



NATIONAL LEADERSHIP GRANTS FOR LIBRARIES

Sample Application LG-259801-OLS-26
Project Category: Planning

University of Pittsburgh

Amount awarded by IMLS: \$196,394
Amount of cost share: \$0

The University of Pittsburgh will plan to develop the Civic Algorithmic Sandbox, a library-based artificial intelligence literacy initiative that helps the public understand how government agencies use data-driven decision tools in services that shape their economic lives. The project is comprised of AlgoHub, an interactive simulation platform; a librarian facilitation training curriculum; and a community knowledge documentation practice focused on public services. The primary target group for planning activities is public library staff at the Western Allegheny Community Library in Allegheny County, Pennsylvania. The second target group is library patrons who will participate in AlgoHub testing and community listening sessions. This project examines the knowledge, skills, and support that public librarians need to facilitate community engagement with government decision systems. This project is focused on how community members currently understand how government agencies use data and algorithmic systems in public services.

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.

Advancing AI Literacy Through Public Libraries: The Civic Algorithmic Sandbox

National Leadership Grants for Libraries: Planning Grant Program

Introduction

The University of Pittsburgh, in partnership with the Western Allegheny Community Library requests \$196,237 from the Institute of Museum and Library Services to support a **two-year planning grant** for the Civic Algorithmic Sandbox (CAS). CAS is a library-based AI literacy initiative that helps the American public understand how data-driven decision tools work in the government services that shape their economic lives. The initiative integrates three components developed through co-design with public library staff and patrons: (1) **an interactive simulation platform called AlgoHub**, (2) **a librarian facilitation training curriculum**, and (3) **a structured practice for documenting community knowledge about public services**.

AI is now embedded in government services ranging from workforce placement and benefits eligibility to housing assistance and child welfare screening (Eubanks, 2018). Americans who understand how these tools function are better positioned to navigate them effectively, make informed decisions about available services, and participate productively in an AI-driven economy (Haesevoets et al., 2024; Alon-Barkat et al., 2023). But awareness remains remarkably low. A 2021 survey of 1,500 Pittsburgh residents found that only 8% had any knowledge of algorithmic systems used by local government agencies (Lin et al., 2023). Public libraries are trusted and accessible civic learning institutions with established roles in public education, digital literacy, and community engagement, making them a strong institutional setting for this work. They, however, currently lack the tools, training, and facilitation frameworks to offer substantive AI literacy programming that goes beyond introductory awareness (Huang et al., 2024). In this project, AI literacy refers to the **practical ability to recognize when algorithmic decision systems are being used in public services, understand in simplified but structurally faithful terms how those systems process information, and discuss their consequences, limits, and points of human discretion**.

We have developed a working proof-of-concept simulation platform, AlgoHub, that demonstrates the technical feasibility of this approach. During the planning period, we will co-design the three CAS components with library staff and patrons through community listening sessions, iterative participatory development cycles, and curriculum piloting. This co-design process is intended to align the resulting tools and training with library workflows, staff capacities, and patrons' expressed learning needs. All tools and findings will be openly available to help other library systems assess what may be adaptable in their own contexts.

NLG-L Program Objective: This **project primarily addresses Objective 1.1: Serve the learning needs of the public through libraries and archives**. CAS addresses it by helping public libraries determine the conditions under which a model is feasible and useful for supporting public understanding of how algorithmic systems operate in government services.

Project Justification

The Challenge: AI in Public Services and the Economic Stakes for American Families

Algorithmic decision systems are already used in a range of government services. Federal, state, and local governments are deploying AI-powered decision systems in workforce development, benefits eligibility, housing assistance, child welfare, and public health, with direct consequences for the economic stability and financial progress of millions of American families. These systems play a growing role in determining who receives job training referrals, how benefits applications are processed, which families receive home-visiting services, and how risk is assessed across a range of public programs.

Americans who understand how these tools function can engage more productively with the agencies that serve them, make better-informed decisions about available services, and **better understand and navigate the public systems that increasingly shape economic opportunity and access to services**. Those who do not are at a disadvantage in navigating systems that increasingly shape their economic futures. Long and Magerko (2020) identify the core competencies of AI literacy as including the ability to recognize when AI is being used, understand the basic logic of how it processes

information, and evaluate the outputs it produces. Ridley and Pawlick-Potts (2021) argue that libraries are uniquely positioned to develop this literacy because they already serve as trusted sites for information literacy and can extend that role to include algorithmic systems. Kong, Dilinika, Nie, Gautam, and Huang (2025) further demonstrate that AI literacy involves not only technical comprehension but also socio-ethical engagement, the capacity to reason about how AI systems affect people and communities in practice. This fuller conception of AI literacy is precisely what a simulation platform grounded in real public services can support.

The White House AI Action Plan identifies enabling AI innovation and adoption as a national priority, and Executive Order 14277, “Advancing Artificial Intelligence Education for American Youth,” calls for practical AI education accessible to all Americans. Public libraries serve every community in the country, they are trusted across demographic and political lines, and they have decades of experience translating complex information domains into practical public programming. What they currently lack are tested tools, facilitation resources, and programming models for doing this work around AI.

The Gap: Library Infrastructure for Public Engagement with Government AI


Despite growing interest, libraries are not yet equipped to help communities engage with the government AI systems that affect their economic lives. The 2023 PLA Public Library Technology Survey found that 61% of libraries responding to an open-ended question about emerging technologies identified AI as a priority area for new programming (Public Library Association, 2024). At the same time, 54% of libraries overall cited limited staff capacity as the greatest challenge they face in providing digital literacy training, and 27% specifically identified lack of staff expertise as a barrier.

This pattern holds among academic library professionals as well. Lo’s (2024) survey of 760 academic library employees found that 70% did not feel prepared to adopt AI tools within the coming year, and nearly 63% disagreed with the statement that they felt adequately prepared to use generative AI tools in their professional work. Confidence in facilitating discussions about AI integration was low across all position levels, with troubleshooting AI tools emerging as the area of weakest confidence. If academic librarians with graduate training and institutional resources report these gaps, the challenges facing public library staff in community-facing settings are at least as significant.

Huang, Gascó-Hernandez, Prasad, and Gil-Garcia (2024) provide a systematic picture of what public libraries are doing around AI. Their environmental scan of 105 AI-related programs across North American public libraries found five categories of activity: open lectures and podcasts, hands-on workshops, seminars and conversations, exhibitions, and makerspaces. Nearly all of these focus on commercial AI tools and general AI awareness. None offer interactive simulations of government decision systems. None include facilitation training for librarians specific to AI in public services. And none incorporate structured processes for community members to document their own experiences. Existing public-library AI programming demonstrates institutional interest, but it has not yet produced tested models for helping patrons examine how government decision systems shape access to services, rights, and opportunities.

We identify three specific gaps that limit libraries’ ability to support public engagement with government AI systems.

Figure 1: Screenshot from the AlgoHub interactive simulation platform (working prototype). The working prototype reduces project risk by showing that the proposed simulation approach is technically feasible. The planning grant will determine what level of fidelity, explanation, facilitation, and community relevance is needed for such a tool to function effectively in public library settings.



Related Stories

Stories about Child Welfare Screening

Case Study: How Allegheny County Uses Predictive Analytics in Child Welfare

An in-depth look at how the AFST was developed, deployed, and evaluated — including community engagement and ethical review processes.

[Read story →](#)

Balancing Algorithmic Precision and Equity in Child Protection

Examining the tension between predictive accuracy and racial equity in automated decision-support tools used in child welfare systems.

First, we found no tested simulation platform for public-library learning about government decision systems in the literature and examples reviewed. Existing AI education resources focus on commercial AI tools such as chatbots or image generators. They do not examine government decision systems that determine benefits eligibility, workforce referrals, and service access. A community member may use ChatGPT but has no way to explore how a risk prediction model weighs data to produce a score that influences a caseworker's decision.


Second, the field lacks facilitation training specific to public engagement with AI in government services. Traditional digital literacy training covers device use, internet navigation, and software skills. Algorithmic literacy requires a different skill set: the ability to help patrons reason about how data inputs become decision outputs, what the boundaries of these systems are, and how to engage with agencies that use them (Ridley & Pawlick-Potts, 2021). We could not find any tested facilitation training of this kind in scholarship.

Third, there is no established method for libraries to capture and organize what community members already know about navigating public services. Community members hold practical knowledge about how government systems function in their lives, knowledge that is valuable both to libraries designing responsive programming and to the agencies seeking to improve their services. However, libraries currently lack established documentation practices and tools for capturing, organizing, and reusing this knowledge. Research on community-inclusive data ecosystems shows that local organizations face barriers in accessing and using data effectively, and that more inclusive ecosystems require educational programming, partnership structures, and mechanisms for bringing community voices forward (Yoon & Copeland, 2020).

Why a Simulation Platform

Government AI systems are fundamentally about process: data goes in, a model applies rules and weights, and an output informs a human decision. Static presentations or lectures can introduce these systems, but they are less effective for showing how changing inputs, rules, or thresholds can alter outputs. A simulation that allows people to adjust inputs, observe how outputs change, and explore how different data points carry different weight produces the kind of hands-on, exploratory understanding that passive instruction cannot (Barocas, Hardt, & Narayanan, 2023). We have previously argued for reconfiguring digital tools toward existing infrastructure and promoting reflection through intentionally paced design, principles drawn from their work integrating digital literacy support into community organizations (Gautam et al., 2024). The same logic applies here: a simulation platform designed for library settings, and built through co-design with library staff and patrons, can be integrated into existing library programming rather than requiring an entirely new service model.

Figure 2: Users can explore how an algorithmic decision-making system processes inputs and generates outcome scores.



Prior interactive educational tools have demonstrated that narrativized simulations can build intuitive understanding of AI mechanisms for non-technical audiences. The Wilson Center's "AI's Anatomy" uses an interactive game to teach about AI thresholds and bias through a healthcare eligibility scenario (Wilson Center, 2019). "Survival of the Best Fit" illustrates bias in recruitment algorithms through a hiring simulation (Csapo et al., 2019). Both are well designed, but both rely on hypothetical scenarios rather than documented public systems that affect identifiable communities. In this project, we are interested in whether grounding a simulation in a real public system, especially one that operates in the region where participants live, can support more situated learning and discussion.

AlgoHub is an educational simulation platform built from publicly documented descriptions of the algorithmic systems. It will not use real case data or interface with county systems. It will provide librarians with configurable teaching modules that allow workshop participants to experiment with how algorithmic systems process information. These modules will be designed so that library staff can adapt activities to different audiences, program formats, and algorithmic systems with limited technical modification. For example, they can use the Allegheny County Family Screening Tool (AFST), which is a data-driven decision support tool used in child welfare intake screening in Allegheny County. It has been the subject of

extensive public documentation including the original methodology report (Vaithianathan et al., 2017), independent academic analysis of its decision-making properties (Chouldechova et al., 2018), a county-commissioned ethical review (Dare & Gambrill, 2017), a third-party impact evaluation (Hornby Zeller Associates, 2018), and an external audit (Gerchick et al., 2023). This depth of publicly available documentation makes the AFST particularly well suited for educational use where we can build a structurally faithful simulation, and patrons can verify what they learn against published reports.

AFST is used here as an initial pedagogical case rather than as the sole focus of the project. It exemplifies how a government AI system works in practice, that is, how data from multiple sources is integrated, how different variables are weighted, and how a numerical output informs but does not replace a human decision. Through co-design with librarians and library patrons, we will identify design requirements for adapting the framework to additional public-service contexts such as workforce services, benefits administration, and housing.

Target Group and Beneficiaries

The primary target group for planning activities is public library staff at the Western Allegheny Community Library in Allegheny County, Pennsylvania. Approximately 10 staff members, including the library director, adult services staff, and community engagement personnel, will participate in needs assessment, co-design workshops, and curriculum piloting during Phase 1. In Phase 2, the project will expand to include an additional 40 library staff from neighboring branches and library systems in the Allegheny County region whose leadership expresses interest in the project based on Phase 1 results. This incremental approach is deliberate. **We want to build trust and demonstrate value at one site, then invite participation from other libraries based on their own assessment of fit.**

The second target group is library patrons who will participate in AlgoHub testing and community listening sessions. During Phase 1, around 30 patrons at Western Allegheny Community Library will participate. In Phase 2, an additional 120 patrons across expanded sites will be engaged. All patron participation is voluntary and compensated. The project is designed for adult public library patrons with varying levels of familiarity with government services, data systems, and AI. Testing sessions will naturally include people with a range of experiences navigating government services, which strengthens the feedback the project team receives about what patrons want to understand and how AlgoHub can support that learning.

These gaps are not specific to Allegheny County. Public libraries across the country are beginning to offer AI-related programming, but as Huang et al. (2024) documented, no existing programs provide interactive simulations of government decision systems, facilitation training specific to AI in public services, or structured ways for community members to contribute their own knowledge. A planning grant is needed to determine, through co-design with library staff and patrons, what a replicable model for informed public engagement with government AI systems would require: what level of simulation fidelity supports learning, what training prepares librarians to facilitate these conversations, what documentation practices capture community knowledge effectively, and what staffing and governance structures make the model adoptable by libraries with different capacities and contexts. The Civic Algorithmic Sandbox addresses this by building and testing these components in one community before proposing broader implementation. Aligning with this goal, we will make the AlgoHub platform, and the facilitation materials developed through this project openly available to the public. This planning grant will determine what level of simulation fidelity, facilitation support, documentation practice, and governance structure are necessary for a library-based model to be feasible, useful, and adaptable beyond the initial site. It will provide early design evidence and materials that other libraries can review and adapt, and **inform future national implementation effort.**

Project Work Plan

The planning grant spans 24 months, organized in three phases. Phase 1 (Months 1-6) centers on community listening, needs assessment, and foundational co-design with library staff and patrons at Western Allegheny Community Library. Phase 2 (Months 7-18) focuses on iterative co-design, development, and testing of AlgoHub and the librarian training curriculum, with incremental expansion to additional library sites based on demonstrated results and expressed interest. Phase 3 (Months 16-24) finalizes the planning deliverables, develops the community knowledge base design, and prepares evidence and specifications for a future national implementation phase. Phase 3 overlaps with Phase 2 so that documentation and synthesis work can proceed in parallel with late-stage testing and curriculum refinement. This phased model will engage approximately 200 total participants, beginning with 40 participants at Western Allegheny and expanding to an additional 160 participants across interested regional library sites. The project addresses four planning questions that will inform the design of a library-based model for public engagement with algorithmic decision systems:

1. What knowledge, skills, and support do public librarians need to facilitate community engagement with government decision systems?
2. What do community members currently understand about how government agencies use data and algorithmic systems in public services, and what do they want to know?
3. What level of simulation detail is sufficient for patrons to meaningfully understand how algorithmic decision systems operate in public-service contexts?
4. What forms of documentation and governance can enable libraries to incorporate community knowledge and lived experience into programming about algorithmic decision systems in public services?

Personnel

The PI, Aakash Gautam, is Assistant Professor in the Department of Information Culture and Data Stewardship (ICDS) at the University of Pittsburgh's School of Computing and Information. He has a background in computer science and human-computer interaction. His research focuses on human-centered design with community organizations, and he has led the design, development, and evaluation of interactive digital systems for nonprofit organizations and adult literacy training centers (Gautam, Gandhi, & Sendejo, 2024). He supervises graduate and undergraduate students in computer science and teaches a core MLIS course on data and information systems. The PI will serve as project director, leading the research and co-design activities, supervising the development team, and coordinating with library partners.

The Co-PI, Eleanor “Nora” Mattern, is Director of the Sara Fine Institute and Teaching Associate Professor in ICDS. Her research focuses on community-centered information work, civic engagement, and the role of libraries in civic data ecosystems. She is the PI of the Civic Switchboard project, funded through an IMLS National Leadership Grant and two IMLS Laura Bush 21st Century Librarian grants, which developed guides, toolkits, and professional learning resources to strengthen partnerships between public libraries and local data intermediaries. This prior IMLS-funded work on positioning libraries as active civic information partners provides direct methodological and institutional knowledge for this CAS project. Mattern will co-lead the design of community listening sessions and the community knowledge base and will contribute to curriculum development drawing on her expertise in how libraries engage communities around civic data.

The development team comprises one graduate student and three senior undergraduate computer science (CS) students participating through the CS Capstone courses, in which students complete technical projects under faculty supervision. This team will lead platform development, conduct usability testing, and support data collection and analysis.

Heather L. Auman, Manager of Library Services at the Western Allegheny Community Library, will serve as site coordinator (please see the attached letter of support). She will facilitate staff participation in needs assessment and co-design activities, coordinate patron recruitment for listening sessions and testing, and ensure that the project design reflects the operational realities of community library service.

A Community Advisory Group (CAG) of five to seven members will be recruited during Phase 1, comprising library staff, patrons from Western Allegheny, and at least one practitioner with direct professional experience in a domain where government AI systems are actively used (such as workforce services, child welfare, or public housing). The CAG will meet quarterly throughout the project duration. They will review all major design decisions, and provide substantive input into platform priorities, curriculum framing, and documentation practices. **The CAG will also conduct a structured review of deliverables against the four planning questions.** We believe this will help provide a more authentic assessment based on both professional and lived experience. CAG members will be compensated in recognition of their expertise.

Phase 1: Community Listening and Needs Assessment (Months 1-6)

Activity 1.1. Project Launch and Partnership Governance (Months 1-2)

The University of Pittsburgh and the Western Allegheny Community Library will convene a launch meeting to finalize governance structures, communication protocols, and shared decision-making processes. This meeting establishes the co-design relationship from the outset: both the university team and the library partner have voice in setting priorities, shaping timelines, and determining what success looks like. The CAG will be recruited and convened. IRB approval will be secured for all activities involving human participants.

As part of launch activities, the project team will **conduct an initial review of existing AI literacy curricula, simulation-based learning tools, and community engagement frameworks relevant to government AI systems**. This review will help the planning work build on rather than duplicate prior efforts. The review will be updated as the project progresses and synthesized in the final planning report.

Activity 1.2. Librarian Needs Assessment (Months 2-4)

We will design and administer a mixed methods needs assessment with library staff at Western Allegheny Community Library. The assessment is structured around Lo's (2024) and our (Kong et al., 2025) frameworks of core AI literacy competencies for library professionals: understanding AI system capabilities and limitations, identifying and evaluating AI use cases, utilizing AI tools effectively, critically assessing AI quality and biases, engaging in informed AI discussions, recognizing data privacy and security issues, and anticipating AI's impacts on library stakeholders. These existing frameworks provide empirically grounded model for what AI literacy means in a library context. This will help ensure that the needs assessment **captures specific competency gaps** that our subsequent curriculum can address.

The assessment will include a structured survey administered to all staff and in-depth interviews with six to eight staff members selected to represent variation in role, programming experience, and prior comfort with technology topics. Lo found that confidence in facilitating AI discussions and troubleshooting AI tools were the weakest areas across all position levels, and that 63% of library professionals did not feel adequately prepared to use AI tools in their work. The needs assessment will determine whether these patterns hold in a public library community-facing context, where the demands are different from academic settings. The findings address the first planning question and will ground the curriculum developed in Phase 2 based on the documented staff realities.

Activity 1.3. Community Listening Sessions (Months 3-6)

In partnership with library staff, the PI, Co-PI, and library staff will facilitate four community listening sessions at Western Allegheny Community Library. These sessions are the first stage of the co-design process with patrons. These sessions serve two goals. First, we will gain an understanding of what people know or do not know. Second, and more importantly, we will position patrons as contributors to the design of CAS.

The sessions are designed in response to a specific gap identified in the literature. Huang et al. (2024) found that while public libraries offer various AI-related programs, none of them incorporate structured processes for community members to contribute their own experiences and knowledge. Lin et al. (2023) found that only 8% of Pittsburgh-area residents were aware of algorithmic systems used by local government. The listening sessions go beyond this baseline by surfacing not just what people know, but what they want to understand. Community members who interact with workforce services, benefits systems, or housing assistance carry practical knowledge about how these systems function in their daily lives, knowledge that is rarely surfaced or valued in the design of AI literacy programs. Gautam et al. (2024) argue for positioning the research team as a "partner on the side" that facilitates rather than drives community-based efforts. We operationalize this posture by creating the space and structure for dialogue, but patrons set the substantive direction by identifying what they want to understand, what concerns them, and what they think a library-based learning tool should help them do.

Sessions will use a structured dialogue format that treats patrons' lived experience navigating public services as a substantive source of knowledge for design. The project serves all library patrons and is designed for any community member interested in understanding how government agencies use data and AI in the public services available to all Americans. Documentation from these sessions will be reviewed by participants before being incorporated into planning outputs.

These sessions directly address the second planning question and allow AlgoHub's design priorities to be shaped by the people it is intended to serve, and not solely by the research team. They also begin to surface the kinds of community knowledge that will inform the design of the documentation practices addressed in Phase 3.

Phase 2: Iterative Co-Design, Development, and Testing (Months 7-18)

Activity 2.1. AlgoHub Co-Design and Iterative Development (Months 7-13)

Drawing on Phase 1 findings, the PI and the development team will refine and extend AlgoHub through a series of rapid co-design and testing cycles. The PI brings a computer science background and prior experience leading the design, development, and evaluation of interactive digital systems for nonprofit organizations and adult literacy training centers (e.g., Gautam et al., 2024; Gautam et al., 2020), providing both the technical depth and the community-oriented design sensibility this work requires.

AlgoHub will support librarian-led workshops in which patrons interact with simplified models of algorithmic decision systems used in public services. Through guided activities, participants will explore how data inputs, model assumptions, and policy choices influence algorithmic outcomes. Users will be able to see, for example, how factors like a family's interaction with various public services influence the tool's assessment, building concrete understanding of how AI-driven government decision tools process information in services that affect families' economic stability (please see Figure 2).

The design of AlgoHub responds to specific limitations in prior interactive AI education tools. The Wilson Center's "AI's Anatomy" and "Navigating Neural Networks" (Wilson Center, 2019) and "Survival of the Best Fit" (Csapo et al., 2019) demonstrated that narrativized simulations can build intuitive understanding of AI mechanisms for non-technical audiences. All three, however, use abstract or hypothetical scenarios (hospital eligibility, image classification, hiring) rather than systems that affect real communities. As the PI's prior research on this topic has documented, abstract examples miss the element of relatability with the participant (Gautam, 2024). When people explore a system that actually operates in the region where they live, using data categories they recognize from their own interactions with public agencies, the learning is grounded in something that matters to them directly. AlgoHub tests this hypothesis through iterative development: does a simulation grounded in a real, local system produce deeper engagement and understanding than abstract alternatives?

Co-design means that community members and library staff are active contributors to design decisions, not passive testers who receive a finished product. Each development cycle includes a build phase, a testing phase with eight to ten community members and three to four library staff, and a revision phase incorporating structured feedback. Participants help determine what features matter, what explanations make sense, what feels confusing or misleading, and what additional government systems the platform should eventually cover. This iterative approach directly addresses the third planning question about simulation fidelity: rather than deciding in advance what level of technical detail is appropriate, the project lets patrons and library staff determine this through repeated interaction with working prototypes.

Activity 2.2. Expanded Testing and Feedback (Months 11-16)

Based on the results from Activity 2.1 and the interest generated by the work at Western Allegheny, the PI and co-PI will invite participation from additional library staff and patrons. The expansion model is incremental, with additional sites invited after Phase 1 findings and early prototypes are shared and discussed with interested library leaders (please see letter of support from Allegheny County Library Association).

During this phase, the project will engage an additional 40 library staff and 120 patrons beyond the 10 staff and 30 patrons involved in Phase 1, bringing total participation to approximately 50 staff, 150 patrons, and 200 participants overall. Sessions will be co-facilitated by the project team and library staff, allowing the team to observe facilitation dynamics alongside platform usability. All participants will be compensated for their time. Data from testing sessions will be shared with the CAG and used to drive subsequent development cycles.

Activity 2.3. Librarian Training Curriculum Co-Design (Months 12-18)

Concurrent with platform development, the project team will work with Western Allegheny Community Library staff and CAG members with library backgrounds to co-design a modular training curriculum for public library staff. The curriculum is not a standalone deliverable but a support structure whose purpose under Objective 1.1 is to ensure library staff can facilitate patron learning about government AI systems.

It is informed by three complementary intellectual frameworks. First, Long and Magerko (2020) and Lo (2024) identify core competencies of AI literacy, including the ability to recognize AI, understand its basic logic, evaluate its outputs, and consider its societal implications. These competencies define *what* the curriculum should help librarians understand so they can facilitate meaningful sessions with patrons. Second, Ridley and Pawlick-Potts (2021) argue that libraries occupy a unique position in the algorithmic literacy landscape because they already serve as sites for information literacy, and that

algorithmic literacy can be integrated into this existing role rather than treated as a new domain. This argument shapes *how* the curriculum is structured. We build on what librarians already know how to do (help people evaluate information, ask good questions, navigate complex systems) rather than asking them to become technical experts. Third, in our prior work, we have shown that meaningful engagement with AI involves not only technical comprehension but socio-ethical reasoning, the capacity to think about how AI systems affect people and communities in practice (Kong et al., 2025). This adds a third dimension where we need to go beyond understanding how AI works and knowing how to facilitate a session. Librarians need the capacity to help patrons engage with the values and choices embedded in AI systems, questions about what data should be used, who benefits and who bears risk, and what recourse exists when a system produces harmful outcomes.

The curriculum addresses three domains grounded in these frameworks: foundational concepts in how AI systems process and weight information (drawing on Long & Magerko and Lo); facilitation skills for leading community sessions that draw on participants' own questions and practical knowledge (drawing on Ridley & Pawlick-Potts); and guidance for helping patrons engage with the socio-ethical dimensions of AI in public services (drawing on Kong et al.). Building on our prior work (Gautam et al., 2024), we will design the curriculum to fit into how public libraries already schedule, staff, and deliver programming. It will be modular so that libraries with different staffing levels, programming schedules, and community contexts can adapt it to their own realities.

A pilot testing with six to eight library staff across four sessions, with structured feedback collected after each, will ensure the curriculum reflects both staff readiness and the realities of community engagement that emerge from Activity 2.2.

Phase 3: Finalization, Documentation Design, and Dissemination Planning (Months 16-24)

Activity 3.1. Community Knowledge Base Design (Months 16-22)

Building on what emerged across community listening sessions and testing, the project team and CAG will design the documentation practices and infrastructure for the community knowledge base component of CAS. Public libraries have long served as what the literature describes as civic information intermediaries: institutions that help communities access, interpret, and act on information about the public systems that affect their lives (Robinson & Mather, 2017; Gascó-Hernandez et al., 2022; Oh & Mon, 2024). The Co-PI's IMLS-funded Civic Switchboard project has demonstrated this role in practice, developing guides and educational materials that help libraries serve as civic data intermediaries in partnership with local data organizations (Mattern et al., 2026; Mattern & Rapchak, 2022). This design is informed by work on community-inclusive data ecosystems, which emphasizes the need for educational supports, partnership structures, and mechanisms that bring community voices forward (Yoon & Copeland, 2020).

Key design questions include what forms of documentation patrons find meaningful and practicable (text, audio, image, or combinations), what governance structures ensure community members retain ownership of their contributions, and what technical infrastructure supports a knowledge base that is accessible, searchable, and useful across different library contexts. In this space, participatory archives scholarship argues for the value of collaborative documenting, storytelling, and memory-making for building more responsive and community-grounded knowledge infrastructures (Bushey, 2023). The project will not publish raw community narratives or documentation session transcripts; public outputs will use aggregated themes, anonymized examples, and participant consent procedures reviewed through IRB.

This activity addresses the fourth planning question and positions the knowledge base as a practical resource for strengthening the connection between communities and the public agencies that serve them.

Activity 3.2. Planning Report and Implementation Proposal Development (Months 22-24)

All planning findings will be synthesized into a comprehensive planning report covering the literature and landscape review, librarian needs assessment findings, community listening findings, AlgoHub design documentation and testing results, the pilot-tested training curriculum, and the community knowledge base design. This report will be reviewed by CAG members and library staff before finalization.

The report will also include a candid assessment of what worked, what did not, and what the project team would do differently, providing honest and useful guidance for other library systems considering similar initiatives. The final months will be used to develop a full National Implementation grant proposal building directly on these findings.

Activity 3.3. Dissemination (Months 18-24)

The project team will present findings and share planning tools through peer-reviewed publications, conference presentations (e.g., ALA Annual, ALISE, iConference, CHI), a practitioner webinar co-hosted with the Western Allegheny Community Library, and direct outreach to library systems in other regions. All tools, curricula, and the planning report will be published under open licenses through the project website. The AlgoHub platform will be released as open-source software (please see the attached Digital Products Plan).

Dissemination is intended to make the project's planning tools, findings, and design decisions usable by other libraries during and after the grant period. Because all materials are openly licensed, other library systems will be able to review, adapt, and build on this work without licensing barriers.

Project Results

By the end, this project will have produced **five concrete deliverables**, that answer the four planning questions and establish the evidence base for a future implementation proposal:

1. Librarian Needs Assessment Report. This report will document the knowledge gaps, facilitation barriers, and training structure preferences of public library staff as assessed against existing AI literacy frameworks. It will include survey and interview data from staff at Western Allegheny Community Library, analysis of how public library staff preparedness compares to the patterns identified among academic library employees, and specific recommendations for curriculum design grounded in documented staff realities. We will make this publicly available through the project website. Other libraries will be able to use both the assessment instrument and the findings to benchmark their own staff preparedness.

2. Community Listening Findings Report. This report will synthesize what patrons across community listening sessions identified as their most pressing questions about government AI systems, what they want to understand about the services that affect their economic lives, and what they want a library-based platform to help them explore. It will also document what forms of engagement patrons found most and least productive, providing practical guidance for other libraries designing community input processes around AI. The report addresses Planning Question 2. Because the listening sessions are co-designed with library staff and position patrons as contributors rather than subjects, the findings will reflect the perspectives of people whose knowledge of government services comes from direct practical experience, a form of expertise that is rarely captured in the AI literacy literature.

3. Iteratively Developed AlgoHub Platform. The planning period will produce a medium-to-high fidelity version of AlgoHub with accompanying design documentation. "Medium-to-high fidelity" means the simulation is functional, interactive, and tested with real users, but may not include every feature or government system envisioned for full implementation. The design documentation will specify what fidelity levels and interaction affordances were found to support learning without oversimplification, what design decisions were made and why, and how patrons and library staff responded to different levels of technical detail. This directly addresses Planning Question 3 and generates design findings that may inform future simulation-based AI literacy work in library and related public-learning settings. AlgoHub's codebase will be released as open-source software so that other institutions can adopt, adapt, or extend it.

4. Pilot-Tested Librarian Training Curriculum. The curriculum will be a modular facilitation training package in three domains: foundational AI concepts relevant to government decision systems, facilitation skills for leading community sessions grounded in patrons' own questions and experiences, and guidance for engaging with the socio-ethical dimensions of AI in public services. "Pilot-tested" means the curriculum has been delivered in multiple sessions with library staff, revised based on structured feedback after each session, and assessed for whether staff report increased confidence and competence in the specific areas identified as gaps in the needs assessment. The curriculum addresses Planning Questions 1 and 3 and is designed to be modular so that libraries with different staffing levels, programming capacities, and community contexts can adapt it to their own realities.

5. Community Knowledge Base Design. This deliverable specifies recommended practices, governance structures, and technical specifications for a community knowledge base integrated with the AlgoHub platform. It addresses Planning Question 4. The design will include specifications for what documentation formats patrons found meaningful, what governance structures ensure community members retain ownership of their contributions, and what technical infrastructure

is needed to make the knowledge base accessible and useful across different library contexts. Because this is a planning grant, the deliverable is a design document and prototype specification so that we can prepare a larger implementation phase.

Formative Assessment

Assessment is built into every phase of the project rather than conducted as a separate, standalone evaluation. Each community listening session, each co-design testing cycle, and each curriculum pilot includes structured feedback collection from participants. The CAG will review all major deliverables against the four planning questions at each quarterly meeting, providing ongoing quality assurance grounded in both professional knowledge and lived experience. This approach treats assessment as part of the co-design process itself where feedback from patrons and library staff will actively shape what the project produces. Each AlgoHub testing session will include a brief pre/post check on whether participants can identify key inputs, decision thresholds, and points of human discretion in the simulated system. Each curriculum pilot will assess whether staff can explain the system in plain language and guide a discussion of its limits and uncertainties.

For each deliverable, the project team will **assess both product quality and process quality**. Product quality asks whether the deliverable is complete, accurate, understandable, and usable by its intended audience: does the needs assessment capture the right competencies, does AlgoHub actually help patrons understand how a government AI system works, does the curriculum leave librarians feeling more prepared? Process quality asks whether the co-design methodology functioned as intended: did patrons and library staff have genuine influence over design decisions, did the incremental trust-based expansion model work, and what would the project team do differently? Both dimensions will be documented in the final planning report.

The planning report will include a dedicated assessment section documenting how participant and CAG feedback influenced each deliverable across the full project period. This section will be candid about what worked, what did not, and what the project learned about co-designing AI literacy tools with libraries and their patrons.

National Relevance and Adaptation

Although this project begins with a pilot partnership at the Western Allegheny Community Library and limited regional expansion, its planning purpose is to determine which components of the model appear adaptable to other public library contexts and which remain locally specific. The planning period will serve as a testbed for developing and refining a replicable planning outputs and design specifications. We will **disseminate these through national and regional library networks**, including the American Library Association, the Digital Public Library of America, and state library systems that coordinate professional development for public library staff. Project outputs will also be shared through practitioner-oriented venues such as webinars, conference workshops, and open repositories to enable libraries to experiment with algorithm literacy programming in their own communities. Also, insights from the planning grant will inform a subsequent multi-site implementation effort involving additional public libraries in diverse geographic and demographic contexts.

The NLG program expects projects to “create models, tools, research findings, services, and partnerships that can be widely used, adapted, scaled, or replicated to extend and leverage the benefits of federal investment.” This project meets that standard through two mechanisms. First, every deliverable is openly licensed and freely available. The AlgoHub platform will be released as open-source software. The needs assessment instrument, the curriculum, the community listening protocol, and the knowledge base design specifications will all be published under open licenses. Any library system in the country can adopt, adapt, or build on this work without cost and without needing to contact the project team for permission. Second, the planning methodology itself is documented as a replicable model. The combination of librarian needs assessment, community listening, iterative co-design, and curriculum piloting is not specific to Pittsburgh or to the Allegheny County context. It is a planning approach that other library systems may adapt and test in their own contexts, adapted to its own community, staffing, and programming realities.

The planning grant will generate the evidence and specifications needed for a future national implementation, where tested tools and curriculum would be deployed through partnerships with library systems in multiple states, building on the openly available materials from the planning phase and the interest generated through dissemination. The incremental growth model that guides the planning period, demonstrating value first, then inviting engagement, will continue to guide implementation.

Schedule of Completion

Advancing AI Literacy Through Public Libraries: The Civic Algorithmic Sandbox

Year 1

Major Activities	MONTHS											
	1	2	3	4	5	6	7	8	9	10	11	12
Phase 1: Community Listening and Needs Assessment (Months 1-6)												
1.1 Project Launch and Partnership Governance												
1.2 Librarian Needs Assessment												
1.3 Community Listening Sessions												
Phase 2: Iterative Co-Design, Development, and Testing (Months 7-18)												
2.1 AlgoHub Co-Design and Iterative Development												
2.2 Expanded Testing and Feedback												
2.3 Librarian Training Curriculum Co-Design												
Community Advisory Group Review (throughout the project duration)												
Community Advisory Group Quarterly Review			Q			Q			Q			Q

Note: Q marks quarterly Community Advisory Group (CAG) review.

Year 2

*Months 1-12 below correspond to project Months 13-24.
Activities continuing from Year 1 are shown in their Year 2 portion only.*

Major Activity	MONTHS											
	1	2	3	4	5	6	7	8	9	10	11	12
[Continuing] Phase 2: Iterative Co-Design, Development, and Testing (Months 7-18)												
2.1 AlgoHub Co-Design and Iterative Development												
2.2 Expanded Testing and Feedback												
2.3 Librarian Training Curriculum Co-Design												
Phase 3: Finalization, Documentation Design, and Dissemination Planning (Months 16-24)												
3.1 Community Knowledge Base Design												
3.2 Planning Report and Implementation Proposal												
3.3 Dissemination												
Community Advisory Group Review (throughout the project duration)												
Community Advisory Group Quarterly Review			Q			Q			Q			Q

Note: Q marks quarterly Community Advisory Group (CAG) review.