics and charts that were appropriate to the data composition, supporting the selection of suitable analytical methods for the survey results. The resulting analysis was completed substantially faster than would have been feasible for a single human analyst. Patterns were consistent across sectors: 76.7% of respondents indicated that new hires were likely or extremely likely to encounter AI-enabled tools within three to five years; 67.7% rated AI literacy as essential or very important for new hires; and the most commonly expected graduate

competencies were responsible use of AI-enabled tools as part of normal work (57.3%) and the ability to interpret or evaluate AI outputs in support of decisions (49.1%). Fewer than a quarter of respondents expected graduates to build or modify AI systems directly, confirming that applied AI fluency and not system development is the primary workforce need.

3.2 Internal Validation

Three stakeholder groups were consulted alongside the employer survey. AI agents were used at this stage to identify preliminary survey questions, to evaluate questions for clarity and precision, to generate synthetic data and identify analysis mechanisms a priority, and to estimate the impact of varying the number of questions in each instrument. AI agents helped refine question clarity, assess alignment with program learning outcomes, and identify potential analysis approaches before data collection.

Faculty within the School of Computer Science (SOCS) were surveyed to assess existing expertise and capacity relative to the proposed curriculum. Agent assistance was used to refine the question set so that each question was explicitly linked to one of the proposed learning outcomes for the program, to better illustrate alignment and gaps. The agent was prompted with the list of learning outcomes and the proposed list of questions and was asked to identify the linkages. Iterative prompting resulted in a more cohesive set of questions that did not over-represent anyone learning outcome.

Prospective student data were gathered at a STEM outreach event in February 2026 and during tours of the university campus in April. AI agent consultation for these surveys was limited to creating and proofing the survey questions.

Guardian surveys were conducted across two channels: a survey at an on-campus outreach event ($n = 17$ complete responses) and a social media outreach ($n = 6$ complete responses), for a combined total of 23 usable responses. AI agents were employed to inspect and combine this data, with the directive to ensure conclusions respected the small n and did not overstate any observations. The agent observed that the data were too few to be combined for quantitative analysis but offered several qualitative observations. For example, one pattern that directly informed program design was that many guardians flagged mathematics or programming prerequisites as a potential barrier. The program has since been deliberately positioned as an applied degree that does not require prior programming experience and that does not rely on advanced high school mathematics.

Historically, institutional financial modelling for new program proposals is produced by the Decision Support office late in the development process, once the program structure is sufficiently stable to support enrolment projections. In an accelerated development timeline, this creates a gap: faculty and administrative stakeholders need to assess financial projections at the beginning of the process before a formal model is available.

An AI agent was used to construct a parameterized financial model in Excel, enabling early financial projections while the official institutional model was not yet available. Inputs included: cohort size, attrition rate, co-op participation, and fee schedule; outputs were projected revenue, cost, and teaching staff requirements. The parameterized structure lets stakeholders adjust assumptions and see financial implications early in the program development process.

The official financial model was received from Decision Support on May 1, 2026. Its projections closely aligned with those produced by the AI-assisted model three months earlier, validating the modelling approach. Institutional processes calibrated to a conventional twelve-month development cycle create information gaps that an agile, AI-assisted team can bridge with appropriate tooling.

3.3 Curriculum Synthesis

Learning outcomes were created by combining results from a multi-agent AI consultation that probed what a forward-looking outcome set should look like, with the draft Program Learning Outcomes (PLOs) produced by a curriculum committee. The Program learning outcomes are listed in section 4.

The consultation ran in two rounds, and both were approximately seven months apart. In October 2025, two generative AI agents, ChatGPT (OpenAI) and Gemini (Google), were speculatively consulted using identical prompt sequences. In May 2026, a third agent, Claude (Anthropic), was consulted using the same sequence as a validation step. Fresh instances were used throughout both consultations to ensure responses reflected model reasoning rather than context-driven responses. ChatGPT and Gemini were initially selected to compare perspectives across different generative AI models, while Claude was consulted in the second round to assess whether the same themes persisted across model families and over time. The goal was to compare results across different model families and across time. The prompt sequence was designed to elicit forward-looking reasoning before requesting curriculum-level outputs. Rather than requesting PLOs directly, the consultation began by asking agents to reason about the evolving role of humans in computing. The sequence used in all three consultations was as follows:

1. **Prompt 1:** Predict the role that humans will have in computing in 5 years.
2. **Prompt 2:** 10 years?
3. **Prompt 3:** What should students be learning in the next 4 years? Suppose a student was starting now in a program that was called AI _____ where the blank could be any of many words.
4. **Prompt 4:** Generate program learning outcomes for these degrees. There will be overlap, so let's say generate 5 PLOs and indicate which of those degrees they are suitable for.

This futures-first prompt progression was intentional. Without it, an agent prompted directly for PLOs produces outcomes shaped by existing programs in its training data. The framing used shifts responses toward predictions about what graduates will need. Despite the seven-month gap and three distinct model families, the PLO sets converged on five themes: critical evaluation of AI outputs; ethical reasoning and institutional accountability; problem formulation and scoping under uncertainty; communication across disciplinary and technical boundaries; and adaptive practice as a durable meta-competency. The themes persisted regardless of how agents named the degree types in the third prompt. The consistency across model families and across time reflects a broadly shared inference about what will be durable and distinctive in applied AI education.

The agent-generated PLO sets were reconciled against three main inputs: the employer survey findings, UoG's institutional learning outcomes, and eight draft PLOs that were developed by a curriculum committee in late 2025. The committee and the agents converged on core competencies. The employer survey confirmed them: responsible use, critical evaluation, data governance, adaptability, and ethical awareness were most consistently expected across sectors. The final PLOs consolidate the committee's eight outcomes into five integrated capabilities, each traceable to the agent consultations, committee work, and employer survey.

Constructing the program sequence required identifying a pool of existing UoG courses that mapped onto the five PLOs, selecting from that pool, and resolving the resulting prerequisite structure into a coherent four-year progression. The first step was an external scan of course listings in comparable programs. This established the course types and level distribution typical of existing applied AI degrees.

The second step was identifying eligible courses within the UoG calendar. An AI-assisted scan of the course catalogue was attempted but produced unreliable results. The calendar is available only as a PDF formatted in two columns, a structure that generative AI tools parsed inconsistently. After reviewing a sample of AI-generated output against the source document and finding substantial errors, the team pivoted to a direct manual scan. Each team member reviewed sections of the calendar and recorded candidate courses in a shared spreadsheet. This was one of the clearer cases in the project where AI assistance was not net-positive.

Once outlines were obtained for the candidate courses, each outline was submitted to an AI assistant with the five PLOs and a prompt requesting a mapping table indicating which PLOs a given course addressed and at what depth. This worked well: the structured output was consistent, and the mappings were defensible. The human team reviewed and adjudicated the mapping table, making corrections where the AI had over-credited surface-level content coverage.

Prerequisite resolution proved to be beyond reliable AI assistance. Attempts to use an agent to construct a four-year schedule that respected all prerequisite chains, credit-hour constraints, and course availability patterns produced outputs with logical errors. The team resolved the prerequisite structure manually, using the AI-generated mappings as a guide.

4.The Proposal

The proposed credential is a four-year honours B.A.I with a major in Human-Centred AI, expected to admit students in Fall 2027. The organizing framework of the degree is the learning outcomes describing responsible AI use, critical evaluation of outputs, AI governance, and human-centred design. The major title reflects the program’s emphasis on evaluating, applying, and governing AI systems within organizational and societal contexts, rather than primarily developing the underlying technologies.

- **PLO 1** Critically analyze and evaluate artificial intelligence systems by identifying their capabilities, limitations, assumptions, and impacts within real-world sociotechnical contexts.

- **PLO 2** Apply statistical, data-analytic, and AI tools to support evidence-based decision-making, including evaluating data quality, provenance, uncertainty, bias, and the validity of AI-supported conclusions.
- **PLO 3** Design and evaluate human-AI interactions that support usability, transparency, explanation, and appropriate human oversight, informed by principles from psychology, human-computer interaction, and the social sciences.
- **PLO 4** Identify, assess, and address ethical, legal, and professional responsibilities associated with the deployment of AI systems, including issues of accountability, risk, governance, and harm mitigation.
- **PLO 5** Communicate AI-related analyses, decisions, and recommendations clearly and effectively to diverse audiences, and apply AI tools responsibly within a chosen domain or professional context.

The program organizes required courses into a foundations stream and a set of restricted elective clusters. Each cluster is organized around one of the five PLO themes and draws from courses across multiple disciplines. Students complete 1.0 credit from each cluster, ensuring PLO coverage while preserving disciplinary breadth. Year 1 establishes AI as a sociotechnical system alongside programming for data, statistical reasoning, UI design, and introductory ethics. This foundational emphasis reflects the 67.7% of surveyed employers who rated AI literacy as essential or very important for new hires. Year 2 deepens applied capability through human-AI interaction, data science, and social implications, and offers the first applied project course. The human-AI interaction course responds to the 57.3% of employers who expect graduates to use AI-enabled tools responsibly, while data science and applied projects address the 49.1% who expect graduates to interpret and evaluate AI outputs. Year 3 addresses organizational contexts, governance, and ethics in depth. Year 4 centres on synthesis, culminating in an AI practicum requiring demonstration of responsible AI application in a real context. Required courses are drawn from computer science, philosophy, statistics, psychology, and management. No single discipline provides sufficient coverage for AI literacy. Students complete 1.0 credit from each of five PLO-themed restricted elective clusters.

5.Development Timeline

Phase 1: Scoping (January 2026)

Program Learning Outcomes were drafted for the major early in January 2026, marking the formal start of the accelerated development phase. The environmental scan of comparable programs ran from January 12 to January 25, 2026, progressing from a broad overview to an in-depth analysis. A project whiteboard was established on January 28 to support collaborative tracking, and biweekly council meetings began on January 27. Stakeholder surveys targeting employers, students, and faculty were developed from January 29 onward, with iterative revisions continuing through April. The first draft of course search and sequencing was completed between January 30 and February 2, and a preliminary draft program took shape from February 2 to February 7.

Phase 2: Parallel Design (February 2026)

February was a period of parallel activity. Personas were created from February 4 to 6. In-house budget modelling was conducted throughout the month and presented to faculty on

February 24. The first draft of the AI major proposal was developed between February 11 and February 25, and AI minor PLOs were finalized by February 25. Community engagement continued during this period, including participation in a daylong session for high school students to learn programming on February 7 and campus tour outreach to support student survey efforts.

Phase 3: Consolidation (March-June 2026)

March brought a consolidation phase. Course development brainstorming occurred on March 4, followed by curriculum resequencing and the launch of AI course proposal writing on March 9. Draft consultation documents for all units whose courses were required were created beginning on March 11. The administrative survey was distributed on March 9, with results received on April 9. CIMS downtime from March 18 to March 26 interrupted progress during this period. The Work Experience (WEX) document was developed between March 23 and April 2.

In early April, a new degree structure to house the program was proposed and approved in principle by senior administrators. In May, the official financial model was received from Decision Support, and it closely aligned with the in-house modelling completed in February, validating the team's earlier projections. The proposal is being internally validated and has an external review scheduled for June 2026.

The compressed timeline required simultaneous management of stakeholder consultation, curriculum design, and administrative approvals, precluding a fully sequential review process and demanding continuous iteration on surveys and course sequencing. The parallel activities were, in part, successful because of the workflow integration with AI agents, demonstrating the viability of an agile, collaborative approach to program development under time and resource constraints.

6.Discussion

6.1 What Worked

AI assistance was most effective where the task involved structured synthesis of large, heterogeneous document sets. During the environmental scan, submitting program descriptions with a fixed prompt template produced consistent, reviewable summaries far faster than manual reading would have permitted. The same pattern held for the employer survey's open-text responses: AI-assisted thematic coding identified recurring themes across 171 responses in a fraction of the time required for manual coding, and the themes survived preliminary human review without major revision.

The multi-agent PLO consultation worked well for a different reason: it externalized the question of what a forward-looking applied AI program should emphasize and surfaced convergent answers from three distinct model families. The convergence itself was analytically useful; it provided a warrant for the PLO set that went beyond the development team's own judgment.

At the process level, AI assistance enabled parallel execution of tasks that would conventionally be sequential. Draft survey questions could be reviewed and refined by an agent while the environmental scan was still underway; preliminary course mappings could be generated while the survey was in the field. The compression of the timeline from a

conventional twelve-plus months to approximately four months is attributable in substantial part to this parallelism rather than to any single AI-assisted task.

The document distillation pattern of using agents to extract relevant content from the growing project corpus provided continuous leverage throughout all phases, not only at the formal synthesis tasks. In general, the prompting of the agent was key. Agents often needed to be explicitly limited with respect to input data (e.g. "use only the information in the document I have uploaded"). In general, it was beneficial to have the prior history available to the agent via setting up projects or coworking folders.

The financial modelling effort proved highly valuable, and the resulting tool has the potential to be adopted more broadly across the college.

6.2 Where AI Fell Short

Not all AI assistance was effective. The most clear-cut failure was the attempt to use generative AI to parse the UoG course calendar. The document's two-column PDF format produced systematic parsing errors, and the institutional specificity of the content exceeded what the model could reliably reconstruct from formatting alone. The lesson is practical: AI-assisted document synthesis performs well on well-structured, web-accessible text and poorly on institutionally idiosyncratic, format-constrained documents.

Prerequisite resolution presented a different limitation. The task required holding a complex, interdependent constraint structure, including prerequisite chains, scheduling patterns, credit loads, and cohort progression, in view simultaneously and generating a schedule free of logical errors. Generative AI tools produced plausible-looking schedules that contained errors on inspection. This appears to reflect a known limitation of current models in multi-step constraint satisfaction rather than a failure of the prompting strategy.

Finally, the program itself has not yet run. This paper documents a design process and a resulting credential structure; it makes no claim about learning outcomes or graduate employment. The question of whether the AI-assisted design process produces programs that are as effective as those developed through conventional processes is an open research question.

6.3 Transparency and Ethics of AI Use

AI tools were used substantively throughout this project, and that use is documented here as a methodological contribution rather than disclosed as a caveat. The tools used were Perplexity Pro, for web-sourced program discovery during the environmental scan; ChatGPT (OpenAI GPT-4 series), for document synthesis and general question-and-answer tasks; Gemini (Google), for the October 2025 PLO consultation; and Claude (Anthropic), for the May 2026 PLO consultation, open-text thematic analysis, and course-to-PLO mapping. All AI-generated content was adjudicated by the research team, and decisions are documented in the relevant sections above. This project has a recursive dimension that warrants acknowledgment: an AI-assisted process was used to design a program in AI literacy. The competency structure the program teaches reflects direct experience with AI tools as design instruments. The program's emphasis on human oversight and output evaluation emerged in part from observing the conditions under which AI assistance failed or required correction during development. The authors followed the conference's AI disclosure requirements. The use documented here is

consistent with the norms of academic transparency: AI tools accelerated and structured the work; they did not author it.

6.4 Limitations

As the program is launching in the Fall of 2027, it is not possible to evaluate the effectiveness of AI-designed curriculum. This includes not being able to assess graduate employment rates and whether the evident demand translates upon graduation as well as the effectiveness of the design of an AI-assisted curriculum. Although the employer survey was distributed to hundreds of employers across multiple sectors, participation was voluntary, and the 171 responses might not fully represent the broader workforce. Future research should compare the outcomes of AI-assisted and traditionally developed curricula, specifically in terms of student learning outcomes, graduate employment and alignment with workforce demands.

7.Conclusion

This paper has described an agile, AI-assisted approach to designing a new Bachelor of Artificial Intelligence degree with a major in Human-Centred AI in a fraction of the usual development time. Grounded in an environmental scan, employer survey, and stakeholder consultations, the resulting program is explicitly aligned with workforce demand for applied AI literacy rather than system development alone. The methodology is transferable, but effective replication depends on conditions similar to those at Guelph, including at minimum one senior academic with curriculum design authority and junior researchers able to sustain structured data collection, as well as access to networks for large-scale survey distribution. What is most readily portable is the methodology itself: structured prompt templates for environmental scan synthesis, a multi-agent consultation sequence for learning outcomes, AI-assisted course-to-PLO mapping, and an agile, parallel workflow, while course sequencing and prerequisite resolution still demand substantial human effort. More broadly, the work illustrates how AI can function as a partner in academic program design, not a replacement for human expertise, and points toward future research on the long-term effectiveness of AI-assisted curriculum development in shaping graduate outcomes and labour market readiness.

8.Acknowledgments

This work was conducted at the University of Guelph with financial support from the College of Computational, Mathematical, and Physical Sciences.

Artificial Intelligence Tool Disclosure

In accordance with conference AI disclosure requirements, the authors confirm that the following AI tools were used during the work and preparation of this paper: Perplexity Pro (program discovery and web research during the environmental scan); ChatGPT/GPT-4/5 (document synthesis, survey instrument refinement, and open-text thematic analysis); Gemini (program learning outcome consultation, October 2025); and Claude (program learning outcome consultation, May 2026; course-to-PLO mapping; proposal and paper drafting assistance). All AI-generated content was reviewed, adjudicated, and, where necessary, revised or overridden by the research team. Final authorial responsibility rests with the human authors.

References

Chee, H., Ahn, S., & Lee, J. (2025). *A competency framework for AI Literacy: Variations by different learner groups and an implied learning pathway. British Journal of Educational Technology, 56(5), 2146–2182.* https://doi.org/10.1111/bjet.13556

Chiu, T. K. F. (2025). *AI literacy and competency: Definitions, frameworks, development and future research directions. Interactive Learning Environments, 33(5), 3225–3229.* https://doi.org/10.1080/10494820.2025.2514372

Duklas, J. (2015). *New program approval practices: A summary of current typologies at Ontario colleges and universities. Duklas Cornerstone Consulting.* https://duklascornerstone.ca/wp-content/uploads/2021/05/An-Ontario-study-on-new-program-approval-practices.pdf

Karataş, F., Eriçok, B., & Tanrikulu, L. (2025). *Reshaping curriculum adaptation in the age of artificial intelligence: Mapping teachers' AI-driven curriculum adaptation patterns. British Educational Research Journal, 51(1), 154–180.* https://doi.org/10.1002/berj.4068

Liu, X., Zhang, L., & Wei, X. (2025). *Generative artificial intelligence literacy: Scale development and its effect on job performance. Behavioral Sciences (2076-328X), 15(6), 811. (186207742).* https://doi.org/10.3390/bs15060811

Long, D., & Magerko, B. (2020). *What is AI Literacy? Competencies and design considerations. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20, 1–16.* https://doi.org/10.1145/3313831.3376727

Ontario Universities Council on Quality Assurance. (n.d.). *Guidance on Timelines for Cyclical Program Reviews. Retrieved May 13, 2026, from* https://oucqa.ca/guide/guidance-on-timelines-for-cyclical-program-reviews/

Polomoshnov, J., Masso, A., & Lobanova, K. (2026). *Comparing Digital, Data, and AI Literacy: A narrative review. Information Polity, 31(1), 13–26.* https://doi.org/10.1177/15701255251401863

Trent University. (n.d.). *New Degree Program Development—Provost & VP Academic. Retrieved May 13, 2026, from* https://www.trentu.ca/vpacademic/quality-assurance/new-degree-program-development

University of Guelph. (n.d.). *New Programs. Retrieved May 13, 2026, from* https://www.uoguelph.ca/quality-assurance/new-programs

Zhao, L., Wu, X., & Luo, H. (2022). *Developing AI literacy for primary and middle school teachers in China: Based on a structural equation modeling analysis. Sustainability, 14(21), 14549.*

How to cite this article:
Sadiq Z.; Kaur S.; McCuaig J. (2026) 'Designing for Workforce Demand: An Agile, AI-Assisted Process for Developing an Applied AI Bachelor's Degree, *International Multidisciplinary Research Journal*, Volume V; November 2026; Page 94-105. DOI: https://doi.org/10.47722/imrj.2001.95