INSTITUTE of Museum and Library SERVICES

NATIONAL LEADERSHIP GRANTS FOR LIBRARIES

Sample Application LG-260439-OLS-26
Project Category: Applied Research

Northern Illinois University

Amount awarded by IMLS:	$617,597
Amount of cost share:	$0

Northern Illinois University, in partnership with Old Dominion University and the University of North Texas, will research whether library-archived university records, made accessible through specialized AI agents, can improve the accuracy and trustworthiness of student academic advising support. This research project is based on a specialized artificial intelligence architecture grounded in library curated university archives. Working with librarians, academic advisors, and artificial intelligence researchers across three institutions, the project will develop an open-source toolkit that uses specialized artificial intelligence agents to provide library-led student advising grounded in archived university records; evaluate its effectiveness through a controlled study across three campuses; and curate a "University DNA" archive of governance and policy records to support long term access and reuse.

Attached are the following components excerpted from the original application.

- Narrative
- Schedule of completion

When preparing an application for the next deadline, be sure to follow the instructions in the most recent Notice of Funding Opportunity for the grant program to which you are applying.

Library Led Student Advising Support Using Specialized AI Agents Grounded in University Records Archives

1. Introduction

Northern Illinois University (NIU), in partnership with Old Dominion University (ODU) and the University of North Texas (UNT), respectfully requests $617,597 in funding from the Institute of Museum and Library Services (IMLS), with $0 in cost share, for an Applied Research project to investigate **ATLAS (Advising Through Library Archived Sources)**, a specialized AI architecture grounded in library curated university archives. This project aligns with NLG-L Objective 1.3 (Provide broad access to and preservation of information and collections through libraries and archives), with Objective 1.2 (Improve community prosperity through libraries and archives) as a secondary benefit. The project addresses *three interconnected crises*: (1) undergraduate non-completion rates approaching 40% nationally, with six-year completion across all institution types remaining below 62% [1], [2], and information gaps and limited access to advising identified as contributing factors [3], [4]; (2) the rapid loss of university knowledge as faculty retire [5] and governance records go unpreserved [6]; and (3) the declining visibility of academic libraries as generative AI tools reshape how students seek information [7]. Working with librarians, academic advisors, and AI researchers across three institutions, the project will develop ATLAS, an open-source toolkit that uses specialized AI agents to provide library-led student advising grounded in archived university records; evaluate its effectiveness through a controlled study across three campuses; and curate a "*University DNA*" archive of governance and policy records to support long term access and reuse. Results will be disseminated through peer-reviewed conferences and journals, national conferences, and practitioner workshops. The research is guided by five questions:

- **RQ1 (Student Outcomes)**: How does library-led AI that connects academic, financial aid, and workforce information affect students' decision-making compared with generic AI tools such as ChatGPT?
- **RQ2 (Information Use and Confidence)**: To what extent do transparent, source-cited recommendations improve information literacy and students' confidence in evaluating and using guidance, particularly for first-generation and rural learners?
- **RQ3 (Preservation Impact):** How does the systematic preservation of shared governance records (e.g., Faculty Senate minutes, curriculum committee documents, course syllabi) affect the accuracy and reliability of AI-generated guidance?
- **RQ4 (Responsible Use)**: Under what conditions does the system provide weak or incomplete guidance, and what review workflows, including human review steps, are needed?
- **RQ5 (Adoption)**: How can open source, library-led AI tools be packaged so that institutions with limited staffing and budgets can adopt them without high technical overhead?

2. Project Justification

2.1 The Navigation Crisis and Its Impact on Student Success

Nearly one in four first-year students do not persist to their second year at any institution, and roughly one in three do not return to the institution where they started [8], [9]. Research on student departure consistently identifies **information gaps** (i.e., difficulty locating, interpreting, and acting on academic and financial policies) as a contributing factor [10], particularly for students with limited access to advising [11]. This problem is especially acute in rural "education deserts," where an estimated 3 million Americans lack a broad-access public institution within 25 miles [12], [13], and where academic advisor caseloads at two-year institutions reach a median of 441 students per advisor [14]. The resulting support deficit is particularly consequential for adult learners (aged 25 and older), who account for 24% of all undergraduates [15] and often must navigate complex return-to-study policies and financial aid rules not designed for non-traditional enrollment patterns [16].

A major pathway from navigation breakdowns to attrition is the disconnect between academic decisions and financial consequences. Dropping a course can reduce a student's enrollment intensity below full-time status, potentially reducing Pell Grant awards (up to $7,395 annually) [17] and other benefits such as Post-9/11 GI Bill funding [18]. These pressures are compounded by transfer friction. **Students lose an average of 43% of credits when changing institutions** [19], and those who retain less than half their credits are less than half as likely to graduate as those who retain 90% or more [20]. *Existing advising capacity cannot keep pace when advisors serve hundreds of students and have limited time to provide the coordinated academic and financial guidance these situations require*.

2.2 Who Is Most Affected

The navigation crisis affects students in different ways. First-generation students, who represent 57% of NIU's entering class [21], often lack informal guidance about how higher education systems work that other students may receive through family experience. Military-affiliated students, who make up nearly 30% of ODU's enrollment [22], must navigate GI Bill rules, transfer credit from military training, and accelerated course formats, and many report anxiety about benefit disruptions when shifting between course formats. Adult learners, who account for 24% of undergraduates nationally and are often returning to higher education after extended absences [23], [24], must interpret return-to-study policies and degree requirements that may have changed since their initial enrollment. This challenge grows when institutions cannot document when requirements changed and how those changes were applied. This is closely related to "***catalog rights***," the principle that students may complete degree requirements under the catalog in effect when they first enrolled. Rural and remote learners in education deserts may have limited access to in-person advising and rely on online information that is incomplete or out of date [13].

These factors often overlap. A single learner may be first-generation, rural, and returning after a long break, compounding the navigation burden. Local capacity constraints further intensify the problem. At NIU, the Computer Science department has one undergraduate advisor supporting more than 500 students, limiting the time available for coordinated academic and financial guidance. At UNT, with more than 46,000 students, many learners transfer from community colleges with credits that may not apply to their intended program, and the institution has set a goal of 90% first year retention by 2030 [25]. Progress toward that goal depends in part on better tools that help students understand program requirements, transfer pathways, and the financial consequences of course decisions. When students turn to generic AI tools for quick answers, they may receive advice that sounds plausible but does not match local policies, a well-documented phenomenon known as *AI hallucination*, in which generated content appears credible but is factually inaccurate [26], [27]. This can create misplaced confidence and lead to costly mistakes, particularly when students lack the domain knowledge to recognize errors.

2.3 The Organizational Memory Crisis

Alongside navigation losses is a rapid erosion of university knowledge. Large-scale retirements among baby boomer faculty are expected to accelerate over the coming decade [28], [29], creating a generational knowledge cliff. Faculty Senate minutes, curriculum committee records, course syllabi, departmental policy webpages, and related governance and policy documents capture the "**University DNA**" that explains why student-facing policies exist and how requirements change over time. Yet large-scale surveys of research libraries indicate that nearly half of archival collections have no online presence and 44% lack any online catalog record [30]. For governance and policy records such as Faculty Senate minutes and curriculum committee documents, which are rarely treated as high-priority special collections, discoverability is likely even lower. When records are incomplete or hard to locate, returning students have fewer tools to document catalog rights, AI systems have less context for interpreting policy intent, and universities lose the ability to trace and learn from prior decisions. This project examines a systematic approach to organizing existing University DNA to create a curated, faculty-reviewed record base that supports AI-assisted student advising. The archive primarily draws on publicly available governance records, course syllabi, and departmental policy pages that directly support student-facing policy and degree-requirement questions. **We hypothesize that the archive can function as a "restore point" for the advising system**. When the system produces weak or incorrect guidance, curated records provide verifiable sources that support correction and help keep recommendations aligned with documented policy history. This preservation also supports catalog rights by making requirement changes easier to trace and

explain, reducing the likelihood that returning students repeat coursework because policy changes were not well documented.

2.4 The Library Relevance Opportunity

As generative AI tools become widely used, academic libraries face renewed questions about their distinctive role on campus. Many routine reference interactions are increasingly handled through commercial AI systems. This shift also creates an opportunity: libraries are well-positioned to steward the authoritative university records and sources that AI systems require to deliver reliable, source-based guidance. Unlike commercial vendors, libraries have established responsibilities for curation, preservation, metadata practice, and access governance for university information resources [31], [32]. **This project examines whether libraries can pivot from answer finders to knowledge stewards and AI oversight roles**, including workflows for designating authorized sources, maintaining curated records, and routing uncertain cases for human review [33], [34]. Research on academic library service points suggests that many questions are directional, logistical, or otherwise transactional rather than complex research consultations [35], [36]. If routine queries can be handled with source citations and clear review pathways, librarians can focus more time on complex consultations, metadata work, and ongoing curation of authoritative sources.

2.5 How This Research Differs from Existing Approaches

Current approaches to AI-enhanced advising fall into two categories with important limitations. Retrieval Augmented Generation (RAG) [37] systems use search to retrieve documents and pass them to language models. These systems can answer factual questions from available documents but are limited when guidance requires reasoning across multiple policy areas, accounting for policy changes over time, or identifying conflicts among sources [38], [39]. Enterprise AI platforms (e.g., Salesforce Education Cloud [40] with Einstein add-ons, and Ellucian's AI-embedded platforms [41]) can integrate data from multiple systems, but they are typically proprietary, provide limited transparency into how recommendations are produced, and require licensing costs and technical integration that many institutions cannot sustain [42]. This project evaluates ATLAS as an integrated system with three interdependent components. **First, the ATLAS multi-agent architecture** examines whether a team of specialized agents can coordinate academic requirements, financial aid rules, and workforce information to address common advising scenarios. **Second, University DNA grounding** examines whether anchoring the system in a library curated archive of governance and policy records improves factual accuracy and produces verifiable reasoning traces that make the basis for each recommendation reviewable. **Third, library-led oversight workflows** examine how libraries can govern the system in practice by designating authorized sources, resolving conflicts across records, and routing low-confidence cases for human review. Together, these components form a single pipeline in which the archive provides the evidence base, the agents synthesize guidance from that evidence, and library oversight maintains quality over time. Evaluating the pipeline end-to-end, rather than as isolated parts, allows the study to test whether combining grounded evidence, agent coordination, and governance yields more reliable advising support than any single component on its own. Table 1 compares the current approach limitations with the ATLAS design

Table 1. Summarizes critical gaps in current systems and how ATLAS addresses them.

Current Limitation	Gap in Existing Systems	How Our Project Fills the Gap
Advisors are overburdened	Advisors manage hundreds of students, leaving limited time for personalized, cross-domain guidance.	Library-led AI provides 24/7 guidance as a "paraprofessional library assistant," reducing routine question volume so advisors can focus on complex cases.
Generic AI lacks campus-specific context	Off-the-shelf LLMs (ChatGPT, Gemini) provide generalized responses that may not match local policies, prerequisites, or financial aid rules.	Specialized agents use secure connectors to retrieve and reason across campus catalogs, financial aid rules, and workforce data.

Proprietary AI is costly and hard to audit	Enterprise platforms require licensing and provide limited visibility into how recommendations are produced.	Open source system with source cited recommendations that show the sources consulted and a clear explanation trail.
No long-term university memory	Universities often lack reliable, searchable archives of past governance and policy records, making policy changes difficult to trace.	The library-curated University DNA Archive preserves key governance and policy records and supports “restore points” when guidance needs correction.
AI outputs are difficult to verify	Students and advisors have difficulty evaluating AI guidance when sources and reasoning are not shown.	Each response includes a citation trail with links to designated sources, timestamps, and a review pathway when certainty is low.
Cross-domain reasoning is limited	Current tools answer questions within one area but do not connect academic choices to financial aid or workforce implications.	Specialized agents coordinate across academic, financial aid, and workforce information to support end-to-end student scenarios.

3. Project Work Plan

3.1 Theoretical Framework

This research is grounded in two complementary frameworks. The Technology Acceptance Model (TAM) [43], [44] posits that adoption of new tools is driven by perceived usefulness (does it help me achieve my goals) and perceived ease of use (can I use it without excessive effort). We hypothesize that AI systems grounded in university records and designated sources will score higher on perceived usefulness because users can verify recommendations against visible citations and linked documents. TAM also suggests that adoption can decline if a system is difficult to interpret. This study will examine the tradeoff between transparency and usability, including whether citation trails and explanations support verification without adding undue effort. The Information Seeking in Context (ISC) framework [45], [46] emphasizes that information seeking is situated and that people bridge gaps in understanding as they act toward goals. First-generation students, military-affiliated students, and adult learners may face larger sense-making gaps because they have less prior exposure to how university policies, requirements, and offices interact [47], [48], [49]. Traditional advising systems often assume students already know what questions to ask; proactive advising research suggests that effective support should instead surface relevant issues before students recognize them on their own [50], [51]. ISC suggests that effective support should also help users recognize missing information and anticipate downstream consequences. We hypothesize that specialized AI agents that proactively surface connections, such as linking a course drop decision to financial aid implications, will support sense-making better than tools that respond only to narrowly framed questions. We will examine whether library-led, source-cited advising support can function as a cognitive scaffold that helps students identify what to check, where to verify it, and when to seek human review.

3.2 System Architecture

ATLAS utilizes existing archives from the participating organizations: NIU [52], ODU [53], and UNT [54]. It uses a modular architecture coordinated by an **Orchestrator** that decomposes user questions and routes subtasks to four specialized AI agents [55]. (1) The **Heritage Agent** retrieves governance and policy records from a library archive using semantic search to support interpretation of policy changes over time. (2) The **Policy Agent** retrieves current catalogs and cross-checks policies across campus offices (Registrar, Financial Aid, Bursar), flagging disagreements among sources. (3) The **Prosperity Agent** integrates Bureau of Labor Statistics workforce data with degree curricula to support career-aligned planning, including pathways to industry certifications and registered apprenticeships, consistent with Executive Order 14278, Preparing Americans for High-Paying Skilled Trade Jobs of the Future [56], [57]. (4) The **Verification Agent** checks outputs from the other agents against Authorized Sources designated by library staff, assigns a ***certainty rating***, and routes low certainty cases to human review through established librarian and advisor workflows.

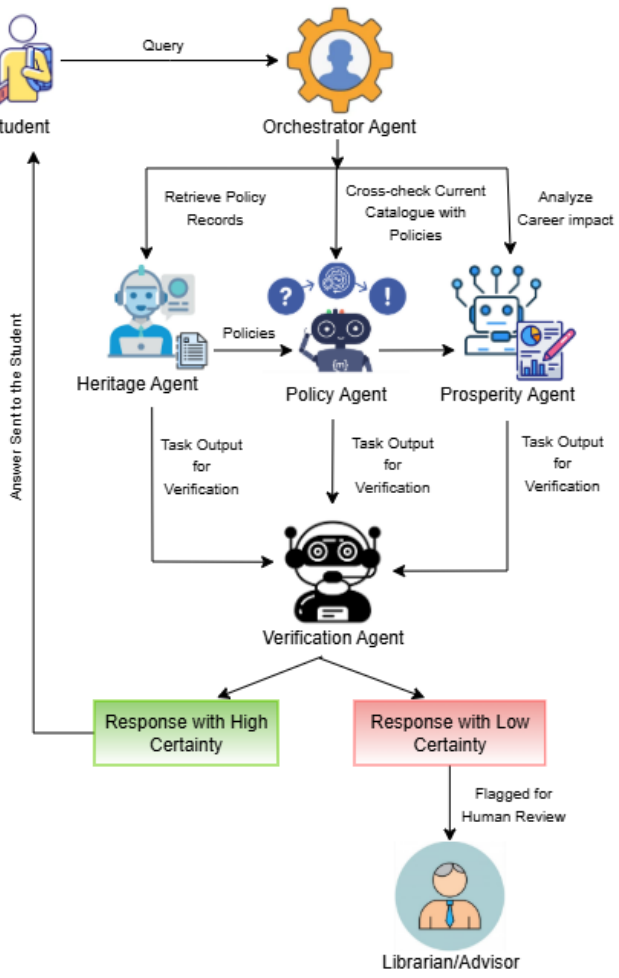

Figure 1. ATLAS System Architecture. The Orchestrator decomposes student queries and routes subtasks to four specialized agents. The Verification Agent checks all outputs against Authorized Sources and routes low certainty cases for human review.

How the Agents Collaborate. When a student submits a query (e.g., "I'm struggling in Calculus II and I am considering dropping it. What happens next?"), the Orchestrator breaks the request into subtasks. The Policy Agent retrieves the university drop and add policy and grade replacement rules. The Heritage Agent checks whether these rules have changed since the student first enrolled, supporting catalog rights. The Prosperity Agent checks whether Calculus II is required for the student's intended career pathway and identifies relevant alternatives. The Verification Agent cross-checks each finding against Authorized Sources, assigns an overall certainty rating, and assembles a single response with citations. **When the certainty rating is low for any component, the system flags that portion for human review and routes it to a liaison librarian or academic advisor**. This workflow mirrors how an advising team approaches complex cases, while providing consistent citations and review pathways at scale. *Please see Supporting Document 4 for detailed scenario-based demonstrations of this workflow, and Supporting Document 6 for the full agent interaction pipeline.*

The system uses **Model Context Protocol (MCP)** [58], an open standard for building reusable connectors between AI agents and campus data sources such as library catalog systems, registrar systems, financial aid systems, and labor market databases [59]. MCP serves as a common interface that allows institutions to connect existing systems to the ATLAS architecture without redesigning the underlying infrastructure. The project will develop and document connectors for four widely used platforms: Ex Libris Alma (library services) [60], OCLC WorldCat Discovery (catalog search) [61], Ellucian Banner (student information) [62], and Ellucian Colleague (student information). These connectors will be released openly with installation and configuration guides to support adoption by other institutions with limited technical effort. The connector approach will be modular so that additional libraries and student information platforms can be supported through future work or contributions from adopting institutions. *Please see Supporting Document 6 for detailed MCP connector specifications.* Each

AI-generated recommendation includes a transparent record showing which documents were consulted (with permalinks to the library archive), the reasoning applied, a certainty rating, and the date each source was last checked. We call this a **Verifiable Reasoning Trace**, a structured citation trail that allows students, advisors, and librarians to see how the system arrived at its answer and to review the supporting records directly. This transparency is the core governance feature of ATLAS, enabling students to practice evaluating information by reviewing the sources and logic behind each recommendation [63], [64], [65].

Privacy and Security. The system operates on university policies and public records, not student records. No personally identifiable student information is used to train models or generate recommendations [66], [67], [68]. Interaction logs are collected in de-identified form for evaluation and auditing. Routine security testing using simulated misuse scenarios will identify vulnerabilities such as prompt injection [69], [70], [71], [72]. Security controls, including authorized source allowlists, injection pattern detection, and confidence gating, will be documented and released as a reusable audit guide [73]. Logs associated with student queries are handled as education records, protected using encryption, and retained in accordance with FERPA requirements [74]. *See Supporting Document 7 for the full privacy-by-design framework and red-teaming protocol.*

3.3 Library Practice Transformation

A significant component of this research examines how library-led AI changes professional library work, studying the shift from traditional reference services toward knowledge stewardship [75]. The project team will manage Authorized Sources designated by responsible campus offices, with library Co-PIs advising on source authority and cataloging standards. When sources disagree, ATLAS applies a documented source order prioritizing official catalog and policy materials. When the system identifies low certainty or conflicting information, it routes the case for human review, first to a liaison librarian and then to an academic advisor. We will document how these workflows operate in practice, the training librarians need, and whether this governance role strengthens the library's visibility among campus partners. This evidence is essential for replication: if libraries are to guide AI-enabled advising at scale, the field needs documentation of effective workflows, staffing needs, and integration with student support functions.

3.4 Research Design

We employ **Design-Based Research (DBR)** [76], [77], an iterative methodology for educational technology research that emphasizes addressing authentic problems in real contexts, collaborating with practitioners, integrating design and research cycles, and producing both scholarly contributions and practical interventions [78], [79]. The research will proceed through three overlapping phases across a 24-month performance period.

Phase 1: **System Development and Refinement (Months 1–9)**. The team will develop the ATLAS architecture and campus connectors and conduct Archive Mapping workshops at each campus. Archive Mapping is a structured process in which the project team, working with librarians, archivists, and academic advisors, identifies common high-impact advising questions, maps those questions to the specific university records that answer them, and assesses which materials to ingest first using a shared appraisal rubric. Each campus will create a University DNA Archive focused on high-value governance and policy records from the past 50 years, including Faculty Senate minutes, curriculum committee records, course syllabi, and departmental policy webpages. The archive is scoped to records that are already digitized, publicly posted, or maintained by campus units. Where additional content preparation is needed, the team will handle scanning, metadata creation, and accessibility formatting to the extent feasible within the defined project scope. All records access will comply with applicable state records management requirements and institutional policies governing archival transfers, copyright, and privacy. The appraisal rubric (developed in Months 2–3) will evaluate candidate records on policy impact, temporal significance, discoverability risk, and AI utility, and will be released as Deliverable D1. The team will also develop and pilot test research instruments, obtain IRB approvals, and conduct preliminary usability testing (n=10 per campus) [80] to refine the system before full deployment.

Phase 2: Controlled Comparative Study (Months 8–18). We will implement a controlled study with three conditions: (a) a Control Group using the current campus website search and human advisor support; (b) a RAG Comparison Group using a keyword search based RAG system [81] with current catalog data only; and (c) an ATLAS Group using the full multi agent system with the University DNA archive and Verifiable Reasoning Traces. Each condition will include 100 students across the partnership (total n=300). Recruitment will target first-

generation (40%), rural (25%), military-affiliated (20%), and adult returner (15%) learners. We will recruit an additional 10% to account for anticipated withdrawal. Students will be randomly assigned using block randomization [82], [83].

Table 2. Three-Condition Controlled Comparative Study Design.

Dimension	Control Group (n=100)	RAG Comparison (n=100)	ATLAS Group (n=100)
Information Access	Campus website search and human advisor	Keyword search RAG chatbot with current catalog data only	Full multi-agent system with University DNA archive
Source Grounding	Whatever the student finds independently	Current catalog and policy pages only	Current catalog plus historical governance and policy records
Transparency	No structured citation trail; student navigates alone	Basic retrieval results shown	Verifiable Reasoning Traces with source citations, timestamps, and certainty ratings
Human Review	Standard advisor appointment (if student seeks one)	None	Low-certainty cases flagged and routed to a librarian or advisor
RQ Addressed	Baseline for RQ1, RQ2	RQ1 (isolates agent effect), RQ2 (citation without archive)	RQ1–RQ4 (full system with archive and verification)
Measures	Task accuracy (0–10), completeness, source verification, TAM surveys, Trust in Automation, Information Literacy Self-Efficacy, semi-structured interviews (n=90)		

Students will complete **three scenario-based decision-making tasks** designed to test multi-domain reasoning relevant to common student navigation challenges. The three tasks are: (1) **Transfer Credit Evaluation**, requiring students to determine credit applicability, prerequisite satisfaction, and time-to-degree impact from a community college transcript [84], [85], [86] ; (2) **Course Drop Impact Analysis**, requiring assessment of academic, financial aid, and progress consequences of dropping a course [87], [88]; and (3) **Catalog Rights Protection**, requiring identification of which degree requirements apply after a multi-year absence [89]. *See Supporting Document 4 for detailed scenario descriptions and system comparison tables.* These tasks represent high-impact navigation scenarios requiring cross-domain reasoning. Each task will be administered identically across all three conditions. Responses will be scored independently by two trained raters using a rubric covering accuracy, completeness, and source verification, with disagreements resolved by a third rater [90]. Task materials and scoring rubrics will be included in the replication package. Data collection will include pre- and post-surveys, task performance data, system interaction logs, and post-task interviews with a subset of participants (n=30 per campus, 90 total).

Survey Instruments. Measures will include the Perceived Usefulness and Perceived Ease of Use scales [91], the Trust in Automation Scale [92] adapted for AI advising contexts, and an Information Literacy Self-Efficacy scale [93]. Demographic surveys will capture first-generation status, rural zip codes, military affiliation, and prior academic experience. All instruments will be pilot tested during Phase 1.

***Phase 3*: Field Deployment, Analysis, and Dissemination (Months 17–24)**. We will deploy ATLAS campus-wide on an opt-in basis at the three partner institutions, tracking outcomes including retention, credit accumulation, time to degree progress, and user satisfaction using an observational comparison design [94]. We will study how source designation, record maintenance, and human review routing can be integrated into sustainable library and advising practice. We will finalize the open source Library Agent Toolkit (D6) through collaborative design workshops with librarians, advisors, and student representatives [95], and conduct national dissemination through webinars, conference presentations, and a national convening of library practitioners [96].

Qualitative Methods. Semi-structured interviews (n=90 total, 30 per campus) will explore how students interpreted and acted on ATLAS recommendations, what shaped confidence in the guidance, whether Verifiable Reasoning Traces supported understanding or added burden, and how the experience compared with working with human advisors. Interviews will be analyzed using thematic coding [97], [98]. Library staff studies will include

observation during human review cases, workflow analysis, and focus groups on implementation challenges [99], [100].

Data Analysis Plan. Quantitative analysis will compare task performance across the three study conditions and participant subgroups using multivariable models to estimate differences in outcomes by condition [101], [102]. Additional models will examine whether interaction with Verifiable Reasoning Traces is associated with information literacy outcomes and user confidence. We will compare the frequency of incorrect or incomplete recommendations between conditions using a review rubric and independent human review. We will also analyze failure cases by categorizing question types and source conditions (e.g., conflicting records, outdated pages) to identify patterns that inform responsible use guidance. Qualitative analysis will use thematic coding with both theory-informed and emergent codes, cross-case analysis, and mixed methods integration [103], [104].

IRB Plan and Recruitment. The project team will submit protocols to the IRBs at NIU, ODU, and UNT in Month 1. The study is expected to qualify for a streamlined review pathway [105], [106]. NIU will serve as the primary reviewing board, with reliance arrangements to support a coordinated multi-site process. Protections include informed consent, secure data handling, FERPA-compliant de-identification, and the right to withdraw without penalty [107]. Recruitment will be conducted through academic advising and student support networks at each institution, with outreach targeting the project's priority learner groups.

3.5 Project Team and Stakeholder Engagement

Principal Investigator Hamed Alhoori (NIU) is an Associate Professor in Computer Science, specializing in digital libraries and AI-driven data analytics, and serves on the board of the International Journal on Digital Libraries [108]. His published work on scholarly information systems, recommendation methods, and large-scale data mining provides the foundation for developing the ATLAS system and the supporting archive. He served as the General Chair of the ACM-IEEE Joint Conference on Digital Libraries (JCDL) in 2025 [109], demonstrating leadership in the digital library field. *Co-PI Sagnik Ray Choudhury (UNT)* specializes in natural language processing with a focus on explainable and interpretable AI methods. His research on explanations in text models and on evaluating language model generalization is directly relevant to the Verifiable Reasoning Trace specification and to the design of outputs that can be reviewed with cited sources. His foundational work with CiteSeerX at Penn State demonstrates experience in building large-scale digital library information extraction systems. He served as the Program Chair for JCDL 2025. *Co-PI Faryaneh Poursardar (ODU)* is a researcher in web archiving, digital preservation, and AI model customization. Her work on digital collection change detection and web resource boundary identification supports the archival workflows that underpin the University DNA Archive. Her expertise in analyzing user-generated content will inform the design of student-facing interactions. *Co-PIs Fred Barnhart, Kimberly Shotick, and Shupei Yuan (NIU)* provide expertise in academic librarianship, student success, and strategic communication. Barnhart, as Dean of University Libraries at NIU, brings leadership and commitment to sustaining the archive beyond the grant period. Shotick's research on the connection between library use and student academic self-efficacy informs the project's approach to building accessible tools and services. Yuan's research on message credibility, information use, and decision making will guide the system's communication approach so that guidance is clear and grounded in cited sources.

Project personnel include three Graduate Research Assistants (one per campus). The PI will coordinate cross-campus workflows, with library Co-PIs advising on metadata standards and source authority. Technical development will be led by the computer science research team in coordination with the library and advising partners. An Advisory Board comprising registrar representatives, financial aid directors, academic advisors, university archivists, student representatives from priority learner groups, and an external AI ethics expert will meet quarterly. Stakeholder engagement is integrated throughout the work plan, including student focus groups for interface co design (Months 2–3), Advisory Board meetings (Months 6, 12, 18), a mid project workshop with community college partners (Kishwaukee, Patrick and Henry, and Del Mar) in Month 10, a national convening of library practitioners in Month 20, and a final validation workshop in Month 22. Community college partnerships are included because transfer students are among those most affected by navigation challenges. The national convening will bring together 30 to 40 library practitioners from varied campus types to review findings, provide feedback on the toolkit, and begin forming a National Library Agent Registry community of practice. To address RQ5, the convening will include a structured evaluation session in which practitioners from non-partner institutions assess the toolkit's documentation, installation process, and training materials. Participants will be recruited through professional networks, including ACRL, ALA, and CNI, as well as through referrals from

Advisory Board members and community college partners, with emphasis on institutions with limited IT staffing. Participants will complete a structured assessment covering perceived ease of installation, clarity of documentation, estimated staffing requirements, and technical readiness. Feedback will directly inform final revisions to the toolkit (D6) before its public release in Month 22. An independent evaluation consultant will conduct a mid-point evaluation (Month 12) and a final evaluation report (Month 23) to assess progress toward project goals, review study recruitment relative to participant representation targets, and recommend priorities for follow-on research and dissemination.

Institutional Contexts Supporting Applicability. The three partners offer complementary contexts: NIU (R2 public, rural serving, 54% first generation, 16,000 students in DeKalb, Illinois, a rural community 65 miles from Chicago), ODU (R1 public, 30% military affiliated, 23,000 students in Norfolk, Virginia, with accelerated 8 week schedules and extensive online programs), and UNT (R1 public, HSI with 27% Latino/a enrollment, 46,000 students in Denton, Texas, with a first year retention goal of 90% by 2030). This mix supports findings that can inform implementation across rural-serving institutions, military-serving campuses, large research universities, and HSI settings, reflecting a wide range of U.S. higher education environments. **Broad Participation Plan**. Recruitment will prioritize first-generation students (40% of participants), rural and remote learners (25%), military-affiliated students (20%), and adult returners (15%). The Advisory Board will include student representatives from these groups. The open source toolkit will be packaged with tiered implementation guidance for institutions of different sizes and staffing levels. Analyses will report results by participant subgroup to document how outcomes vary across groups.

4. Project Results

Intended Deliverables

The Digital Heritage Exhibit (D7) directly addresses the IMLS emphasis on America's 250th anniversary, documenting how democratic deliberation in Faculty Senates shaped educational access over half a century (see Supporting Document 5). The project will produce eight deliverables within the 24-month period:

- **D1 (Month 12):** University DNA Collection Policy Template, including an appraisal rubric and metadata schema, released under a CC BY license.
- **D2 (Month 14):** Prosperity Metadata Profile linking Bureau of Labor Statistics (BLS) workforce data with academic curricula, submitted to the Dublin Core Metadata Initiative for community review.
- **D3 (Month 18):** Open source connectors for widely used platforms (Ex Libris Alma, OCLC WorldCat Discovery, Ellucian Banner, and Colleague), released on GitHub with installation and configuration guides.
- **D4 (Month 18):** Verifiable Reasoning Trace Metadata Standard defining a structured citation format that records sources consulted, reasoning applied, confidence scores, and source verification timestamps.
- **D5 (Month 20):** Human in the Loop Audit Protocol and Training Curriculum for librarians and advisors, shared through professional channels including the Association of College and Research Libraries (ACRL) and the National Academic Advising Association (NACADA).
- **D6 (Month 22):** Library Agent Toolkit, an open source, tiered implementation package including source code, configuration guides, and hardware recommendations for institutions of different sizes.
- **D7 (Month 20):** Digital Heritage Exhibit, "Evolution of Academic Opportunity: 50 Years of Civic Life," documenting how shared governance shaped educational access in commemoration of America's 250th anniversary.
- **D8 (Month 24):** Library AI Governance Model, a practical blueprint for libraries as stewards of AI-supported university knowledge.

Library Agent Toolkit (D6). The toolkit will include the complete system source code; installation and configuration guidance for three tiers of capacity (Tier 1: small colleges using hosted cloud services; Tier 2: regional universities using on premise hardware; Tier 3: large research universities with full IT infrastructure); sample connector configurations for each supported platform; workflow templates for designating Authorized Sources, resolving conflicts between sources, and managing escalation workflows; and scenario based training materials for librarians and academic advisors. This tiered approach supports adoption across institutions of different sizes and staffing levels, including campuses with limited technology budgets.

Digital Heritage Exhibit (D7). The exhibit will highlight pivotal moments in the evolution of academic opportunity through archived governance records, including the expansion of vocational programs in the 1970s following the GI Bill's impact, the development of student degree requirement protections in the 1980s, the growth of transfer pathways in the 2000s, and contemporary efforts to support first-generation, military-affiliated, and adult learner populations. The exhibit will be hosted on the NIU library website and promoted through partnerships with the DeKalb Public Library, reaching both academic and community audiences.

Expected Research Findings. The study is designed to test three primary hypotheses. First, we hypothesize that ATLAS will improve students' performance on complex navigation tasks compared with generic AI tools and standard campus search, with estimated gains of 25 to 40% in task accuracy and the largest improvements among first-generation students and adult returners. Second, we hypothesize that Verifiable Reasoning Traces will improve users' confidence in evaluating guidance and strengthen information literacy outcomes. Third, we hypothesize that adding historical governance records will reduce inaccurate guidance relative to systems grounded only in current catalog content. The study is equally designed to identify conditions under which the system performs poorly, and failure cases will be treated as a primary research output. We anticipate that analysis of system limits will identify recurring cases where human review is needed, including conflicting guidance across sources, questions depending on real-time information, ambiguous policy language, and undocumented advisor practices. These results will be synthesized into a Limitations and Responsible Use framework. In addition to system outcomes, we expect findings on library practice transformation, including the competencies and workflows involved in library-led governance of AI-supported guidance, and whether this governance role changes how campus stakeholders perceive the library's contribution to student success.

Dissemination and National Impact. Academic dissemination will include peer-reviewed journal submissions and conference presentations spanning digital libraries, NLP, higher education technology, and advising practice. Target journals include College & Research Libraries, Journal of Academic Librarianship, JASIST, and the International Journal of Artificial Intelligence in Education. Planned venues include ALA Annual, ACRL, EDUCAUSE, CNI, NACADA, JCDL, LAK, and computational linguistics conferences such as ACL and EMNLP. The PI (Alhoori) served as General Chair of JCDL 2025, and the Co-PI, Choudhury (UNT), served as Program Chair; both Alhoori and Choudhury are serving on the JCDL 2026 organizing committee [110]. Leveraging this sustained leadership in the premier international digital library venue, the team will organize a half-day workshop at JCDL 2026 to demonstrate the ATLAS tools and engage librarians, archivists, and information scientists in hands-on exploration of AI-augmented academic advising using digitized university governance records. This workshop will require no additional funding, as it will utilize the existing conference infrastructure. Beyond JCDL, the three-institution ATLAS team will host virtual workshops open to the broader library and archives community, using institutional videoconferencing platforms already available at NIU, ODU, and UNT. These online sessions will target practitioners who may not attend research conferences, including academic advisors, registrars, and institutional researchers at community colleges and minority-serving institutions. Additionally, each partner campus will organize a local on-campus event to engage its own library staff, advisors, and students in pilot testing and feedback sessions. We anticipate 6 to 8 conference presentations, complemented by practitioner articles in American Libraries and Library Journal. The Digital Heritage Exhibit will be promoted through campus and public library partnerships, including the DeKalb Public Library. The project advances a replicable model for library-led student success support in an AI-shaped environment. Open-source software, documentation, and training materials will support adoption. A National Library Agent Registry will sustain a community of practice among adopting institutions.

America's 250th Anniversary Alignment. This project contributes to the America250 commemoration by preserving governance records that document how academic policies evolved, presenting that history through a Digital Heritage Exhibit (D7) hosted on the NIU library website with DeKalb Public Library partnerships, and hosting public workshops connecting archival AI with civic engagement themes. *See Supporting Document 5 for the full Commemoration Plan*.

Sustainability. Technical sustainability is supported through permissive open source licensing, modular design, and reusable connectors. NIU will host the shared technical infrastructure for at least five years after the grant ends. The toolkit runs on commodity hardware or cloud services at a modest annual cost with no licensing fees. A National Library Agent Registry will serve as a shared hub for documentation, connector updates, and implementation notes. Training materials will be shared through established professional channels, and the team will pursue follow-on funding to support continued maintenance beyond the initial award.

SCHEDULE OF COMPLETION

Lead Institution Responsibilities

Institution	Primary Responsibilities
NIU (Lead)	Overall project management; system development and software engineering; digital archiving; heritage exhibit; toolkit release; hosting infrastructure
ODU	University DNA collection guidance and archival standards; military-affiliated student recruitment; web archiving expertise; MCP connector testing
UNT	Multi-agent AI architecture; NLP and explainability research; labor market data integration; reasoning trace specification

Year 1: September 2026 – August 2027 (Months 1–12)

Activity	Lead	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12
1.1 Submit IRB protocols (all sites)	NIU	X	X										
1.2 Develop multi-agent architecture	UNT, NIU	X	X	X	X	X	X	X	X	X			
1.3 Develop MCP connectors	NIU, UNT			X	X	X	X	X	X	X			
1.4 Archive Mapping workshops	All		X	X	X								
1.5 Prepare & ingest University DNA	All			X	X	X	X	X	X	X	X	X	X
1.6 Develop research instruments	NIU				X	X	X						
1.7 Student focus groups (UI co-design)	All		X	X									
1.8 Advisory Board Meeting #1	NIU						X						
1.9 Usability testing (n=10/campus)	All							X	X				
1.10 Begin comparative study recruitment	All								X	X	X		
1.11 Comparative study begins	All									X	X	X	X
1.12 Pre-surveys and scenario tasks	All									X			X
1.13 Community college partner workshop	NIU										X		
1.14 Publish Univ. DNA Policy (D1)	ODU												X
1.15 Mid-point external evaluation	Ext.											X	X
1.16 Advisory Board Meeting #2	NIU												X
1.17 Quarterly milestone reviews	NIU			X			X			X			X

Year 2: September 2027 – August 2028 (Months 13–24)

Activity	Lead	M13	M14	M15	M16	M17	M18	M19	M20	M21	M22	M23	M24
2.1 Comparative study continues	All	X	X	X	X	X	X						
2.2 Post-task interviews (n=30/campus)	All	X	X	X									
2.3 Publish Prosperity Metadata (D2)	ODU		X										
2.4 Field deployment (campus-wide opt-in)	All					X	X	X	X	X	X		
2.5 Library staff workflow studies	All					X	X	X	X				
2.6 Release MCP Connectors (D3) on GitHub	NIU						X						
2.7 Release Reasoning Trace Spec (D4)	UNT						X						
2.8 Advisory Board Meeting #3	NIU						X						
2.9 Data analysis (quant. & qual.)	All				X	X	X	X	X	X			
2.10 Release HITL Audit Protocol (D5)	UNT								X				
2.11 Launch Digital Heritage Exhibit (D7)	NIU								X				
2.12 National convening of practitioners	NIU								X				
2.13 Manuscript preparation & submissions	All						X	X	X	X	X	X	
2.14 Final validation workshop	All										X		
2.15 Release Library Agent Toolkit (D6)	NIU										X		
2.16 Final external evaluation report	Ext.											X	
2.17 Publish Governance Model (D8)	NIU												X
2.18 Final project reporting to IMLS	NIU												X
2.19 Quarterly milestone reviews	NIU			X			X			X			X