Agentic AI Workflows for Teaching and Learning, Research, and Community Service: A Cross-Disciplinary Framework for Tertiary Education Adoption

Okundaye, B.O.
Department of Computer Science Education
Federal College of Education (Technical), Akoka
benjamin.okundaye@fcetakoka.edu.ng

Abstract

The emergence of agentic artificial intelligence (AI) systems capable of autonomous goal pursuit through planning, tool use, persistent memory, and multi-step execution represents the most profound transformation of institutional practice since the advent of the internet. Unlike prior generations of AI tools that respond to discrete prompts, agentic systems can initiate, coordinate, and complete entire workflows with minimal human intervention. For the modern tertiary education institution, which fulfils simultaneous obligations in teaching and learning, knowledge production, and community service, this transformation is not merely a question of efficiency but of fundamental capability expansion. This paper presents a comprehensive, cross-disciplinary scholarly analysis of agentic AI deployment across all three institutional functions. Drawing on verified, peer-reviewed literature from 2022 through 2025, the paper examines: (1) agentic workflows in Teaching and Learning, including instructional content generation, AI-assisted essay marking using rubrics, and adaptive pedagogical applications; (2) agentic AI in Research, with emphasis on literature synthesis, Zotero integration via the Model Context Protocol (MCP), research proposal pipelines, and citation verification; and (3) Community Service applications encompassing neighbourhood security and home automation. A dedicated section addresses ethical risks — including academic integrity threats, algorithmic bias, surveillance overreach, and equity of access — and proposes a governance framework grounded in institutional best practices. The paper concludes that agentic AI, deployed responsibly within a framework of institutional governance, AI literacy, and human oversight, represents an extraordinary opportunity to fulfil the university's three-fold mission with unprecedented scale and quality.

Keywords: Agentic AI, Large Language Models, Claude Model Context Protocol, NotebookLM, Automated Essay Scoring, Research Automation

Introduction

Artificial intelligence has progressed with extraordinary speed from narrow-task automation to sophisticated agentic systems capable of browsing the internet, writing and executing code, managing files, calling external APIs, and coordinating fleets of sub-agents to accomplish complex, long-horizon goals. This progression is, as Acharya et al. (2025) describe, a 'qualitative leap' in which agents plan, use tools, maintain memory, and self-adjust to changing environments to accomplish multi-step, multi-goal tasks that rule-based systems fundamentally cannot achieve. The broader literature — including the foundational Agent4EDU framework of Dai et al. (2024), the evolutionary analysis by Kamalov et al. (2025), and the workflow-theoretic work of Jiang et al. (2025) — converges on a shared conclusion: agentic systems represent not an upgrade to existing AI tools but an architectural paradigm shift with transformative potential across all domains of human intellectual activity.

For universities whose missions span teaching, research, and community service, the implications are profound. Institutions that equip their graduates with agentic AI fluency — and deploy these tools responsibly — will have a decisive advantage in fulfilling their educational, research, and civic mandates.

Defining Agentic AI

For the purposes of this paper, 'agentic AI' is defined following Hosseini and Seilani (2025) as systems that combine autonomy, reactivity, proactivity, and learning ability, representing a field-wide shift from Copilot (assisted) to Autopilot (autonomous) modes of operation. Four architectural pillars undergird all agentic systems (Kamalov et al., 2025; Jiang et al., 2025): Reflection — the capacity to review, critique, and revise one's own outputs; Planning — multi-step goal decomposition, task sequencing, and contingency management; Tool Use — structured interfaces to external systems including search engines, code interpreters, APIs, databases, and file systems; and Multi-Agent Collaboration — coordination between specialised agents each contributing domain-specific expertise to collective goals.

The foundational agentic reasoning paradigm, ReAct (Reason + Act), was introduced by Yao et al. (2023) at ICLR 2023. ReAct prompts LLMs to generate both verbal reasoning traces and actions in an interleaved manner, enabling dynamic planning that updates as new information is gathered from external tools. This paradigm underlies virtually all contemporary production agentic frameworks.

Scope and Structure

This paper is structured across seven substantive sections covering Teaching and Learning (Section 2), Research (Section 3), Community Service (Section 4), the Enabling Technology Ecosystem (Section 5), Ethics and Governance (Section 6), and an Implementation Roadmap with Conclusions (Section 7).

Agentic AI for Teaching and Learning

Teaching and learning constitute the primary public-facing mission of the university. Agentic AI systems offer transformative potential at every stage of the instructional cycle, from curriculum design through to student assessment and adaptive feedback.

Preparation of Instructional Content

The traditional model of instructional content preparation places the full cognitive burden of research, synthesis, structuring, and formatting on individual academic staff. Agentic AI workflows can dramatically reduce this burden without diminishing pedagogical quality. Jiang et al. (2025), in their Agentic Workflow for Education (AWE) framework grounded in the von Neumann Multi-Agent System architecture, distinguish AWE from traditional LLM-based linear interactions by enabling scalable, personalised, and collaborative task execution through dynamic, nonlinear workflows. This architecture is directly applicable to instructional content generation: rather than a single prompt requesting a lecture, an AWE system orchestrates multiple specialised agents — a literature search agent, a curriculum mapping agent, a Bloom's Taxonomy alignment agent, a graphics agent, and a formatting agent — each contributing to a holistic, quality-assured output.

In practice, the workflow proceeds as follows. An instructor specifies a topic, learning objectives, target student level, and preferred pedagogical approach. The orchestration layer (n8n or Make) activates a Claude Sonnet 4 or Opus 4 agent that queries the Zotero MCP server for relevant prior readings, invokes a web search tool for recent developments, and passes all gathered context to the primary generation model. The result — a fully structured, referenced, Bloom's-aligned lesson plan with lecture notes, in-class activities, discussion questions, and assessment ideas — is then delivered to the learning management system (LMS) via API. Acharya et al. (2025) confirm that agentic systems can synthesise standards, student data, and teacher feedback to propose weekly lesson plans, curate materials, and align them with institutional goals while learning from past cycles.

The AWE proof-of-concept case study by Jiang et al. (2025) demonstrated that agentic workflows for automated assessment item generation produce items statistically comparable to real exam questions, validated at the 95% confidence level. Complete question-bank generation pipelines integrate Claude Opus 4 for item generation, a Bloom's Taxonomy alignment checker ensuring cognitive diversity, rubric and model answer generation sub-agents, and a plagiarism screen against existing question banks. Google's NotebookLM (powered by Gemini) provides a complementary capability: it converts uploaded course documents into Audio Overviews, structured Slide Decks, and Video Overviews, transforming source material into accessible visual presentations (Wikipedia, 2025; TechCrunch, 2025).

AI-Assisted Essay Marking Using Rubrics

State of the Art in Automated Essay Scoring

Automated Essay Scoring (AES) has been transformed by frontier LLMs. Pack et al. (2024), in a rigorous study published in Computers and Education: Artificial Intelligence, investigated the validity and reliability of four LLMs (PaLM 2, Claude 2, GPT-3.5, and GPT-4) for scoring ESL student writing against a holistic rubric. GPT-4 demonstrated the strongest reliability, with Pearson correlation coefficients significant at the 0.001 alpha level. Xiao et al. (2024), in 'Human-AI Collaborative Essay Scoring: A Dual-Process Framework with LLMs,' demonstrated that while LLMs do not surpass state-of-the-art grading models in raw performance, they exhibit notable consistency, generalisability, and explainability and that AI-assisted human graders outperformed either acting alone. The multi-agent AutoSCORE framework (Wang et al., 2025), in which each rubric dimension is assessed by a dedicated sub-agent before scores are aggregated, demonstrated improved consistency compared to end-to-end LLM scoring. Ke and Ng (2024), in their IJCAI-24 paper on recent successes in AES, caution that LLMs' interpretation of rubrics can differ significantly from human interpretation, underscoring the necessity of calibration, anchor exemplars, and human moderation.

The Rubric-Based Agentic Marking Pipeline

A complete rubric-based essay marking pipeline for institutional deployment integrates the following components: (a) Pre-processing student submissions are ingested and standardised, with an initial plagiarism screen run via Turnitin or Copyleaks API; (b) Rubric injection the full rubric including dimensional descriptors, band criteria, and weightings is injected into the agent's structured prompt; (c) AI scoring Claude Opus 4 generates structured JSON output containing numerical scores per rubric dimension, qualitative comments citing specific textual evidence, an overall feedback paragraph, a suggested overall grade, and prioritised improvement recommendations; (d) Consistency verification a second-pass agent compares scores against pre-marked anchor exemplars, flagging anomalous scores for human review; (e) Bias screening an equity audit sub-agent checks for systematic score disparities across student demographic groups (Pack et al., 2024); and (f) Report generation a formatted feedback report is delivered per student via the LMS API, with aggregate class-level analytics delivered to the instructor.

The explainability advantage of LLM-based AES is significant. As Xiao et al. (2024) observe, LLMs can generate detailed natural language feedback on student writing — the primary mechanism through which assessment supports learning. Zheng et al. (2025), in their QwenScore+ framework, further demonstrated that rubric-aware Chain-of-Thought (CoT) prompting combined with reinforcement learning from human feedback (RLHF) significantly outperforms baseline LLMs on feedback quality metrics including BLEU, ROUGE-L, and cosine similarity against expert-written feedback.

Academic Integrity Integration

The integration of AI into assessment raises serious academic integrity concerns. Recent analysis indicates AI-assisted cheating incidents in higher education increased by nearly 400% from 2022–23 to 2024–25 (IRJMETS, 2025). The marking pipeline incorporates Turnitin or Copyleaks similarity screening, AI content detection as one signal among many, and a Claude-assisted contextual review distinguishing proper quotation, paraphrasing without citation, and genuine originality. As the EDUCAUSE AI Ethical Guidelines (Georgieva & Stuart, 2025) emphasise, AI integrity tools must be embedded within institutional policies that communicate clear expectations and provide proportional, evidence-based consequences.

Additional Teaching and Learning Applications

Acharya et al. (2025) identify personalised learning at unprecedented scale as the primary benefit of agentic AI for students. The appropriate deployment model is augmentation: AI agents that monitor LMS performance data, library engagement, and forum activity to recommend resources, generate tailored practice exercises, and alert instructors to emerging student welfare concerns — while preserving students' autonomy. Acharya et al. (2025) and Hosseini and Seilani (2025) both identify 24/7 AI teaching assistants as the most immediate high-value application of agentic AI in higher education, grounded in course materials stored in a vector database for accurate, hallucination-resistant responses. For international students and students with disabilities, agentic workflows can generate translated versions of course materials in 95+ languages, plain-language summaries, audio descriptions of visual content, and closed captions for lecture recordings via OpenAI Whisper API with Claude post-processing. Agents connected to student record systems, LMS analytics, and library data can generate weekly at-risk student reports identifying patterns that predict dropout risk with sufficient lead time for effective intervention (Acharya et al., 2025).

Agentic AI for Research

Research the systematic generation of new knowledge is the domain where agentic AI's capacity for information synthesis, structured multi-step reasoning, and external tool integration has perhaps the most immediate and measurable impact.

Literature Discovery and Synthesis

The traditional systematic literature review requires weeks of manual searching across multiple databases, screening thousands of abstracts, and synthesising findings. An agentic literature review pipeline compresses this process to hours while maintaining full reproducibility through automated logging of all search strings, databases queried, and inclusion/exclusion decisions. A production pipeline demonstrated by FreeCodeCamp (2025) shows how to reduce roughly six hours of manual review into a five-minute automated process using parallel collection from multiple academic APIs

(Semantic Scholar, OpenAlex, arXiv, PubMed), normalisation and deduplication, structured AI extraction of key findings, relevance and quality scoring, and delivery to Google Sheets with an automated narrative summary.

Zotero Integration via the Model Context Protocol

The Model Context Protocol

The Model Context Protocol (MCP) was announced by Anthropic in November 2024 as an open standard for connecting AI assistants to data systems including content repositories, business management tools, and development environments (Anthropic, 2024a). MCP resolves the 'N×M data integration problem' through a universal, JSON-RPC 2.0 based interface in which MCP Servers expose tools and resources that any MCP-compatible Client can discover and invoke. By March 2025, OpenAI had officially adopted MCP across its products including the ChatGPT desktop app, and Anthropic subsequently donated MCP to the Agentic AI Foundation under the Linux Foundation, signalling its transition to a community-governed open standard (Wikipedia, 2025). For universities, MCP is strategically critical: any AI system can interact with institutional data sources through a single standardised protocol, eliminating the need for bespoke integrations for each AI tool.

Zotero MCP Server and Research Pipeline

Multiple Zotero MCP server implementations now enable AI-powered semantic search over a researcher's library, retrieval of full citation metadata, extraction of PDF annotations and highlights, collection management, and bibliography generation in any citation style (GitHub, 2025a; GitHub, 2025b; PulseMCP, 2024). The practical integration, documented by Magalhães (2025) and the Nounai Librarian guide (2025), involves installing the zotero-mcp package and adding server configuration to Claude Desktop. Once configured, direct prompts such as 'Summarise the papers in my Zotero library about machine learning in education' retrieve, synthesise, and respond grounded in the researcher's own verified literature.

Integrating the Zotero MCP server with n8n orchestration and Claude Opus 4 enables a complete end-to-end research proposal generation pipeline. Drawing on n8n's documented multi-agent research architecture (n8n, 2025a), the pipeline proceeds through requirements intake, literature synthesis

(Zotero MCP + web search agents), methodology drafting, funding alignment (retrieving funder strategic priority documents via web fetch agent), bibliography assembly verified against CrossRef for DOI accuracy, and a final consistency review generating a submission-ready document. Citation verification operates by checking each in-text citation against the Zotero library, resolving DOIs via CrossRef, retrieving abstracts to confirm relevance through semantic similarity, and flagging citations that do not substantively support the citing claim for human review — directly addressing the hallucination problem that has made academic adoption of LLM-generated text hazardous (Kamalov et al., 2025).

Data Analysis and Research Synthesis Tools

Claude Code (Anthropic, 2025b) is an agentic coding tool that operates in the terminal, understands the codebase, and enables complex data analysis through natural language commands: reading and inspecting data, identifying appropriate statistical tests, writing and executing analysis code in Python or R, interpreting results, generating publication-quality visualisations, and drafting results sections in academic prose. Internally, Claude Code was adopted by 80% of Anthropic engineers by day five of release, with pull request throughput increasing 67% (Agent Factory, 2025). Google's NotebookLM (powered by Gemini), purpose-built for researcher use, enables uploading of up to 300 sources and interaction with an AI grounded exclusively in those sources — eliminating hallucination from outside the uploaded corpus. Features include Audio Overviews, Video Overviews, Deep Research integration for autonomous web research, and inline citations linking every claim to specific source passages (Wikipedia, 2025; TechCrunch, 2025; Google Workspace, 2025).

AutoGen (Wu et al., 2023), developed by Microsoft Research, enables construction of collaborative AI research teams in which specialised agents engage in structured multi-agent conversation to accomplish complex research tasks. CrewAI (2024) offers a complementary role-based framework with agents assigned defined roles, backstories, and tool access. Both frameworks rely on the ReAct pattern (Yao et al., 2023), in which each agent alternates between reasoning and acting through tool calls, with transparent reasoning chains providing human-interpretable explanations — a critical feature for institutional trust and governance.

Agentic AI for Community Service

Universities have a moral and civic obligation to the communities in which they are embedded. Agentic AI systems offer new tools for fulfilling this obligation at scale — from enhancing neighbourhood safety to empowering residents through intelligent home systems and multilingual information services.

Neighbourhood Security Systems

The global home security systems market was valued at USD 72.4 billion in 2025 and is estimated to reach USD 109.4 billion by 2030, driven primarily by AI integration (Digitalholics, 2026). AI-enabled security cameras can now distinguish between incidental movement and a human security threat, reducing false alarms by up to 80% compared to traditional motion sensors (SafeHome.org, 2025). For 2025 and beyond, the primary trend is AI as an intelligence layer over existing surveillance cameras, access controls, and IoT sensors — improving situational awareness and enabling security personnel to focus on high-priority tasks (RealTime Networks, 2024).

A university-deployable community security system based on open-source components and privacy-by-design principles integrates: edge detection using IP cameras connected to edge computing nodes running lightweight detection models (YOLOv8, MobileNet SSD); OpenCV-based multi-camera management and event-driven triggers; n8n workflow orchestration that filters false positives, applies Claude-based contextual reasoning to identify patterns, and routes confirmed events to appropriate stakeholders; and community notification via Telegram Bot or WhatsApp Business API while maintaining privacy through data minimisation and resident consent management.

The deployment of AI-powered surveillance raises serious ethical concerns requiring explicit governance. A privacy-by-design architecture must incorporate: no facial recognition capability; local processing with no raw video transmission to cloud servers; community governance of data access and retention; independent oversight by a Community Safety Advisory Committee; transparency through publicly available technical specifications; and sunset provisions requiring periodic community reauthorisation. These requirements align with the principles articulated by Georgieva and Stuart (2025) in the EDUCAUSE AI Ethical Guidelines framework.

Home Automation and Smart Environment Systems

Home automation platforms have evolved from simple scheduled timers to sophisticated AI-orchestrated environments. An agentic home management system built on Claude API, connected to Home Assistant (the leading open-source home automation platform with 3,000+ integrations) via n8n or Make, enables residents to interact in their native language to control any connected device, create complex automation routines expressed in natural language, query energy consumption and receive optimisation recommendations, and receive proactive maintenance alerts (Axe Fire & Security, 2024).

For elderly or disabled residents living independently, passive wellbeing monitoring systems using PIR sensors, door contacts, and activity monitors (without camera-based monitoring) enable Claude to analyse activity patterns against each resident's personal baseline, generating welfare alerts when deviations suggest potential concern. All monitoring data is processed locally, with notifications delivered only to the resident's chosen contact with explicit consent — consistent with GDPR, Nigeria's NDPR, and equivalent data protection frameworks worldwide. Universities with community engagement mandates can additionally deploy Claude-powered chatbots in 95+ languages on community websites and WhatsApp, grounded in verified local authority documents updated automatically via n8n web-scraping pipelines.

The Enabling Technology Ecosystem

This section surveys the key enabling technologies referenced throughout this paper, organised by functional category, as a technology selection guide for practitioners.

Large Language Model Foundations

Claude Sonnet 4 (Anthropic) offers balanced capability and speed, ideal for teaching workflows, rubric-based marking, research synthesis, and community chatbots, with a 200K token context window. Claude Opus 4 (Anthropic) provides the highest reasoning capability in the Claude family, achieving state-of-the-art results on leading benchmarks (Anthropic, 2025a), and is best suited for complex research tasks, multi-step proposal writing, and nuanced analysis. Claude Code / Agent SDK (Anthropic, 2025b; 2025c) provides an agentic coding and general-purpose agent harness supporting subagents, hooks, background tasks, and parallel workflows. NotebookLM (Google Gemini) offers source-grounded research synthesis, Audio/Video Overviews, Deep Research, and Slide Decks, with user data not used for model training (Google Workspace, 2025). AutoGen (Wu et al., 2023) and CrewAI (2024) provide multi-agent research frameworks. Llama 3.1 405B (Meta, open weights) offers on-premises deployment for data-sensitive environments.

Orchestration Platforms

n8n is an open-source, self-hostable workflow automation platform combining visual building with custom code, 400+ integrations, and native AI capabilities, with 180,500+ GitHub stars and 200,000+ community members (n8n, 2025b). For universities, the self-hosted deployment model is critical for data protection compliance. n8n's AI Agent nodes provide native LLM integration through LangChain, enabling single agents, multi-agent systems, and multi-agent teams within a visual, no-code environment (n8n, 2025a; Perficient, 2025). Human-in-the-loop checkpoints can be inserted at any workflow stage.

Make (Integromat) is a cloud-based alternative with 1,500+ application integrations, appropriate for administrative workflows not involving sensitive student data. The Model Context Protocol (MCP), introduced by Anthropic in November 2024 and now maintained by the Agentic AI Foundation (Linux Foundation), provides a universal standard enabling any AI system to interact with institutional data sources without bespoke integrations (Anthropic, 2024a; Wikipedia, 2025). Available MCP servers immediately relevant to academic practice include Zotero MCP, GitHub MCP, Google Drive MCP, Slack MCP, PostgreSQL MCP, and Filesystem MCP (PulseMCP, 2024).

ReAct: The Foundational Agentic Reasoning Pattern

The ReAct framework, introduced by Yao et al. (2023) at ICLR 2023 (arXiv:2210.03629), remains the foundational architectural pattern of virtually all production agentic systems. ReAct's core insight is that LLMs should generate both verbal reasoning traces and task-specific actions in an interleaved manner: reasoning traces help the model induce, track, and update action plans, while actions allow it to interface with external knowledge bases or environments. Empirical results demonstrate that ReAct outperforms chain-of-thought reasoning alone on tasks requiring external information retrieval, while the explicit reasoning trace provides human-interpretable explanations — a critical feature for institutional trust and governance.

Ethics, Risks, and Governance Frameworks

The Ethical Stakes

The scale and speed of AI adoption in higher education is itself a risk indicator. Recent data indicates 92% of UK students now use AI in some form, up from 66% in 2024, and 88% have used generative AI for assessments, with AI-assisted cheating incidents up nearly 400% between 2022–23 and 2024–25 (IRJMETS, 2025). Over 60% of educators in advanced economies were using AI in their classrooms by 2023 (World Economic Forum, 2025). Georgieva and Stuart (2025) identify core ethical tensions as: authorship and attribution; assessment validity; data privacy; equity of access; and preservation of human academic relationships. Tan and Maravilla (2024) invoke Smuha's (2022) Trustworthy AI principles of fairness, accountability, and transparency as the ethical bedrock for AI in education, while cautioning against assuming technological innovations are inherently beneficial.

Risk Assessment Summary

Academic integrity risks (AI cheating up 400%) are mitigated through process-based assessment, oral examinations, clear institutional AI use policies, and AI detection as one signal among many (Georgieva & Stuart, 2025; Tan & Maravilla, 2024). Assessment bias: LLMs may score differently across demographic groups (Pack et al., 2024) is mitigated through bias audits before deployment and human moderation. Data privacy risks from student data processed by third-party APIs are addressed through on-premises deployment (n8n, open-source LLMs) and GDPR/NDPR compliance (Azevedo et al., 2025). Surveillance overreach in community security systems is mitigated through privacy-by-design architecture, no facial recognition, and community governance. Equity of access concerns — AI tools widening gaps between well-resourced and under-resourced institutions — are addressed through open-source-first strategies (n8n, open-weight LLMs, NotebookLM free tier) and consortium licensing. Misinformation and hallucination risks are addressed through citation verification pipelines, grounded generation (NotebookLM, RAG architectures), and mandatory human expert review gates (Kamalov et al., 2025).

Proposed University AI Governance Framework

Synthesising frameworks proposed by Azevedo et al. (2025), Georgieva and Stuart (2025), and Tan and Maravilla (2024), this paper proposes the following governance structure. An Institutional AI Governance Committee, a multi-stakeholder body including academic staff, students, IT, legal, ethics, community partners, and senior management, oversees policy development, technology review, annual audit, and appeals. An AI Policy Framework covers acceptable use in teaching, assessment, research, and community engagement, including definitions of permissible and prohibited AI assistance in student work and data processing requirements. An AI Ethical Review Process extends institutional ethics review to cover AI-generated content, autonomous data collection, and community deployment, with particular attention to equity impacts and vulnerable population protections. An AI Literacy Programme provides mandatory training for all academic and professional staff and optional certification for students. Transparency and Accountability provisions require an annual public AI Use Transparency Report covering deployed systems, purposes, outcomes, and incidents, with accessible complaints and appeals mechanisms. All AI deployments are subject to two-year reauthorisation requiring evidence of compliance, equity impact assessment, and community satisfaction review.

Implementation Roadmap and Conclusions

Phased Implementation Roadmap

Phase 1: Foundation (Months 1–3): Deploy Claude API institutional access; establish MCP servers for Google Drive and LMS; train 50 academic staff in prompt engineering and agentic workflow basics; pilot AI-assisted lecture note generation in five departments; constitute AI Governance Committee.

Phase 2: Assessment Integration (Months 4–6): Deploy rubric-based AI marking pilot with human moderation in three programmes; build Zotero MCP integration with citation verification; launch 24/7 student AI tutoring chatbot grounded in course materials; publish institutional AI Policy Framework.

Phase 3: Research Pipelines (Months 7–12): Deploy full research proposal pipeline via n8n + Claude MCP; integrate NotebookLM for research synthesis; launch Claude Code workshops for data analysis; develop community security pilot in one partner neighbourhood with full privacy governance.

Phase 4: Community Scale-Up (Year 2): Expand community security and home automation to all partner areas; deploy multilingual community information chatbots; establish open-source MCP server repository; publish first Annual AI Transparency Report.

Phase 5: Maturity and Leadership (Year 3+): Annual bias audits; model upgrade cycles; community needs assessments; contribution to national AI policy frameworks; publication of research findings on institutional AI adoption outcomes.

Conclusions

This paper has presented a comprehensive, empirically grounded framework for the deployment of agentic AI across the three core functions of the modern tertiary education institution: teaching and learning, research, and community service. The analysis demonstrates that this transformation is not merely quantitative — though the productivity gains documented in the literature are substantial — but qualitative: institutions can now offer genuinely personalised learning at scale, conduct research with a speed and comprehensiveness previously unattainable, and engage with communities through services that were simply not possible before.

The foundational theoretical work of Yao et al. (2023) on ReAct reasoning architectures, Wu et al. (2023) on multi-agent collaboration, Dai et al. (2024) on agentic educational frameworks, and Jiang et al. (2025) on the Agentic Workflow for Education establishes that agentic AI in academic contexts is not speculative it is a documented, rapidly maturing field with peer-reviewed evidence of efficacy. The parallel developments in AES documented by Pack et al. (2024), Xiao et al. (2024), Wang et al. (2025), and Ke and Ng (2024) demonstrate that rubric-based AI marking can be deployed with validity and reliability comparable to human assessment while providing dramatically superior feedback volume and consistency.

The enabling technology ecosystem n8n, Make, Claude Code, the Claude Agent SDK, NotebookLM, Zotero MCP, the Model Context Protocol, OpenCV, AutoGen, and CrewAI is production-ready, increasingly open-source, and deployable by institutions without access to technology-company resources. The barrier to adoption is not primarily technical but institutional: the need for clear governance frameworks, updated academic integrity policies, and AI literacy across the academic community. The ethical and governance frameworks proposed by Georgieva and Stuart (2025), Azevedo et al. (2025), and Tan and Maravilla (2024) provide a strong evidence base for constructing these frameworks responsibly.

Universities that invest now in the governance infrastructure, human capability, and technical integration required to deploy agentic AI responsibly will have a measurable advantage in fulfilling their educational, research, and civic mandates. The age of agentic AI has arrived. The defining question for universities is not whether to engage, but how to engage with the rigour, equity, and ethical commitment that have always been the hallmarks of the academy at its best.

References

Agent Factory. (2025). Claude Code origin story. Panaversity Agent Factory. https://agentfactory.panaversity.org/docs/General-Agents-Foundations/general-agents/origin-story

Anthropic. (2024a). Introducing the Model Context Protocol. Anthropic News. https://www.anthropic.com/news/model-context-protocol

Anthropic. (2025a). Introducing Claude Opus 4. Anthropic News. https://www.anthropic.com/news/claude-opus-4

Anthropic. (2025b). Claude Code overview. Claude Code Documentation. https://docs.claude.ai/en/docs/claude-code/overview

Anthropic. (2025c). Building agents with the Claude Agent SDK. Anthropic Engineering Blog. https://www.anthropic.com/engineering/building-agents-with-the-claude-agent-sdk

Axe Fire & Security. (2024). The future of home security: Tech trends in 2024. https://www.axefiresecurity.com/blog/tech-trends/future-of-home-security-trends-in-2024

Azevedo, R., Robert, A., & McCormack, M. (2025). Institutional policies on artificial intelligence in higher education: Frameworks and best practices for faculty. New Directions for Adult and Continuing Education. https://doi.org/10.1002/ace.70013

Case Western Reserve University. (2025, December 22). Google NotebookLM receives major updates. University Technology [U]Tech. https://case.edu/utech/about/utech-news/google-notebooklm-receives-major-updates

CrewAI. (2024). CrewAI: Framework for orchestrating role-playing AI agents [Software]. GitHub. https://github.com/crewAIInc/crewAI

Dai, L., Shi, J., Zhao, J., Wu, X., Xu, R., Jiang, Y.-H., & He, L. (2024). Agent4EDU: Advancing AI for education with agentic workflows. In Proceedings of the 2024 3rd International Conference on Artificial Intelligence and Education (ICAIE 2024). ACM. https://doi.org/10.1145/3722237.3722268

Daily Dose of DS. (2026, January 17). Implementing ReAct agentic pattern from scratch. https://www.dailydoseofds.com/ai-agents-crash-course-part-10-with-implementation/

Digitalholics. (2026, January 20). AI home security: Complete guide. https://digitalholics.com/ai-home-security-complete-guide/

Effortless Academic. (2025, December 15). NotebookLM review 2025: AI tool for researchers. The Effortless Academic. https://effortlessacademic.com/notebook-lm-googles-newest-academic-ai-tool/

EDUCAUSE. (2025, June 24). AI ethical guidelines. EDUCAUSE Library Resources. https://library.educause.edu/resources/2025/6/ai-ethical-guidelines

FreeCodeCamp. (2025, March). How to build an AI-powered research automation system with n8n, Groq, and academic APIs. FreeCodeCamp News. https://www.freecodecamp.org/news/build-an-ai-powered-research-automation-system-with-n8n-groq-and-academic-apis/

Georgieva, M., & Stuart, J. (2025, June). Ethics is the edge: The future of AI in higher education. EDUCAUSE Review. https://er.educause.edu/articles/2025/6/ethics-is-the-edge-the-future-of-ai-in-higher-education

GitHub. (2025a). 54yyyu/zotero-mcp: Zotero MCP — Connects your Zotero research library with Claude and other AI assistants via the Model Context Protocol [Software repository]. https://github.com/54yyyu/zotero-mcp

GitHub. (2025b). kujenga/zotero-mcp: Model Context Protocol (MCP) server for the Zotero API, in Python [Software repository]. https://github.com/kujenga/zotero-mcp

Google Workspace. (2025). NotebookLM: AI-powered research and learning assistant tool. https://workspace.google.com/products/notebooklm/

Hosseini, S., & Seilani, H. (2025). The role of agentic AI in shaping a smart future: A systematic review. Array, 26, 100399. https://doi.org/10.1016/j.array.2025.100399

IRJMETS. (2025). Ethical AI governance in higher education. International Research Journal of Modernization in Engineering, Technology and Science. https://www.irjmets.com/upload_newfiles/irjmets70800007778/paper_file/irjmets70800007778.pdf

Jiang, Y.-H., Shi, J., Tu, Y., Liu, T.-Y., Zhuang, X., Hu, H., Li, R., & Jia, R. (2025). Agentic workflow for education: Concepts and applications [Preprint]. arXiv. https://arxiv.org/abs/2509.01517

Kamalov, F., & colleagues. (2025). Evolution of AI in education: Agentic workflows [Preprint]. arXiv. https://arxiv.org/abs/2504.20082

Ke, Z., & Ng, V. (2024). Automated essay scoring: Recent successes and future directions. In Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI-24) (pp. 8969–8977). https://www.ijcai.org/proceedings/2024/0897.pdf

Lab.nounai-librarian.com. (2025, December 22). Zotero meets Claude: AI-enhanced research workflow. Lifelong Research. https://lab.nounai-librarian.com/en/zoteroclaude-2/

Magalhães, A. (2025, June 29). Building a Zotero MCP server for research management. Tuts Blog. https://arturmagalhaes.com/python/research/mcp/zotero/2025/06/29/zotero_server.html

n8n. (2025a). Research AI agent team with auto citations using OpenRouter and Perplexity [Workflow template]. n8n Template Library. https://n8n.io/workflows/2607-research-ai-agent-team-with-auto-citations-using-openrouter-and-perplexity/

n8n. (2025b). AI agentic workflows: A practical guide for n8n automation. n8n Blog. https://blog.n8n.io/ai-agentic-workflows/

Ohio State University, ASCode Distance Education Office. (2025). Agentic AI in higher education. ASCode OSU. https://ascode.osu.edu/news/agentic-ai-higher-education

Pack, A., Barrett, A., & Escalante, J. (2024). Large language models and automated essay scoring of English language learner writing: Insights into validity and reliability. Computers and Education: Artificial Intelligence, 6, 100234. https://doi.org/10.1016/j.caeai.2024.100234

Perficient. (2025, July 4). Intelligent agents with n8n: AI-powered automation. Perficient Blog. https://blogs.perficient.com/2025/07/04/intelligent-agents-with-n8n-ai-powered-automation/

PulseMCP. (2024). Zotero MCP server by Aaron Taylor. PulseMCP Directory. https://www.pulsemcp.com/servers/kujenga-zotero

RealTime Networks. (2024, December 9). Emerging AI security trends for 2025. https://www.realtimenetworks.com/blog/artificial-intelligence-trends-in-security

SafeHome.org. (2025, August 4). AI-powered home security: 2026 predictions. https://www.safehome.org/security-systems/ai-powered-security/

Tan, R., & Maravilla, J. (2024). Shaping integrity: Why generative artificial intelligence does not have to undermine education. Frontiers in Artificial Intelligence, 7, 1471224. https://doi.org/10.3389/frai.2024.1471224

TechCrunch. (2025, November 13). Google's NotebookLM adds 'Deep Research' tool, support for more file types. https://techcrunch.com/2025/11/13/googles-notebooklm-adds-deep-research-tool-support-for-more-file-types/

Wang, Y., Ding, Z., Wu, X., Sun, S., Liu, N., & Zhai, X. (2025). AutoSCORE: Enhancing automated scoring with multi-agent large language models via structured component recognition [Preprint]. arXiv. https://arxiv.org/abs/2509.21910

Weaviate. (2025, July 25). AI agent workflow automation with n8n and Weaviate. Weaviate Blog. https://weaviate.io/blog/agent-workflow-automation-n8n-weaviate

Wikipedia. (2025). Model Context Protocol. Wikipedia. https://en.wikipedia.org/wiki/Model_Context_Protocol

Wikipedia. (2025). NotebookLM. Wikipedia. https://en.wikipedia.org/wiki/NotebookLM

World Economic Forum. (2025, May). See why EdTech needs agentic AI for workforce transformation. WEF Stories. https://www.weforum.org/stories/2025/05/see-why-edtech-needs-agentic-ai-for-workforce-transformation/

Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., Awadallah, A. H., White, R. W., Burger, D., & Wang, C. (2023). AutoGen: Enabling next-gen LLM applications via multi-agent conversation [Preprint]. arXiv. https://arxiv.org/abs/2308.08155

Xiao, C., Zhou, C., Liang, M., Yang, L., Feng, X., & Shi, S. (2024). Human-AI collaborative essay scoring: A dual-process framework with LLMs [Preprint]. arXiv. https://arxiv.org/abs/2401.06431

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. In International Conference on Learning Representations (ICLR 2023). https://arxiv.org/abs/2210.03629

Zheng, Q., & colleagues. (2025). Explainable AI for education: Enhancing essay scoring via rubric-aligned chain-of-thought prompting (QwenScore+). International Journal of Modern Physics C. https://doi.org/10.1142/S0129183125420136