



NATIONAL LEADERSHIP GRANTS FOR LIBRARIES

Sample Application LG-260361-OLS-26
Project Category: Planning

University of Alabama, School of Library and Information Studies

Amount awarded by IMLS:	\$199,994
Amount of cost share:	\$0

The University of Alabama, School of Library and Information Studies will design and pilot a Library AI Ethics and Literacy Network on Social Media to help libraries extend AI education into the online spaces where the public is increasingly encountering AI. The project studies barriers libraries face in reaching the public online, surveys librarians and community members, builds a prototype content hub, and tests a cross-library social media network with pilot libraries. Deliverables include planning tools, pilot-tested workflows, and a rollout framework. Expected outcomes are practical, low burden approaches for library-led AI ethics education on social media and a field-ready framework that, if validated, could support libraries nationwide, especially small, rural, and under-resourced libraries and the communities they serve. This project will assess librarian and public needs, pilot a structured social media presence, develop and test a shared content hub prototype, and create a scale-up plan for the project.

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.

Amplifying Library Voices in AI Ethics and Literacy Through a National Social Media Network Design

1 Introduction

The University of Alabama, through the School of Library and Information Studies in the College of Communication & Information Sciences, requests \$199,994 in IMLS funding for a 2-year planning grant to design and pilot a coordinated, library-led network for AI ethics and literacy communication on social media. The project strengthens libraries' visibility and credibility in public AI discourse by reducing the burden of content creation and supporting joint outreach. Libraries are trusted community institutions and already contribute to AI education through programs and training, but much of that work remains local and event-based, limiting its reach. Although library perspectives do appear in these social media spaces, they remain fragmented, limited in reach, and rarely originate high-visibility content. This project will (1) assess librarian and public needs, (2) pilot a structured social media presence, (3) develop and test a shared content hub prototype, and (4) create a scale-up plan for a future Library AI Ethics & Literacy Network on Social Media. *Confirmed collaborating organizations and advisors* include the Howard County Library System, Birmingham Public Library, Harvard Library, and EBSCO Information Services. Additional library-association partnerships will be developed during the project period. PI Li's leadership in the Association for Information Science and Technology (ASIS&T) (currently Chair-Elect of its Special Interest Group (SIG) on AI Governance) provides a strong national professional network for practitioner engagement and dissemination. The project aligns with **NLG-L Objective 1.1**, serving the public's learning needs in AI, digital, and media literacies. It also supports **Objective 1.3** through expanded online public access to library-produced AI educational resources and addresses the OLS priority theme of Technology & AI (**EO 14277**). PI Yuan Li, Ph.D. (Human–AI Interaction), will lead the project with Co-PIs Bharat Mehra, Ph.D. (EBSCO Endowed Chair in Social Justice), and Jack Lipei Tang, Ph.D. (Advertising & Public Relations).

2 Project Justification

2.1 The Need: Libraries' Trusted Role as Educators and the AI Learning Gap in Online Spaces

Artificial intelligence is entering everyday life faster than many people are being prepared to understand it. People increasingly encounter AI and begin forming views about it through online environments, especially social media, where information circulates quickly with little context [1–3]. Although these platforms are accessible and immediate, they are poorly suited to careful explanation, critical evaluation, and sustained education. As a result, public understanding of AI is often fragmented and uneven, with few sources of clear, consistent, and trustworthy guidance.

The gap extends beyond technical instruction alone. Public understanding of AI now also involves questions of fairness, bias, privacy, transparency, accountability, appropriate use, labor, and the limits of automated systems [2, 4–6]. People are increasingly expected to interpret both technical claims and social consequences in fast-moving information environments. National and international discussions have begun to recognize this broader need through growing calls for inclusive public AI education, workforce AI literacy, and community-based support [7]. However, relevant guidance remains dispersed across policy documents, school initiatives, library workshops, nonprofit materials, news coverage, and scattered online posts, leaving too little sustained, accessible, and visible public-interest guidance in the spaces where many people already encounter AI.

Libraries can fill this gap. They are among the most trusted institutions in the country for information, learning, and community guidance [8–12]. Their mission aligns closely with the needs of this moment. Libraries already contribute to AI ethics and literacy through programming, training, community education, and professional guidance [13–20]. Traditional high-quality library resources exist, much of this work remains local, event-based, tied to particular institutions, or directed toward professional audiences [21–24]. Social media changes the conditions of public contact by concentrating attention where questions, confusion, curiosity,

and concern emerge in real time [1]. By making trusted library guidance visible on these online platforms, libraries can help people more easily, quickly, and reliably improve their understanding of AI ethics and literacy.

Small, rural, and otherwise under-resourced libraries face the challenge most acutely. These institutions may have strong local trust and direct knowledge of community needs, but they often lack the staff time, communication workflows, and social media infrastructure needed for sustained public AI education in the online spaces where their communities are learning about AI [21, 25, 26]. Existing research shows that library social media activity is often centered on events, announcements, and general outreach, with limited staff time for strategic content development [21–24]. Studies also point to infrequent posting patterns, information anxiety, and burnout among librarians responsible for social media work [25, 26]. These conditions help explain why trusted library institutions still play a limited role in sustained online interpretation of emerging technology issues.

The project also builds on the PI's prior computational research on library engagement with generative AI discourse on social media. That work found that library accounts were present but limited in reach, weakly connected across institutions and regions, and often amplified through external commentators rather than through original library-authored interpretation [27, 28]. The findings suggest that library expertise already exists in online AI discourse, but it does not yet function as a visible, connected, and sustained educational voice. (see Figure 1).




Figure 1: Repost network visualization from AI-related library discourse on X/Twitter [27]. Red nodes represent the top 10 viral content creators, non-library users whose posts achieved mass amplification. Dark blue nodes represent the most active library professionals and library organizations. Light blue nodes are a sample of other reposting participants. The hub-and-spoke pattern shows that high-reach AI content originates almost entirely from outside the library profession: library-affiliated accounts appear as amplifiers clustered around external hubs rather than as originators of visible discourse. Library nodes are also sparsely connected to one another, forming small, locally oriented clusters rather than a cohesive cross-institutional network, illustrating why library expertise, though present, has not yet coalesced into a coordinated public-facing educational voice on AI.

2.2 Relationship to Existing Models, Standards, Scholarship, and Practice

IMLS's FY2025–26 investment of over \$4.1 million in AI literacy supported eight projects, but that funding focused primarily on librarian training and K–12 programming. A review of the IMLS Awarded Grants Search database (FY2015–2025) likewise suggests that social media strategy and generative AI have received comparatively limited attention within this portfolio. Professional competency frameworks have strengthened internal readiness by clarifying expectations around AI knowledge, ethics, and institutional responsibility [13, 16]. Our planning grant complements that foundation by shifting attention from internal preparation and local programming to public AI education in the online spaces where many people now learn about AI.

Prior research shows that library social media activity is often concentrated on announcements, events, and outreach rather than sustained public interpretation, in part because most libraries have limited staff time and limited strategic content capacity [21, 22, 24]. WebJunction's 2024 survey found that most public libraries spend fewer than five hours per week on social media [21]. Recognizing these constraints, we ask under what conditions a shared, lower-burden model can make participation more feasible, especially for small, rural, and otherwise under-resourced libraries. Rather than asking each participating library to maintain its own stream of original AI content, the project tests a common workflow in which libraries contribute local materials, review adapted content, reuse common resources, and participate in cross-library amplification. The workflow is designed for realistic staff capacity and treats each participating library as a node in a distributed AI education network rather than as a stand-alone educational social media operation.

Several established public communication models also inform the project's design. The CDC uses coordinated social media messaging to deliver public health education where health misinformation circulates [29]. The USDA Cooperative Extension Service has delivered university-based knowledge to communities through local educators for over a century, a model of distributed public education that directly informs this project's design [30]. The News Literacy Project's RumorGuard places media-literacy educational content directly in the information streams where misinformation spreads [31]. These models demonstrate the value of coordinated messaging, distributed participation, and in-stream public education. This project adapts those ideas to the library-AI context by testing how libraries can organize local expertise, review, attribution, and pooled content workflows into an organized public-facing model.

The project is also grounded in the complementary expertise of the project team. **PI Li** brings expertise in human–AI interaction, computational social media analysis, and AI-related public discourse. In a computational analysis of AI- and library-related posts on X/Twitter, she found that library accounts produced fewer than 15% of high-reach posts, and that most circulation came through retweets by external technology commentators rather than original library-authored interpretation [27]. Active library accounts were often locally oriented, loosely connected across institutions, and frequently lacked infrastructure support for sustained content production [27, 28]. **Co-PI Mehra's** long-standing work on social justice, community-centered librarianship, and underserved populations strengthens the project's focus on inclusion, trust, and equitable participation across diverse library settings. His expertise is especially relevant to governance, equitable access, and how such a model can serve smaller, rural, and under-resourced libraries. **Co-PI Tang's** expertise in social network analysis, digital communication, and online engagement supports the study of platform dynamics, amplification pathways, and network structure. This expertise will inform the analysis of how library-created content circulates, how participating libraries connect across the pilot network, and which coordination practices prove most effective.

2.3 Intended Beneficiaries

The immediate beneficiaries are public library staff and public library systems (approximately 9,000 systems serving communities nationwide [32]) that want to extend their educational outreach into online spaces but lack the staff time, communication workflows, and social media infrastructure to do so consistently. The need is especially acute for small, rural, and under-resourced libraries: for instance, a rural public library with 1–2 full-time-equivalent (FTE) staff that currently posts only event announcements on Facebook would gain access to vetted, platform-ready educational content it can share in minutes, without building a content strategy from scratch. Although this planning grant centers on public library contexts, where the trust advan-

tage and capacity constraints are most pronounced, the pilot deliberately includes limited participation from academic libraries (e.g., through advisory member Meg McMahon of Harvard Library) as learning partners to test cross-context transferability. Special libraries and additional academic settings can adapt outputs in future phases. The project design was informed by pre-proposal conversations with prospective pilot libraries and advisory members. Feedback from the Howard County Library System and the Birmingham Public Library confirmed that the participation cap, common-content model, and focus on AI ethics topics reflected their staff's expressed interests and capacity constraints (see also **Letters of Support**).

The broader beneficiaries are community members who need to learn about AI in everyday life. For instance, workers navigating AI-driven workplace changes, parents evaluating AI tools used in schools, and older adults encountering AI-generated health information need trustworthy, accessible guidance, as well as the field-level organizations that may coordinate future rollout. Because this is a planning project, the primary public benefit at this stage is preparatory: establishing whether libraries can extend their educational reach into the online environments where AI learning needs now arise. By producing evidence on public trust, librarian readiness, platform fit, staffing burden, governance needs, and feasible workflows, the planning work reduces uncertainty for state library agencies, associations, and consortia considering coordinated action, providing the profession with an evidence-based foundation for a more connected, public-facing approach.

3 Project Work Plan

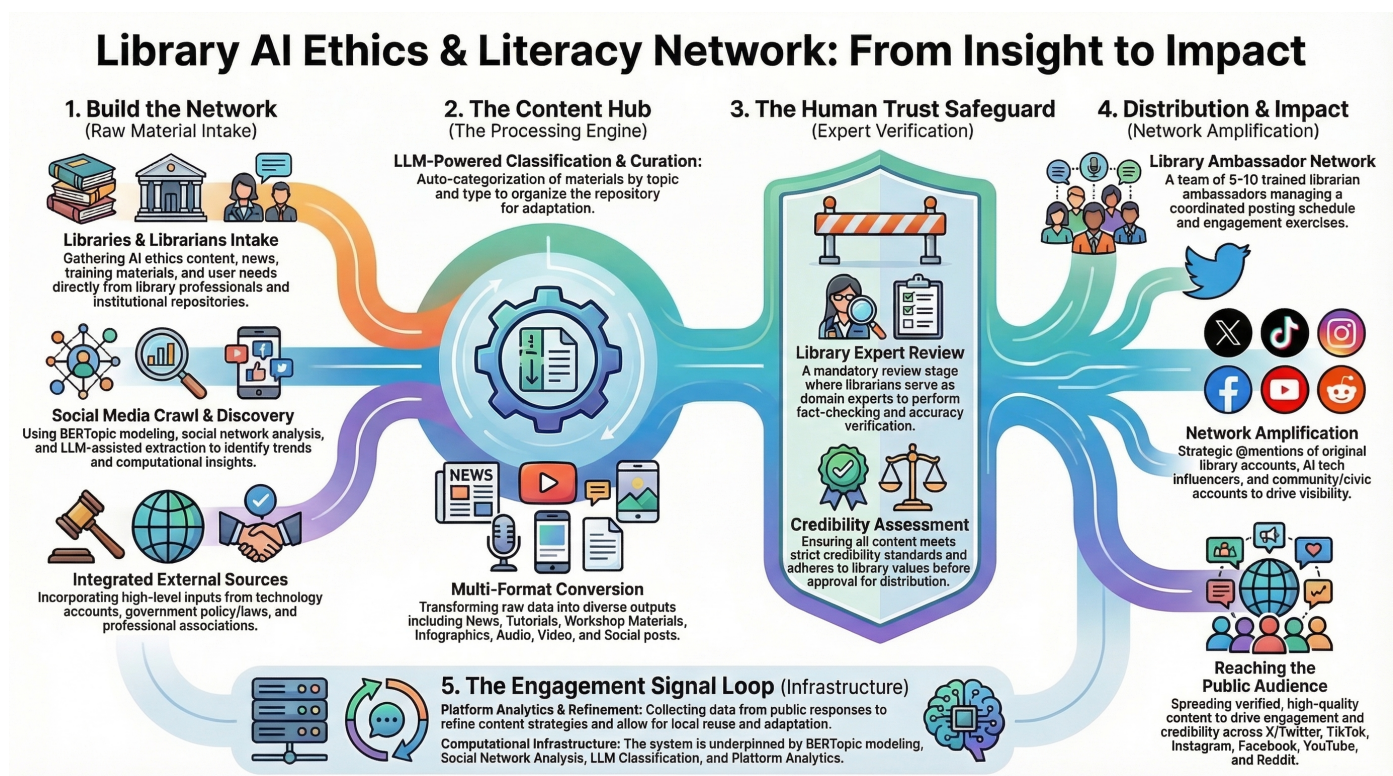


Figure 2: Conceptual model of the proposed Library AI Ethics & Literacy Network on Social Media. The workflow proceeds through five stages: (1) building the network through content intake from libraries, social media, and external sources; (2) processing through the Content Hub via LLM-assisted classification and multi-format draft conversion; (3) human review through librarian expertise and credibility assessment; (4) distribution and impact through the participating librarian network and coordinated amplification; and (5) an Engagement Signal Loop providing platform analytics and computational feedback to refine content strategy.

3.1 General Questions

The planning work is organized around four **general questions (GQs)**. **GQ1:** What learning needs, trust patterns, and participation conditions should guide the design of a library-affiliated AI ethics and literacy network? Through dual surveys of librarians and the public, this question establishes the design baseline (trust patterns, learning needs, participation conditions, and institutional readiness) that shapes all subsequent design decisions. **GQ2:** Under what conditions can a coordinated library-affiliated social media presence increase public engagement with library AI ethics and literacy resources, and how do public audiences perceive library-affiliated sources as educators relative to other providers? GQ2 assesses the extent to which a bounded, coordinated presence on one or two platforms can improve public engagement with these resources and examines how public audiences evaluate its trustworthiness under realistic field conditions. **GQ3:** What content intake, review, and transformation workflow would allow libraries to turn existing local materials into reusable AI ethics and literacy learning resources with manageable effort? Using the hub prototype, this question tests how effectively local materials (workshop content, FAQs, policy guidance, case examples) can be collected, reviewed, and adapted into such reusable learning resources with manageable effort. **GQ4:** What governance, staffing, stewardship, and partnership structure would be required to implement and sustain the network beyond the planning period? Drawing on evidence from the surveys, pilot, and hub prototype, this question identifies the oversight, staffing, moderation, and partnership structures required for sustainable operation beyond the planning period.

3.2 Project Components

The project is organized around **four Project Components (PCs)**, each tied to a core GQ and a set of planning deliverables. **Figure 2 illustrates how these components integrate into a unified workflow, from content intake and hub processing to expert review, organized distribution, and engagement feedback.** The project is phased over two years, with Year 1 focused on gathering evidence and Year 2 focused on synthesis and scale-up planning; a project website launches by Month 3 to provide public-facing project information (see *Schedule of Completion*).

PC1: Needs Assessment and Design Baseline. PC1 establishes the design baseline for the project through two online surveys: one for library workers and one for U.S. adults. The library survey examines current AI-related programming and communication practices, use of social media for public engagement, staff confidence in discussing AI-related topics, perceived barriers to sustained participation, and conditions under which involvement in such a model would be manageable. The public survey examines where adults learn about AI online, what AI topics they most want to understand better, and how they judge the credibility of library-based AI content.

For the library survey, we will recruit at least 200 library respondents through partner associations, state library networks, and professional listservs. For the public survey, we will recruit approximately 200 U.S. adults through an online panel service (e.g., Prolific), with sampling quotas designed to capture variation across age, region, and rural or urban residence (see *Budget Justification*). The combined recruitment target is approximately 400 respondents (≥ 200 librarians, ≥ 200 public), with a minimum analysis threshold of $n \geq 300$. The goal is structured planning evidence, not a definitive national opinion study. Both instruments are piloted for clarity before full deployment.

Our analysis focuses on identifying patterns that matter for scale-up planning. These include differences in trust, topic priorities, barriers to participation, platform use, and willingness to engage with library-based AI content. Open-ended responses are reviewed qualitatively to capture recurring concerns, expectations, and examples that may not appear in fixed-response items. The needs assessment produces **two deliverables**. The first is the Survey Package: two validated survey instruments (one for library workers, one for U.S. adults) and a codebook with variable-level definitions and missing-data conventions, completed after fielding closes (Month 10). The second is the Needs and Readiness Report, which synthesizes baseline findings into key design requirements—audience segments, priority content areas, participation constraints, and support needs—that guide the pilot and hub prototype (Month 15).

PC2: Pilot Social Media Network Design. In PC2, we will pilot a coordinated library-affiliated social media

presence on one or two platforms. Platform selection is guided by PC1 survey data on where librarians and the public encounter AI content, weighted against four operational criteria: API stability and data-access terms, content portability across platforms, audience demographics matching target populations, and compatibility with staff capacity at pilot libraries. Preliminary platform research and identification of participating librarians begins in Year 1 (Months 8–12), guided by emerging PC1 data, so that the formal pilot can launch efficiently once the Needs and Readiness Report is complete (Month 15). Our goal is to test how a shared network can connect library-affiliated participants, extend public AI education to online audiences, and support lower-burden public engagement around AI learning. The pilot is a bounded feasibility test designed to determine what kinds of participation are workable, not to build a public audience at scale.

The pilot network will include the project account structure, a small cohort of participating librarians, and simple coordination routines for sharing and amplification. Within this social media network, content moves in two directions: participating librarians submit or signal content for cross-library amplification, such as local events, workshop materials, or AI-related guidance, while also drawing from the common feed to re-post, adapt, or use for local communication. All amplified content credits the original creator. The two-way structure tests the extent to which coordinated amplification can increase visibility while reducing the burden of creating original content.

To support the pilot, we will analyze the surrounding platform environment using a computational analysis pipeline built on the PI's prior work in topic modeling and social network analysis. This pipeline identifies relevant AI-related conversations, recurring topics, influential accounts, and interaction patterns that shape how content circulates. Eventually, the analysis pipeline is released as **a reusable open-source deliverable** (see *Digital Products Plan*).

We will recruit a cohort of 10 librarians from partner networks and pilot sites, targeting at least two from rural settings, two from small-urban public libraries, and representation across public and academic library types. They receive brief training in core AI topics, plain-language communication, platform practices, and common participation guidelines. The 12-week pilot requires approximately one hour per week (~1 hr/wk) to reflect realistic field conditions. Throughout the pilot, participating librarians help us assess what kinds of sharing, reposting, adaptation, and local framing are manageable in practice, and what forms of support would be needed for broader rollout.

We will assess the pilot network using both platform measures and participant feedback. Measures include impressions, engagement, amplification actions, follower growth, response patterns, and participating librarians' reflections on workload, usability, and local relevance. **Two deliverables** emerge from this component: the Social Media Playbook & Participating Librarian Training Curriculum, which codifies coordination workflows, platform protocols, orientation modules, and time-management guidelines developed during the pilot, and the Analysis Pipeline described above. We will assess pilot feasibility using two thresholds: a retention rate of $\geq 75\%$ of the initial cohort through the full pilot period, and a mean manageability rating of ≥ 3.5 on a 5-point scale in post-pilot surveys.

PC3: Hub Prototype Design. PC3 develops and tests a limited hub prototype (a lightweight content management system or structured repository, not a production-ready platform) that functions as a learning content pipeline and common repository for the pilot. PC3 tests the extent to which a shared hub model can reduce the content-creation burden and support ongoing contributions, reuse, and local adaptation across participating libraries.

We will use the prototype hub to collect, organize, tag, and track library-AI materials through a simple three-step workflow: (1) gather AI-related content from participating libraries' organizational and individual professional social media accounts, along with selected materials from trusted sources such as library associations, technology companies, public-access education organizations, and state and federal government sources; (2) use large language model tools as drafting support to adapt selected materials into platform-ready formats for the pilot network, such as short explainers, question-and-answer posts, carousel text, myth-and-fact posts, captions, and other reusable multimedia content; and (3) require participating librarians to review all AI-assisted content before distribution. An editorial protocol requires human review and attribution labeling of every AI-assisted draft, with prompts and outputs logged for reproducibility.

The hub prototype processes at least 20 pilot cases through the full intake-to-publication workflow, providing concrete evidence on the extent to which a shared content infrastructure can reduce the original-content burden that currently prevents many libraries from extending their educational outreach into online spaces. This evidence feeds directly into the component's deliverable, the Hub Design, which comprises the prototype itself, a design document, a draft content taxonomy, a workflow specification, and at least 20 sample cases processed through the full workflow (see *Digital Products Plan*).

PC4: Implementation Planning and Governance. PC4 brings together findings from the needs assessment, pilot network, and hub prototype to define the requirements for a future rollout phase. We will focus on the organizational, editorial, and governance conditions required for the model to operate responsibly and at a scale useful to participating libraries.

Key planning questions include (1) who should host or steward the network, (2) how content submission and review should be managed, (3) what moderation and attribution practices are needed, (4) what ethical guidelines should govern public-facing AI-related communication, and (5) what staffing, training, coordination, and partnership structure would be required to support the model beyond the planning period. These questions are addressed using evidence from the earlier components, advisory input, targeted consultation with partner associations and participating libraries, and structured comparisons with existing federated governance models (e.g., the Digital Public Library of America and HathiTrust).

Three deliverables capture these findings: (1) The Feasibility Report provides a candid, evidence-grounded assessment of what worked, what did not, and under what conditions. (2) The Implementation Plan outlines recommended scope, governance structure, staffing assumptions, partner roles, review and moderation practices, risk considerations, and preliminary budget ranges. (3) The Partnership Model and Roster documents which organizations are positioned to co-lead or adopt the network at scale.

Each component will be assessed against readiness thresholds, including survey response adequacy, pilot retention and manageability targets (see PC2), hub workflow completion rates, and advisory endorsement at Gate 3, before advancing into the follow-on model; components that do not meet thresholds are flagged for redesign or deferral. Across all components, the project produces eight deliverables: the Survey Package and Needs and Readiness Report (PC1); the Social Media Playbook & Participating Librarian Training Curriculum and Analysis Pipeline (PC2); the Hub Design (PC3); and the Feasibility Report, Implementation Plan, and Partnership Model and Roster (PC4) (see *Digital Products Plan* and *Schedule of Completion*).

3.3 Risk Mitigation

Four risks bear on planning quality. First, low survey response will be addressed through staged recruitment across multiple channels (associations, listservs, state networks) and public-survey incentives. Second, platform-access disruption will be mitigated by designing the pilot to be portable across two candidate platforms. Third, participating librarian attrition will be managed through over-recruitment by 20% and modular training that supports mid-pilot onboarding. Fourth, if pilot libraries produce fewer hub submissions than projected, the team will supplement with curated public-domain AI educational materials from ALA, WebJunction by OCLC, and partner associations to maintain workflow testing.

3.4 Staff, Advisory Structure, and Partnerships

Staff Roles. PI Li leads survey design, Institutional Review Board (IRB) oversight, hub prototype development, and computational content discovery, and serves as project manager. Co-PI Mehra leads rural library outreach, advisory group development, and partnership cultivation. Co-PI Tang co-leads the public survey, social media strategy, engagement assessment, and development of the participating librarian training curriculum. A Graduate Research Assistant (GRA) with web development and content management skills supports computational analysis, data annotation, and hub development; an Undergraduate Research Assistant (URA) assists with survey administration, qualitative coding, hub workflow, and training materials. PI assignments by component: Li (PC1, PC3), Tang (PC2), Mehra (PC4); all PIs contribute to PC4 synthesis (see *Budget Justification* and *Resumes*). During the academic year, the PIs contribute effort through their

regular faculty appointments; computationally intensive tasks are batched into the summer months when the GRA is available.

Advisory Structure. The project convenes an advisory group of 6-8 members representing library associations, public library leaders, academic library practitioners, and pilot library staff. Co-PI Mehra leads advisory group development, with representation from rural and under-resourced library settings. The group includes at least one rural library director, one state library agency representative, and one communication practitioner to ensure breadth beyond academic and large-system perspectives. The advisory group meets quarterly via video conference and provides written feedback at three design gates aligned with major project milestones: Gate 1 (Month 4) reviews survey instruments and pilot design; Gate 2 (Month 10) reviews the hub taxonomy and pilot workflow; and Gate 3 (Month 20) reviews a preliminary Implementation Plan informed by available pilot and hub evidence through that point. At each gate, the group provides a documented go, revise, or defer recommendation, recorded in a decision memo and used to guide subsequent project activities.

Partnerships. PI Li's sustained leadership in ASIS&T since 2022, and now Chair-Elect of ASIS&T's SIG on AI Governance, provides a direct institutional channel for survey distribution through ASIS&T membership networks, practitioner recruitment from the AI governance research community, and dissemination at the ASIS&T Annual Meeting.

Building on PI Li's and Co-PI Mehra's existing professional networks, the project develops partnerships with national and state-level library associations to support survey distribution, recruitment of participating librarians, and dissemination. Partner-development visits in Year 1 (budgeted under program travel) establish initial partnership agreements by Month 3 and formalize memoranda of understanding by Month 6. A project website launches by Month 3 to serve as the public-facing information point for the project, including project goals, team members, partnership opportunities, and deliverables as they become available. Two pilot library partners have confirmed participation (see Letters of Support): the Howard County Library System (Tonya Aikens, President and CEO) can host hub usability testing and contribute staff time for participating librarian workflow development; the Birmingham Public Library (Katie Vainrib) can provide pilot cases, local adaptation testing, and formative feedback on content relevance across library types. Meg McMahon (Harvard Library) and James Waller (EBSCO Information Services) can participate in the project and serve as advisory members, contributing expertise in academic library user experience and information services platform design, respectively. Jia Zhao (University of Alabama, Department of Mathematics) can advise on machine learning, and responsible AI methodology, supporting computational approaches. Joni K. Hayward Marcum (James Madison University) contributes practitioner perspectives on digital literacies and information literacy frameworks, bridging academic research with library-based digital education practice. Pilot library staff provide formative input through day-to-day participation in hub and pilot activities, including brief monthly check-ins with the project team. Two formal feedback sessions (approximately Months 15 and 19) collect structured assessments of hub usability, participating librarian workflows, and local relevance; the resulting findings inform the Gate 3 preliminary review (Month 20), with complete results incorporated into the final Implementation Plan (Month 24).

External Evaluator. An external evaluator is selected through a competitive process by Month 2 and conducts a developmental evaluation across three deliverables: a formative midpoint memo (Month 12), an interim brief to the advisory group (Month 18), and a summative evaluation report (Month 24). The evaluation addresses three core questions: whether the survey instruments yield valid and actionable planning data, whether the pilot and hub produce sufficient feasibility evidence to support scale-up decisions, and whether the hub workflow is practical for libraries with varying capacities. The evaluator has access to advisory gate materials and decision memos to inform the assessment, but the advisory group makes the gate decisions. Evaluator findings are incorporated into the final Implementation Plan (see *Budget Justification* and *Resumes*).

3.5 Evaluation and Demonstration of Impact

Evaluation is integrated across all four project components to assess feasibility, refine design decisions, and document the model's potential impact on public AI education in library contexts.

(1) For PC1: Needs Assessment and Design Baseline, indicators include response rates across library types and demographic strata, distribution of key constructs (trust, topic priorities, participation barriers), and convergence of open-ended responses on recurring needs. (2) For PC2: Pilot Social Media Network Design, we will use mixed methods to evaluate both platform-level performance and participant experience. Participant indicators include librarians' reported workload, perceived usability of coordination routines, perceived fit with local communication practices, and perceived value of the network for extending AI education. Using brief post-exposure survey items, we will also assess whether a small sample of audience respondents rate library-based AI content as useful and trustworthy. These data are gathered through brief logs, surveys, and a structured debrief at the end of the pilot period (see *Performance Measurement Plan*). (3) For PC3: Hub Prototype Design, evaluation focuses on workflow viability and content quality. We will track the number and type of source materials processed through the hub, the proportion successfully transformed into platform-ready formats, and the time and steps required for intake, adaptation, and review. Participating librarians assess a sample of hub-generated or hub-assisted content on dimensions such as accuracy, clarity, alignment with library values, appropriateness for public audiences, and ease of local adaptation. These findings demonstrate the extent to which a reusable hub infrastructure can realistically reduce original content burden and support reusable, trustworthy public-facing AI resources. (4) For PC4: Implementation Planning and Governance, we will solicit structured feedback from the advisory group, pilot partners, and the external evaluator on governance, staffing, partnership, and moderation assumptions. Each design gate (Months 4, 10, and 20) includes explicit criteria and a documented go/revise/defer decision. The evaluator's midpoint, interim, and summative reports provide an independent assessment of progress and the quality of planning.

4 Project Results

Deliverables. The planning work produces eight deliverables that collectively address the learning gap identified in this proposal: libraries are among America's most trusted educators, yet they remain largely absent from the online spaces where public understanding of AI is forming. Table 1 lists all eight deliverables with the short reference names used throughout this narrative.

#	Deliverable	Short Name	PC
<i>Evidence products</i>			
1	Validated survey instruments and codebook	Survey Package	PC1
2	Needs and Readiness Report	Needs Report	PC1
3	BERTopic/SNA analysis pipeline	Analysis Pipeline	PC2
<i>Pilot-tested designs</i>			
4	Social Media Playbook & Participating Librarian Training Curriculum	Playbook & Curriculum	PC2
5	Hub Prototype and Design Document	Hub Design	PC3
<i>Planning-to-implementation products</i>			
6	Feasibility Report	Feasibility Report	PC4
7	Implementation Plan	Implementation Plan	PC4
8	Partnership Model and Roster	Partnership Model & Roster	PC4

Table 1: Eight project deliverables organized by functional category.

(1) Evidence products establish a dual-perspective design baseline and provide reusable computational tools for identifying where library-led public AI education can have the greatest impact within broader AI discourse. (2) Pilot-tested designs translate pilot findings into operational specifications that other institutions can adopt. (3) Planning-to-implementation products synthesize what worked, under what conditions, and what a follow-on phase would require. Activities for each deliverable are described in the Work Plan; delivery dates are specified in the *Schedule of Completion*; and formats, licensing, metadata standards, and preservation commitments are detailed in the *Digital Products Plan*.

How results address needs. Quality indicators for each deliverable are specified (see *Performance Measurement Plan*), and the external evaluator will independently assess the quality of the outputs at three points

during the grant period. (1) The Needs and Readiness Report provides the first dual-perspective evidence base linking public AI learning needs with librarian readiness for cross-library social media engagement; these data do not currently exist at this level of specificity. (2) The Social Media Playbook and Hub Design Document give libraries and consortia operational models they can adapt without replicating the full planning process. (3) The open-source BERTopic/SNA analysis pipeline extends the PI's published methodology into a reusable tool for identifying where library-created content can have the most impact within broader AI discourse [27, 28]. (4) The Feasibility Report provides a candid, evidence-grounded assessment of what worked, what did not, and under what resource and staffing conditions. Such candid assessments offer the planning evidence most useful for informing future decisions.

Replicability and Adequacy. To ensure that results are adaptable and usable beyond the pilot setting, we have designed the project around variation rather than a single ideal context. The pilot involves 10 participating librarians representing different library types (public and academic), community settings (urban, suburban, rural), and staffing levels. The surveys will recruit across geographic and institutional diversity. The bounded scope is appropriate for a planning grant: the goal is to determine what model could work at a broader scale, not to validate it at full scale. All deliverables emphasize transferable design principles and clearly stated adoption assumptions (library size, staff time required, platform conditions) so that other institutions can judge which elements are adaptable to their own setting. Where appropriate, outputs are presented in modular form (executive summaries, workflow diagrams, taxonomies, checklists) to support selective adoption by associations, consortia, state library agencies, and individual libraries (see *Digital Products Plan*).


Dissemination. We will disseminate results through both scholarly and practitioner channels over 24 months. Conference presentations will target ALA Annual in Year 1 for practitioner feedback and ASIS&T in Year 2 for scholarly dissemination of the project model and emerging findings; ALISE and iConference will be pursued as additional venues as manuscripts develop (see *Budget Justification*). Practitioner dissemination will include webinars through partner associations and state library networks, briefs through channels such as WebJunction and Library Journal, and direct briefings to participating and interested libraries. Public-facing summaries of key findings, including practical AI literacy tips and myth-versus-fact graphics, are shared through the pilot social media channels developed during the project and amplified through partner libraries' existing social media networks, modeling the approach the project seeks to scale. Final products are also shared through association webinars and repository deposits in Months 22–24. All products are deposited in open-access repositories (the UA Institutional Repository for reports and instruments, GitHub and Zenodo for code, and OSF for survey data) within 60 days of grant close (see *Digital Products Plan*).

Sustainability. The benefits of the project continue after the grant period through institutional commitments and a clearly defined pathway to future work. The University of Alabama hosts all final products in its institutional repository at no cost beyond the grant period, with permanent URLs and metadata indexing. Code repositories are archived on Zenodo with DataCite DOIs for long-term discoverability and community maintenance. All products carry open licenses (CC BY 4.0 for documents and MIT for software) so that any library, association, or consortium can use, adapt, and redistribute them without permission or cost (see *Digital Products Plan*). Beyond product preservation, this planning project is designed to produce the evidence, designs, and partnership commitments needed to inform a future decision about whether to pursue a follow-on IMLS implementation proposal. The Implementation Plan includes recommended scope, governance, staffing, partner roles, evaluation design, and a preliminary budget for the follow-on phase. Should follow-on funding not materialize, the network model is designed for lightweight continuation: participating librarians can sustain joint posting through their existing social media accounts at no additional cost, partner associations can incorporate the playbook into ongoing digital literacy programming, and the open hub design allows any consortium or state library agency to adapt the model locally. The Partnership Model & Roster documents which organizations are positioned to co-lead or adopt the network at scale.

Schedule of Completion

Column numbers and M_n references denote project months (M1 = September 2026).


YEAR 1 (September 2026 – August 2027)



Key Milestones and Products

Description	Due
Advisory formation; IRB approval; project website launch	M3
MOU finalization with partner libraries	M6
Survey fielding complete ($n \approx 400$)	M8
Validated survey instruments and codebook	M10
Formative midpoint evaluation memo	M12

YEAR 2 (September 2027 – August 2028)



Key Milestones and Products (continued)

Description	Due
Platform selection decision	M14
Needs & Readiness Report	M15
Hub prototype and design document	M20
Social media playbook and participating librarian training curriculum	M22
Feasibility Report	M22
Implementation Plan	M24
Partnership model and roster	M24
BERTopic/SNA analysis pipeline and code release	M24