aiX Faculty Fellows Program

Final Report — Spring 2026

Columbia University

May 2026

Tian Zheng

Professor of Statistics, Columbia University
Faculty Lead, aiX Faculty Fellows Program

Executive Summary

The Spring 2026 aiX Faculty Fellows Program brought together 20 faculty fellows and 15 student interns from 14 schools and departments across Columbia University to design and build AI-integrated education tools grounded in disciplinary expertise. Over the course of a single semester, 17 distinct projects were developed. Thirteen teams presented at a public showcase on May 6, 2026. Five projects reached classroom-ready status; seven are in active beta testing. Twelve of 15 surveyed fellows plan to continue their work into the summer or next academic year.

The program operates as part of the Empire AI Consortium, a multi-institution initiative across New York State. Through monthly meetings with peer institutions, the program faculty lead confirmed that the challenges identified at Columbia—cost, data privacy, and the need to update institutional practices—are shared across higher education. These discussions also reinforced a consensus that faculty community building and faculty governance of AI in education are essential priorities.

The program generated clear evidence of faculty confidence growth, meaningful cross-disciplinary community, and pedagogical innovation at scale. External recognition has started to take place: one fellow was invited to present her aiX project at a national conference, another received support from his school to develop a pilot proposal, and a third secured additional funding from her school to move her project forward. The program's highly hands-on model, from weekly studio hours with interns, monthly faculty check-ins, to individualized mentoring, proved critical to sustaining momentum across projects at very different stages of development.

The program also surfaced consistent, addressable challenges: the need for dedicated intern support at a one-to-one ratio, stronger structural onboarding, institutional policy clarity on AI tool use and funding, and formal recognition for faculty participation.

Key Recommendations

Based on faculty and intern feedback, the following priorities emerged for sustaining and strengthening the program (see Section 6 for full recommendations):

- Prioritize a one-to-one intern-to-faculty ratio and introduce phased project milestones.
- Provide foundational AI literacy resources and structured technical onboarding.
- Advocate for course-load credit or formal recognition for participating faculty.

- Clarify university AI policy guidance and establish a funding pipeline for API and computation costs.
- Develop shared data governance and IRB guidance templates for pilot projects.

This report provides a full account of the program's origins, design, activities, outcomes, and recommendations. It is intended to serve as both a record of what was accomplished and a proposal for this work to continue.

Table of Contents

1. Background: Why a Faculty Fellowship for AI Education?

1.1 The Challenge

Artificial intelligence is reshaping how knowledge is produced, analyzed, and communicated across every academic discipline. Students encounter AI tools in internships, labs, and industry settings well before faculty develop frameworks for engaging these technologies critically and responsibly. Yet the educational structures of higher education from curricula to pedagogy to assessment have been slow to respond in a coherent way. There is a growing tension between how disciplines have traditionally defined expertise and how that expertise is evolving in practice. Every field represented in this fellowship—from law to nursing to language instruction to chemical engineering—faces discipline-specific questions that AI raises: What counts as expertise when AI can generate drafts, analyze images, or simulate conversations? What should be automated versus practiced? How is professional judgment cultivated when the tools themselves are changing?

These are not questions that can be answered by any single office, technology platform, or training program. They require the deep subject-matter understanding, intellectual curiosity, and sustained inquiry that define faculty work.

1.2 The Vision: Curriculum as the Lever for Change

The aiX Faculty Fellowship is grounded in a specific intellectual argument: that meaningful curricular reform, driven by faculty expertise and anchored in disciplinary purpose, is the central lever for institutional change in the age of AI.

Discussions about AI in higher education often conflate several distinct conversations: AI as a pedagogical tool (tutoring systems, automated grading), AI as an operational concern (admissions, advising), AI as a curricular question (what to teach and how disciplines are shifting), and AI as a societal challenge (preparing graduates for work and citizenship). Without clarity about which conversation is being had, stakeholders talk past one another.

The aiX program focuses deliberately on the curricular and disciplinary layer. It treats pedagogy as a site of faculty inquiry and creates the conditions for faculty to lead the intellectual work of integrating AI into their fields. The premise is that *what faculty need to teach is also what they need to learn*, and that the process of designing AI-integrated learning experiences is itself a form of disciplinary renewal.

1.3 Why Faculty-Led?

Faculty sit at the intersection of research, curriculum, and pedagogy. They are the stewards of disciplinary knowledge. Throughout the history of higher education, meaningful curricular change, from the rise of computational methods to the integration of data science, has emerged from faculty grappling with new ideas, methods, and societal demands.

AI raises questions that are fundamentally disciplinary: questions about evidence, judgment, automation, and professional identity that no external expert can answer on behalf of a field. The aiX Faculty Fellowship is designed to give faculty the structured space, technical support, and peer community they need to take intellectual ownership of these questions, i.e., to approach AI not as users of innovation, but as leaders of educational change.

This approach proved, in practice, more engaging and more grounded in expertise than prescriptive models of AI professional development. Faculty did not adopt recommended tools; they built pedagogical innovations rooted in their own intellectual traditions and addressed problems they cared about.

1.4 Program Origins

The aiX Faculty Fellows Program grew out of a recognition, during the 2024–2025 academic year, that the pace of change in AI had outstripped any one person's or any one discipline's capacity to respond alone. The program's faculty lead, Tian Zheng, Professor of Statistics, initially planned to use her sabbatical year to explore what artificial intelligence means for statistical education. But conversations with colleagues across fields revealed a shared sense of urgency and uncertainty. Faculty across the university were trying to make sense of how AI was reshaping their disciplines, their teaching, and what students need to learn—but there was no shared space to work through these questions together.

The fellowship emerged from that recognition at a DSI retreat in June 2025. A beta version of the program concept was developed during Fall 2025 in partnership with the Office of the Provost, the Center for Teaching and Learning (CTL), and the Data Science Institute (DSI). A parallel effort to secure external funding through a FIPSE (Fund for the Improvement of Postsecondary Education) grant application laid out the program's rationale, design, and resource model in detail. Although that application was not funded, the proposal crystallized the program's intellectual framework and informed the design of the Spring 2026 pilot.

CTL has been an invaluable partner throughout, providing pedagogical expertise and event support. The fellowship is designed as a connecting layer—linking urgently needed faculty development with the university's broader AI strategy—and as a model for how faculty-led programs can complement and strengthen the work of centers for teaching and learning. CUIT offered consultation as part of the kick-off event and through their office hours and emails. A&S central, the Vice Dean of AI initiatives of A&S, and the Department of Statistics provided additional administrative support, in terms of project management, intern hiring and event logistics.

1.5 The Empire AI Consortium

The aiX Faculty Fellowship operates within a broader statewide effort. Columbia is a member of the Empire AI Consortium—a partnership of 11 public and private academic institutions across New York State, including CUNY, SUNY, Cornell, NYU, RPI, the University of Rochester, and the University at Buffalo. In 2025, the consortium received more than $1 million from Google.org to advance the effective integration of artificial intelligence into higher education. Unlike traditional educational research conducted within a single institution, Empire AI enables member campuses to test, evaluate, and share findings rapidly across the consortium, with results—including digital tools, curricula, and teaching materials—shared widely to ensure that best practices in AI pedagogy are adopted and scaled.

As the aiX program faculty lead, Professor Zheng attended monthly meetings with AI education leaders from across the consortium throughout the 2025–2026 academic year. These conversations produced a strong and growing consensus on several points that directly inform this report's recommendations. First, faculty community building is one of the most effective models for AI education development. Second, faculty must be involved in evaluating and building AI tools, not merely adopting them; this involvement serves both curricular oversight and faculty's own development in AI literacy across disciplines. Third, the challenges that surfaced at Columbia are not unique: the cost of AI tools, data privacy concerns, and the urgent need to update institutional policies and governance structures to keep pace with technological change were shared by every project in the consortium.

Another observation from these meetings was the diversity of institutional homes for AI education initiatives across the consortium. Some programs were led by centers for teaching and learning, others by data science or AI institutes, and still others by leaders within the provost's or president's office. This variety reflects the fact that AI in education does not fit neatly into any single administrative structure—and suggests that

the most effective approaches will require coordination across institutional boundaries. Among the peer programs, NYU's faculty learning communities model was particularly inspiring and closely aligned with the aiX Fellowship's emphasis on faculty-led, discipline-grounded collaboration.

The consortium's collective experience reinforces a central conclusion: more faculty involvement and governance in AI education is essential. The aiX Fellowship's faculty-led, discipline-grounded approach is well-aligned with the direction emerging from this statewide effort, and Columbia is well-positioned to serve as a model within it.

2. Program Overview

2.1 Design and Structure

The aiX Faculty Fellows Program ran from January through early May 2026. The program was structured around three elements: a kick-off event, sustained cohort engagement, and a culminating public showcase.

The full-day kick-off event on January 30, 2026 introduced faculty and interns to the program's design, Columbia's existing AI resources, and design frameworks for curricular reform. Extended introductions in which each fellow shared their project vision and plan were widely cited in post-bootcamp reflections as the single most valuable part of the day.

Bi-weekly cohort meetings throughout the spring semester provided a forum for peer-led tool demonstrations, pedagogical discussions, and project troubleshooting. Student interns joined faculty fellows in cluster groups organized around shared disciplinary themes or pedagogical goals. These sessions were consistently cited as among the most valuable elements of the program.

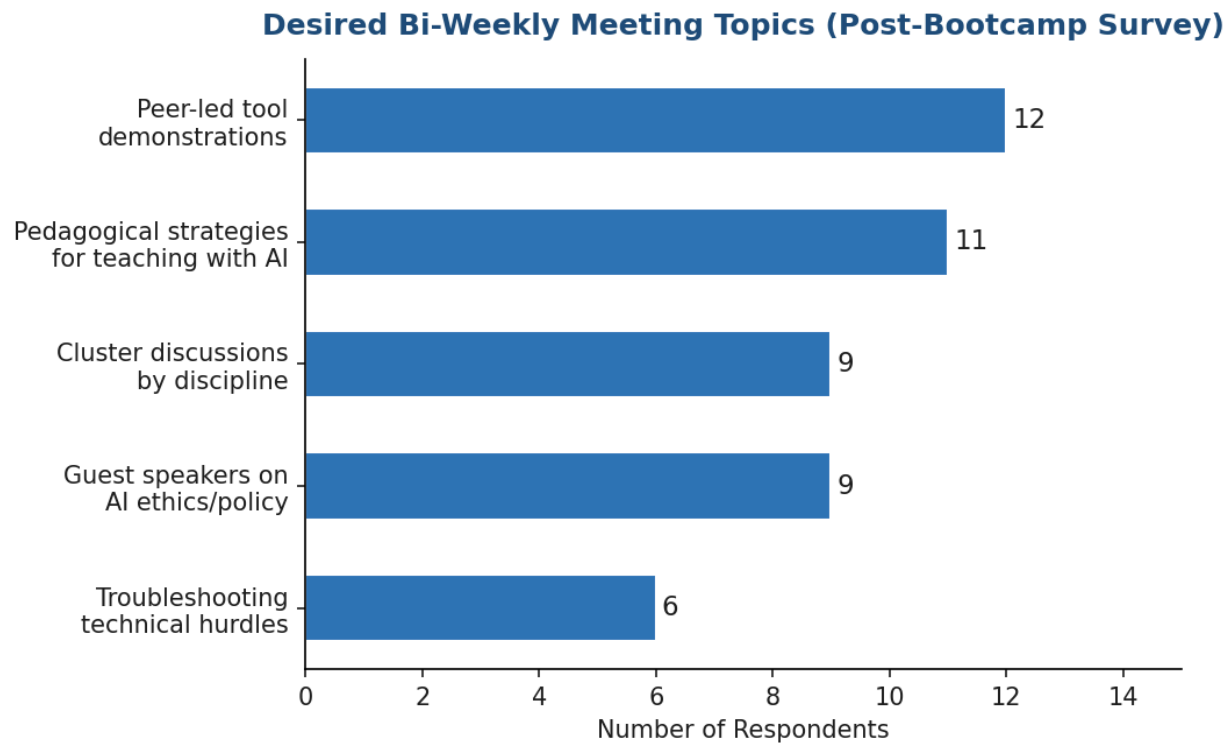

Figure 1. Faculty-requested topics for bi-weekly meetings (post-bootcamp survey, n=16, multi-select)

The final public showcase on May 6, 2026, attended by university leadership including representatives from the Provost's office, school deans, and department chairs, served as the semester's culminating event, with 13 teams presenting live demonstrations of their work. Recordings may be shared upon request.

Between the bootcamp and the showcase, the program maintained a rhythm of structured touchpoints designed to keep projects on track while respecting faculty time. Weekly studio hours between the program's faculty lead and the student interns provided a regular forum for project updates, technical guidance, and mentoring. It allows the faculty lead to stay closely engaged with every project's progress without requiring additional time from faculty.

Communication was sustained through weekly email updates from the program's faculty lead to all participants, sharing information learned from meetings, reminders, and relevant developments from the broader AI education landscape. Slack served as a secondary communication platform; while it was used primarily for private one-on-one messaging between faculty and interns, it also hosted some group discussions and quick troubleshooting exchanges.

The program's operations depended on a small but essential support infrastructure. Ms. Katie Roehrick, the part-time program coordinator, was instrumental in keeping program activities on schedule—managing event logistics, coordinating with participants, and working with the A&S HR department on the hiring and onboarding of student interns. The Department of Statistics provided valuable administrative support for program events, including room reservations, catering, and logistical coordination. The program's faculty lead also handled manual review and approval of all intern time reports, a small but cumulative administrative responsibility that underscores the need for dedicated administrative support as the program scales.

The program also attracted investment from individual schools. Recognizing the importance of AI education in legal training, the program admitted two faculty fellows from the Law School. Ben Liebman, Vice Dean for Intellectual Life at Columbia Law School, provided additional funding to support the hiring of a dedicated intern for the law projects—an early signal that schools across the university see value in this kind of faculty-led AI development and are willing to invest in it.

2.2 Cohort at a Glance

Faculty fellows	20 across 14 schools and departments
Student interns	15 (undergraduate and graduate)
Projects developed	17
Presented at showcase	13
Classroom-ready	5
In active beta	7
Plan to continue	12 of 15 surveyed faculty
Urgency rating	All 20 rated AI reform at least "moderately urgent"

2.3 The Faculty

The 20 faculty fellows represented a wide range of titles, disciplines, and levels of prior AI experience. The cohort included tenured professors, lecturers, and clinical faculty from both the Morningside and CUIMC campuses, bringing together expertise in law, public health, architecture, language instruction, clinical medicine, engineering, psychology, and business. Faculty entered the program with strong pedagogical vision but varied widely in technical confidence—a diversity that the fellowship was designed to accommodate.

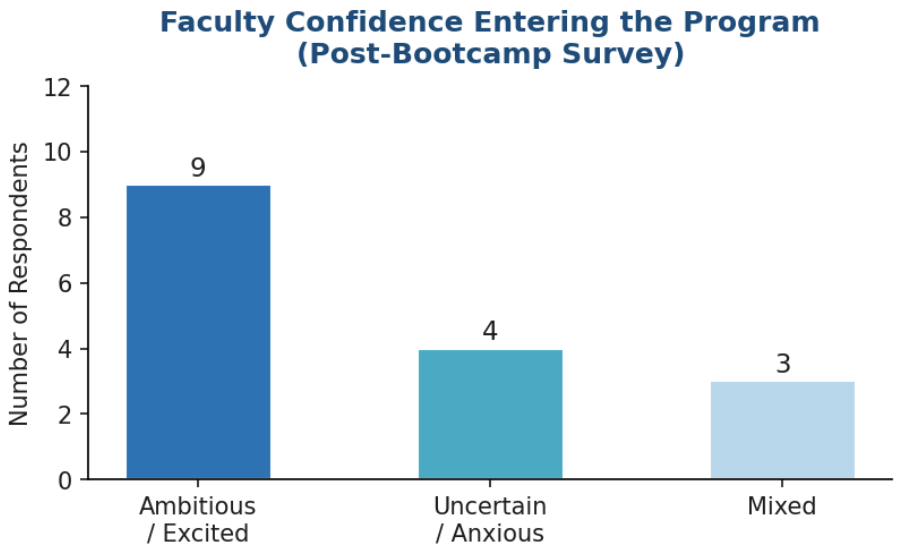

Figure 2. Faculty self-reported confidence level entering the program (post-bootcamp survey, n=16)

2.4 The Student Interns

Fifteen student interns joined the program, drawn primarily from SEAS graduate and undergraduate students in Computer Science and Business Analytics, alongside students from Teachers College, Columbia College, GSAS/Statistics, SIPA, and GSAPP. Interns served as technical builders, researchers, data analysts, and thought partners. The dominant collaboration pattern was a "Translator" model—faculty provided academic vision and learning objectives while interns translated those into technical solutions—though several pairs developed a more equal "Co-Pilot" dynamic.

Many of the interns brought prior industry experience evaluating AI tools—a practical perspective that complemented the academic orientation of the fellowship. The mix of disciplines among the intern cohort proved to be one of the program's unexpected strengths. Students from Statistics brought methodological rigor and an instinct for evaluating claims about AI performance. Those from Teachers College brought educational theory and a learner-centered design sensibility. Interns from GSAPP's digital design program brought interface and interaction expertise. Together with computer science and engineering students, this disciplinary diversity created a collective intelligence within the design studio that no single department could have assembled on its own, and that enriched the projects in ways neither faculty nor interns fully anticipated at the outset.

3. The Projects

The 17 projects developed during the Spring 2026 semester span an extraordinary range of disciplines, pedagogical approaches, and technical implementations. What unites them is a shared commitment to building AI tools that serve genuine learning objectives, designed by the faculty who understand those objectives most deeply. The following descriptions are organized by thematic cluster.

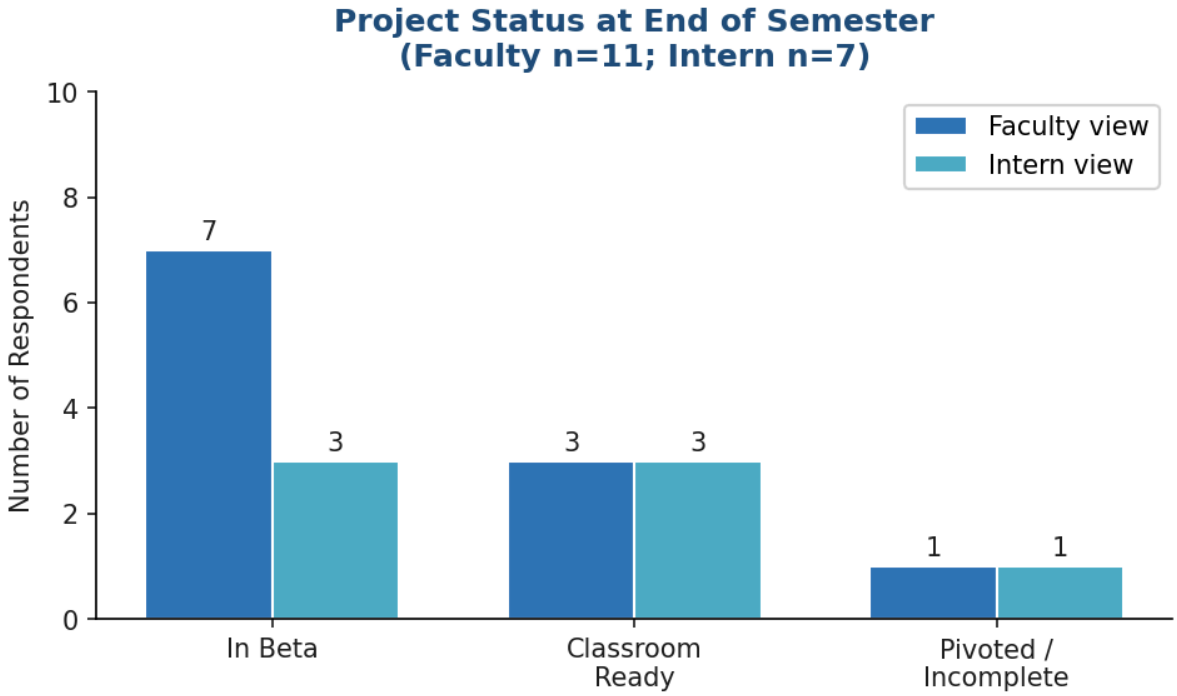

Figure 3. Project status at end of semester, as reported by faculty and interns (faculty n=11; intern n=7)

3.1 Legal Education

Talia Gillis — Law School: Contracts AI Chatbot, Law & AI Assignments, and Socratic Chatbot

Professor Gillis's team pursued three simultaneous goals: a RAG-powered contract analysis chatbot for first-year Contracts students, data-driven assignments for a new Law and AI survey course, and a Socratic chatbot prototype for practicing the cold-call format central to legal pedagogy. The contract chatbot allows students to collectively upload real-world contracts and query them using natural language—asking, for example, which contracts contain arbitration clauses. The Law and AI assignments use publicly available datasets, including the COMPAS recidivism tool, to explore how technical decisions in AI pipelines affect fairness outcomes, without requiring prior coding knowledge. The Socratic chatbot remains a work in progress; the team

recognized that productive Socratic dialogue may require debating open questions rather than testing factual recall.

Status: In Beta. Intern: Henrique Brito Sampaio (Computer Engineering, SEAS).

Rebecca Wexler — Law School: LLM Parameter Exploration Tool

Professor Wexler proposed a tool that would allow users to explore and toggle parameters of a large language model—a pedagogically rich concept for a law course examining how technical decisions shape AI outputs. Her project was delayed for most of the semester by an unresolved question: how to configure or train AI tools using copyrighted content.

Status: Original project redirected; new cross-disciplinary collaboration with Statistics underway. Intern: Jenny Leana Fotso Ngompe (Computer Science, SEAS).

3.2 Health Sciences

Natalya Pasklinsky — School of Nursing: Global Nursing Virtual Instructor

Professor Pasklinsky's team built a mobile-first platform to support nursing students on international clinical placements. The centerpiece is an AI virtual instructor avatar, built with HeyGen, that guides students through culturally competent care specific to their clinical region. Italy was the first fully completed module. The avatar can be configured to play any role—tutor, clinical preceptor, or regional cultural guide who helps students navigate their surroundings, answer cultural questions, and access local resources. The platform also integrates emergency contacts, regional reference documents, and a searchable resource library. The project addresses a real gap: students on international placements often lack access to the on-site mentorship available in domestic rotations.

Status: In Beta. Intern: Khushi Panjabi (Business Analytics, SEAS/CBS).

Sharon Perelman & Zohaib Ahmad — Dental Medicine & Medicine: AI Radiology Tutor

This cross-disciplinary team built a dual-track AI radiology tutor covering chest X-rays (medicine) and dental panoramic imaging (dentistry). The tool uses a Socratic feedback loop: students assess an image and record their findings before any AI input, then engage in guided back-and-forth that surfaces missed findings without simply providing answers. A common memory layer tracks student errors across sessions, allowing the system to resurface similar cases for targeted practice. The dental module includes a drag-and-drop exercise for correctly mounting X-ray series—a critical clinical skill that is frequently undertaught.

Status: In Beta. Interns: Xixi Chen (Operations Research, SEAS) and one additional volunteer with explainable AI expertise (Arjun Balaji, SIPA).

Dr. Dennis Mitchell, the Dean of the College of Dental Medicine, attended the final showcase and subsequently secured funding from the school to support Professor Perelman's continued work with her intern over the summer. This was a concrete signal of institutional interest generated directly by the fellowship. Professor Ahmad similarly reported that his school has recently secured AI tools related to his project, confirming the timeliness of the project and the growing institutional readiness to invest in faculty-led AI education.

3.3 Language Education

Eric Matheis — French: Oral Skills Practice Tutor

Professor Matheis and intern Ray Tang built a three-module AI oral skills platform: vocabulary (image-based object identification with real-time correction), pronunciation (minimal pair discrimination exercises), and conversation (speech-to-speech dialogue with live grammar and article correction). The conversation module was identified as the highest-value component, integrating all skills and supporting fluency development in a low-stakes environment. Professor Matheis noted that AI imperfection in language tools can sometimes function as a pedagogical feature rather than a bug—students learn to notice and correct errors, including the AI's.

Status: In Beta. Intern: Ray Tang (Computer Science, SEAS).

Professor Matheis has since developed a formal proposal to pilot the AI oral skills agents in Elementary French I and II during the 2026–2027 academic year, with the program's faculty lead's mentorship on configuring the resources needed. The pilot would serve the sections he instructs directly, with a potential future audience of over 300 students across the two courses.

Barbara Spinelli — Italian: AI-Mediated Writing Portfolio

Professor Spinelli's project is grounded in a theory of distributed cognition: when digital tools are integrated into language learning, the cognitive process is shared between human and tool, and the writer shifts from "producer" to "designer." Her course asks students to submit three drafts of each composition—the first using only traditional resources, the second with AI assistance, and the third with a reflective writing portfolio documenting their mediation strategies. Intern Fangyi Lin conducted quantitative and qualitative analysis of survey data, finding significant increases in AI literacy and confidence, and an evolution in how students perceive and use AI tools.

Status: In Beta. Intern: Fangyi Lin (Statistics, GSAS).

Professor Spinelli's experience also illustrates the program's capacity to accommodate faculty whose teaching schedules prevented regular attendance at monthly meetings. Through customized support from the program's faculty lead, Spinelli remained fully engaged and productive throughout the semester. Her work has since gained international recognition: she presented the project at an invitation-only conference at the University for Foreigners of Perugia, organized jointly with the Association of Language Testers in Europe (ALTE), with representatives from EUROCALL, the Council of Europe, Cambridge International Education, and leading Italian scholars. The presentation was well received, and a publication in the conference proceedings is forthcoming.

3.4 Social Science and Policy

Rumela Sen — SIPA: Critical Lens

Professor Sen built Critical Lens—a Socratic AI coaching tool that helps graduate students refine claims, define concepts, and surface unspoken assumptions before writing policy memos, op-eds, or research papers. The tool offers two personas: Mentor and Peer, the latter particularly valued by mid-career professionals. It aims to extend the kind of office-hours dialogue typically available only to the most engaged students, particularly serving international students who face language or cultural barriers to seeking help.

Status: Classroom Ready. Intern: Xinyu Jiao (Digital Design, GSAPP).

Lisa Dale — Climate School: The Tutor

Professor Dale developed The Tutor to address the "tail problem" in large lectures of 200 or more students: a significant minority who either find content too basic or cannot keep up, compounded by diverse linguistic and cultural backgrounds. The tool engages students in a 10-minute oral conversation about assigned readings before each discussion section—not as an assessment, but as a low-stakes warm-up to improve participation quality. A paid pilot study with approximately 40 recruited students was launching at the time of the showcase.

Status: In Beta. Intern: Benjamin Appio (Clinical Psychology, Teachers College).

3.5 STEM Education

Anthony Vanky — GSAPP: Urbanist AI

Professor Vanky's project is among the most pedagogically ambitious in the cohort: a custom platform asking Urban Planning students to build and train their own computer vision dataset through an urbanist lens. Students define categories (including abstract ones like "ennui"), photograph urban spaces, and upload images to a shared, privacy-protected database. The platform runs analytics on the collective dataset to reveal biases in age, gender expression, and skin tone. A dramatic "delete all" button, demonstrated live in class, prompts discussion of the right to be forgotten. Facial blurring and EXIF stripping are built directly into the workflow.

Status: Classroom Ready. Intern: Kartik Kumar Gounder (Computer Science, SEAS).

Alex Urban — Chemical Engineering: AI Coding Companion

Professor Urban's team built an AI coding companion for engineering students working on numerical methods and domain-specific assignments. The tool requires students to articulate their objective, approach, and constraints before accessing any code generation—a "gatekeeper" phase validated by an LLM to ensure meaningful prior engagement with the problem. Students then write validation tests in plain English, which are automatically converted to code. The tool was built using a cost-effective model and cost approximately five cents in API usage during testing.

Status: In Beta. Intern: Arjun Vaidya (Computer Science/ML, SEAS).

Matt Lamb & Teresa Janevic — Epidemiology: Socratize

This team built Socratize (socratize.app)—a platform that converts domain experts' tacit expertise into structured skill files that guide an AI agent in domain-specific tasks. Adversarial agents challenge the expert's knowledge base during authoring, surfacing ambiguities and gaps before the material reaches students. A teaching mode restricts the AI from completing tasks outright, instead scaffolding student reasoning step by step. Intern Tata (Chaipat Tirapongprasert) moved from a behind-schedule start to publishing Socratize as a fully deployed web application, an npm package, and a GitHub repository.

Status: Classroom Ready. Intern: Chaipat Tirapongprasert (Astrophysics, Columbia College).

Chris Chen — Chemical Engineering: LEPC Town Hall Simulation

Professor Chen built an AI-powered roleplay simulation of a Local Emergency Planning Committee (LEPC) public hearing. Students present a proposed chemical plant—placed in a deliberately challenging location adjacent to a school, a fireworks depot, and a retirement community—to a panel of AI-generated community stakeholders with conflicting concerns. By giving each student the opportunity to practice with multiple

character configurations before a live in-class panel, the tool democratizes an experience that previously only some students received.

Status: Classroom Ready. Intern: Khushi Panjabi (Business Analytics, SEAS/CBS).

Jaehyuk Choi — Mathematics of Finance: AI Coding Companion

Professor Choi developed an AI coding companion for students in the Mathematics of Finance master's program, supporting numerical methods and quantitative assignments. The project shares infrastructure with Alex Urban's chemical engineering tool, with both supported by intern Arjun Vaidya. Professor Choi indicated plans to conclude the project by the end of summer, working independently on final refinements.

Status: In Beta. Intern: Arjun Vaidya (Computer Science/ML).

3.6 Psychology

Sarah DeMoya & Patricia Lindemann — Psychology: AI Historical Figures and Patient Case Bots

Lecturers DeMoya and Lindemann developed AI bots embodying historical figures in psychology—Piaget, Milgram, and Skinner—that respond to student prompts through the lens of their own theories. A complementary set of directed patient case simulations supports the psychopathology unit. A second strand, led by Lindemann and intern Deborah Lee, involves developing a new course: Learning in the Age of AI, which interrogates how AI affects cognition, learning, and the practice of education. Lee's contributions to the course syllabus—drawing on philosophy, psychology, and AI ethics—were described by Professor Lindemann as substantive and generative.

Status: Classroom Ready. Intern: Deborah Lee (Psychology in Education, Teachers College).

3.7 Business

Sheena Iyengar — Business School: Think Bigger AI

Professor Iyengar, one of the world's leading researchers on choice and innovation, presented Think Bigger AI—a structured choice-mapping tool that helps students generate innovative ideas by breaking problems into subproblems and drawing strategies from familiar, adjacent, and out-of-box domains. The semester's work added a voice interface ("Scout") to the existing web app, making the tool more accessible and naturalistic. Students reported that while the concept resonated strongly, the choice map interface remained cognitively demanding—a design challenge the team is actively addressing.

Status: In Beta. Intern: Casey Sim (Business)

Angela Lee — Business School: AI-Integrated Business Course

Professor Lee, a Professor of Professional Practice and expert in venture capital and entrepreneurship, registered as a faculty fellow and attended the bootcamp. Her project was not ready for curriculum development. She plans to continue project development over summer.

Status: Planning.

3.8 Epidemiology

Allison Aiello — Mailman School: AI-Integrated Epidemiology Course

Professor Aiello developed an AI-integrated epidemiology course covering LLM literacy, prompt engineering, agentic tools, and responsible AI use in research. Intern Shivangi Kumar contributed significantly to course slide development, built a live Custom GPT tutorial for students, and helped map agentic AI tools to the epidemiology research pipeline. Professor Aiello chose not to present at the showcase but submitted a project status survey indicating plans to continue developing the course. She noted that the fellowship was a significant help as she navigated teaching this new course for the first time.

Status: Classroom ready. Intern: Shivangi Kumar (Computer Science/AI-ML, SEAS).

4. What We Learned

4.1 Faculty Confidence and Perspective Shift

The end-of-semester review (n=12) revealed substantial growth across the cohort. The most common confidence gains were in communicating about AI's disciplinary impact, identifying which tools have genuine pedagogical value versus novelty value, and writing effective prompts. Several faculty noted that before the fellowship they lacked a framework for evaluating AI tools; by the end, they felt equipped to make principled distinctions.

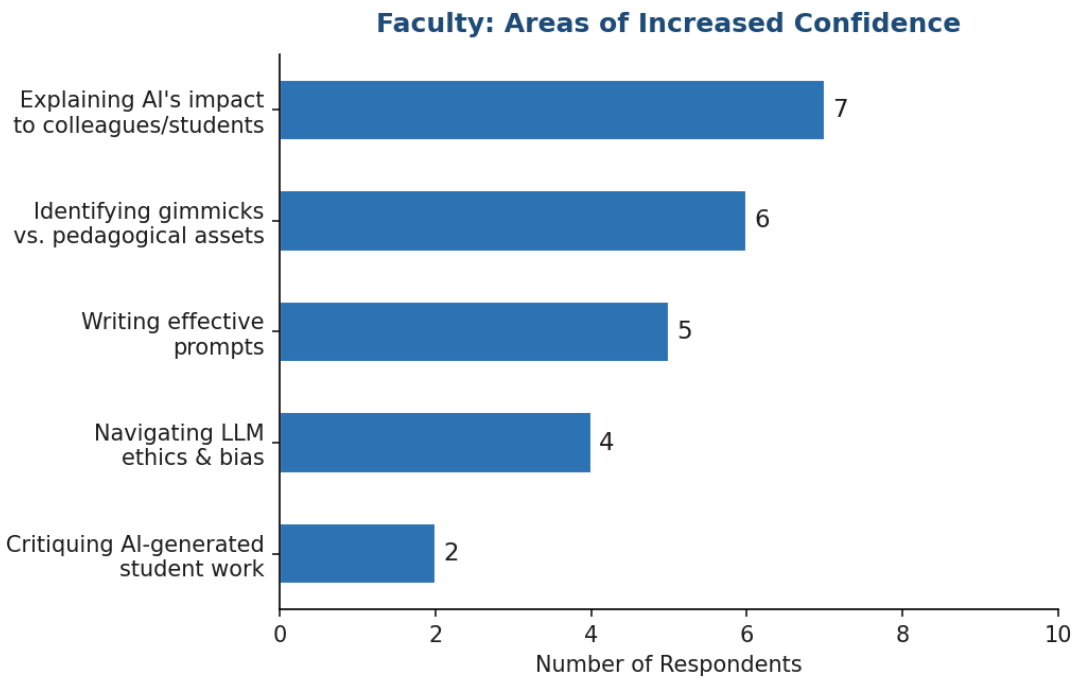

Figure 4. Faculty self-reported areas of increased confidence (end-of-semester review, n=12, multi-select)

Faculty overwhelmingly view AI curricular reform as an urgent necessity, with nine of 12 respondents rating it “very urgent” or “extremely urgent.”

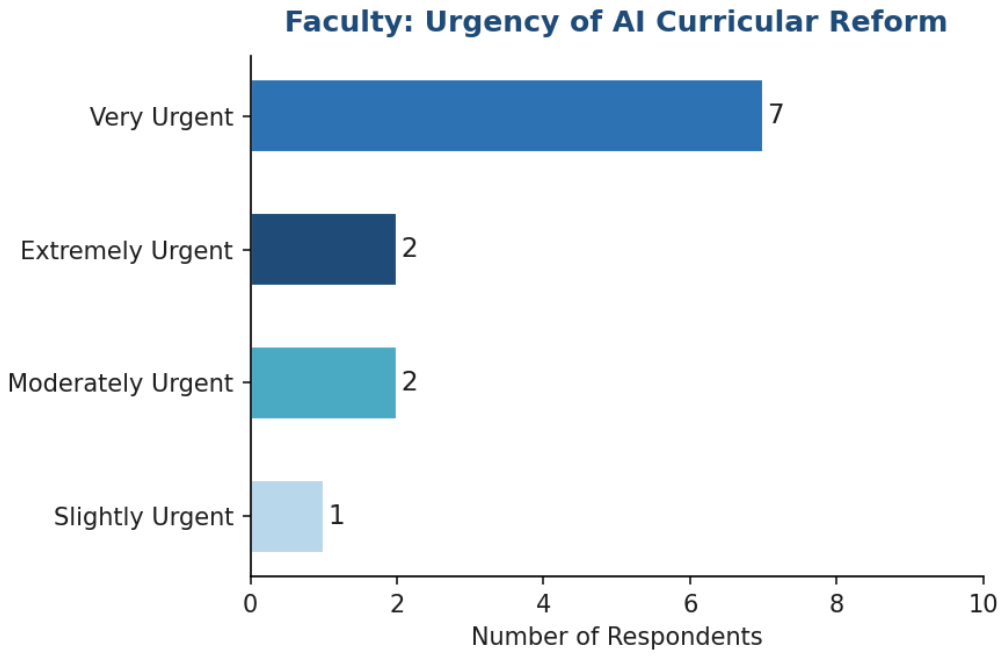

Figure 5. Faculty ratings of urgency for AI curricular reform in their field (end-of-semester review, n=12)

Seven of 11 respondents reported that their view of AI's possibilities expanded significantly during the fellowship. Two had their existing plans validated, and two became more focused on specific, safe use cases. No faculty member reported becoming more skeptical or disengaged—a notable finding given the diversity of starting positions.

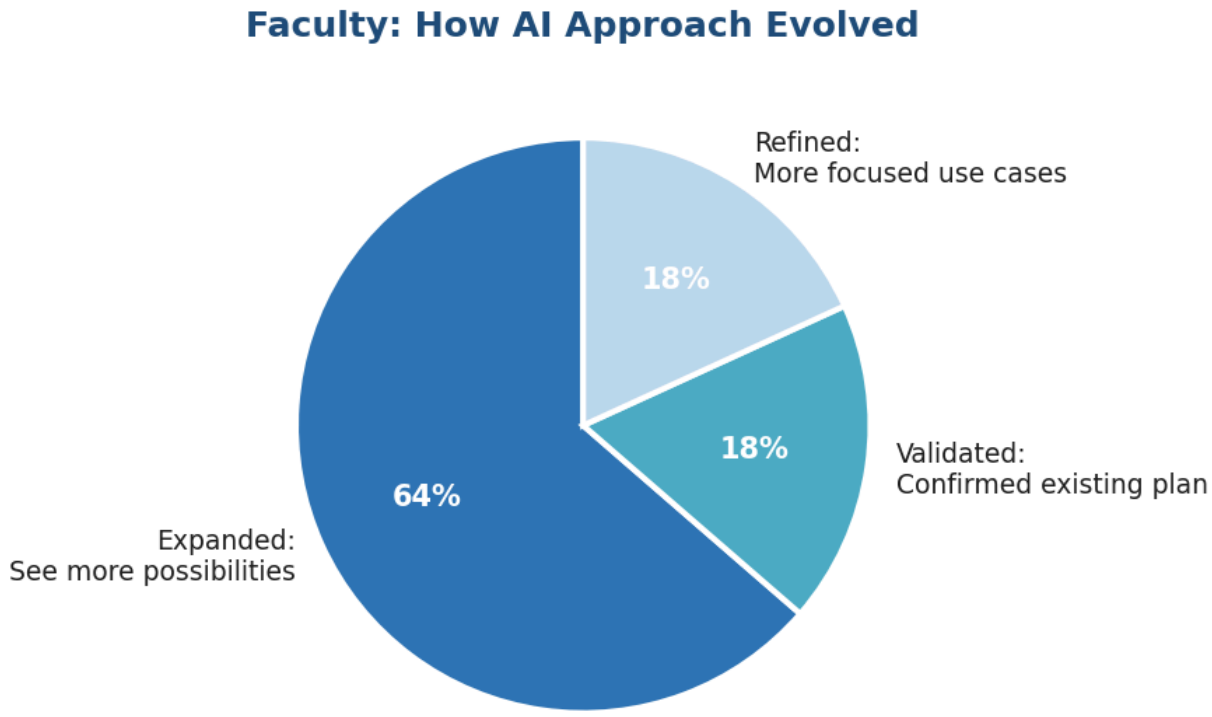

Figure 6. Faculty report on how their approach to AI evolved during the fellowship (n=11)

> *"I learned a tremendous amount about both the opportunities for pedagogical innovation that AI presents and AI's technical and practical limitations, all within an intellectually and personally supportive community."*
>
> *"I now believe that integration of AI is going to lead to breathtaking innovations that universities are uniquely suited to foster."*

4.2 The Intern Experience

The 15 student interns grew technically and professionally in ways many did not anticipate. Beyond technical skills, interns described meaningful shifts in how they think about AI. Several arrived believing AI was straightforwardly powerful and left with a more nuanced view emphasizing the necessity of human oversight. The experience challenged assumptions on both sides: interns expected faculty to be resistant and were

surprised by their genuine curiosity; faculty were surprised by how much interns contributed intellectually, not just technically.

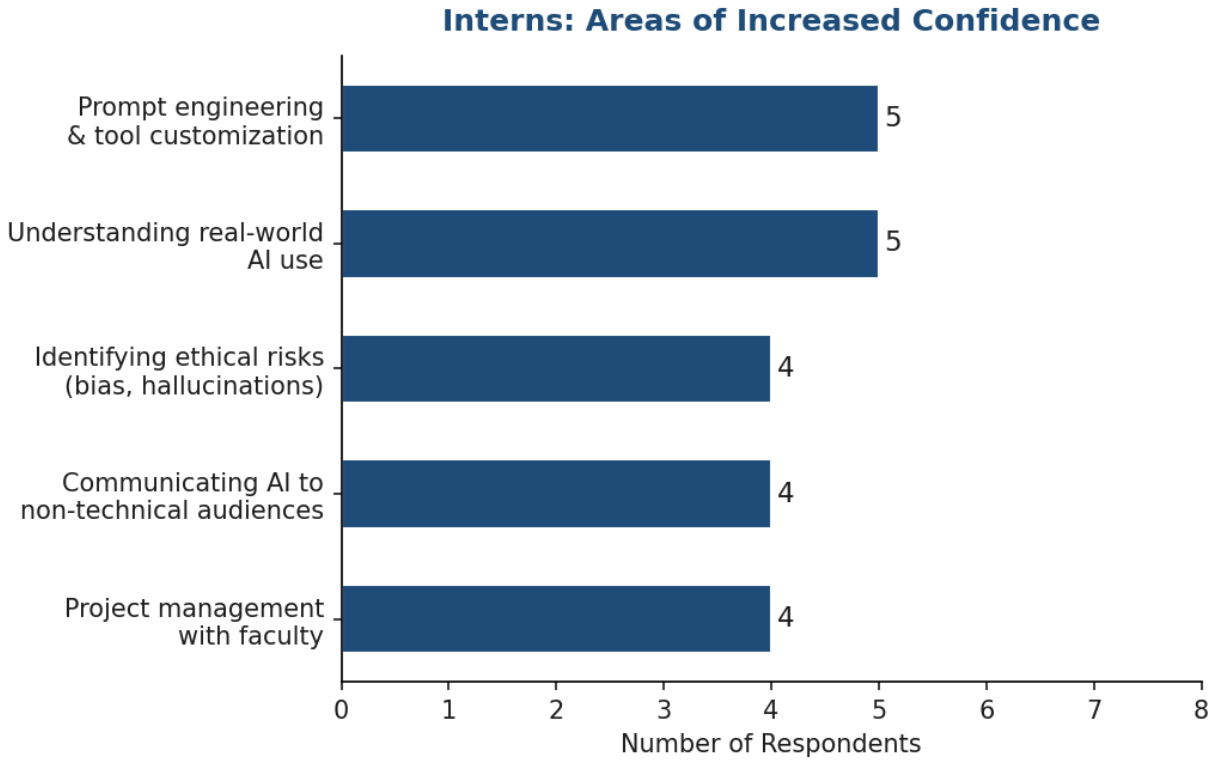

Figure 7. Intern self-reported areas of increased confidence (end-of-semester review, n=7)

> *"You will probably not even know half of what you can do. I came out of this internship learning so much about myself."*
>
> *"The fellowship expanded my thinking about AI in ways that were deeply inspiring and motivating."*

4.3 Cross-Disciplinary Community

The fellowship community received strong ratings: six faculty rated the fellowship "High Impact" for professional connections; five rated it "Moderate Impact." Faculty particularly valued learning what other disciplines were doing and how they were thinking about AI's implications for their fields. The most consistent feedback was that time for cross-disciplinary conversation was valuable but too limited—project work tended to crowd out community-building.

4.4 Technical and Ethical Landscape

The cohort's 17 projects collectively demonstrate the breadth of technical approaches available for AI in education. Projects used models from OpenAI (GPT-4o), Anthropic (Claude), and Google (Gemini), deployed on platforms ranging from Vercel and Supabase to Google Cloud and Streamlit. Socratize's bring-your-own-key architecture stands out as a privacy-forward model that avoids centralizing student data.

No project reported collecting or retaining audio recordings of student interactions. No project reported using student interaction data to train AI models. Multiple projects implemented row-level security, hashed authentication, and encrypted API key storage. FERPA was identified as the primary applicable regulation across most projects.

The copyright and AI tool configuration question that affected Rebecca Wexler's original project highlights the need for clearer institutional guidance. While the fellowship helped redirect her work into a productive cross-disciplinary collaboration, faculty more broadly need clear, official guidance on what AI tools they may use, how they may configure them, and whether copyrighted content can be used in that configuration. Providing this guidance before the next cohort begins would remove a significant source of uncertainty.

4.5 Challenges

The program also surfaced consistent challenges that should be addressed in future iterations. Technical friction—including API costs, account setup, and software bugs—was the most commonly cited remaining challenge. The need for a one-to-one intern-to-faculty ratio was the single most consistent recommendation from both faculty and interns; shared interns experienced divided attention and role ambiguity. Scheduling conflicts, difficulty translating pedagogical goals into technical specifications early in the semester, and waiting for institutional approvals were recurring friction points.

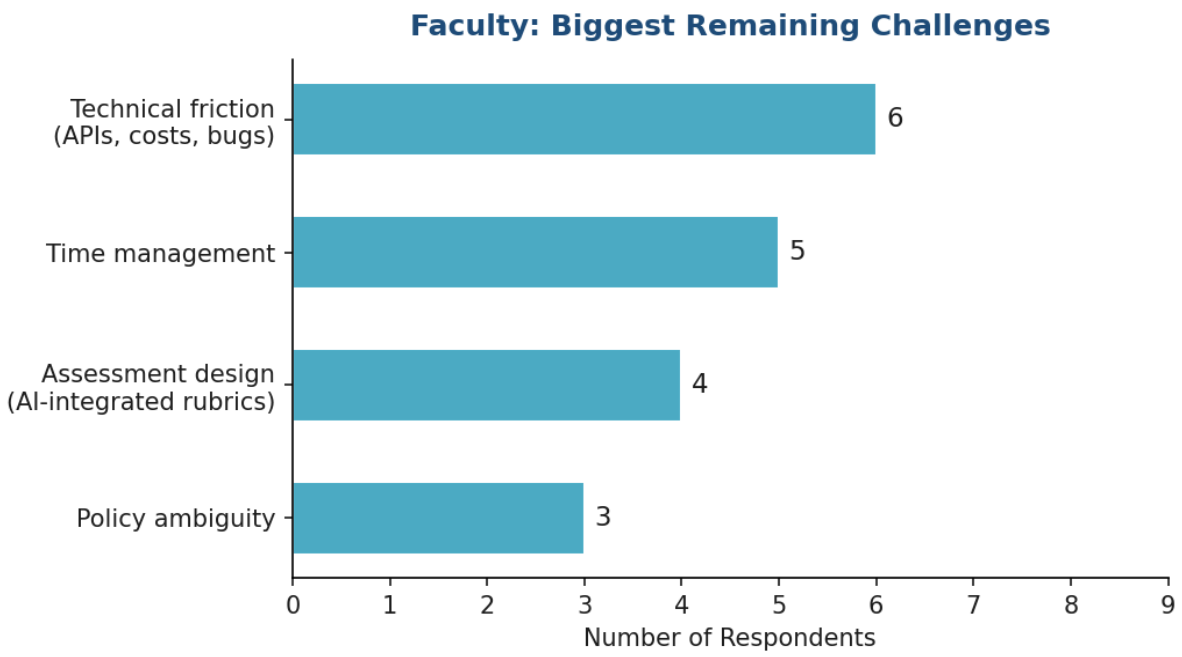

Figure 8. Faculty self-reported biggest remaining challenges (end-of-semester review, n=12, multi-select)

> *"I thought it would be easier to set up generative AI tools with significant pedagogical value in the humanities and social sciences. The cost—in time and money—is much greater than I was aware of."*

5. Continuity and Resource Needs

Twelve of 15 surveyed faculty fellows plan to continue their projects beyond the Spring 2026 semester. Five projects have planned or anticipated Fall 2026 classroom deployments: the French oral tutor, The Tutor, Critical Lens, the LEPC simulation, and the AI historical figures and patient case bots.

Several fellows are actively seeking funding to sustain and scale their work. Eric Matheis has developed a formal pilot proposal with the program's faculty lead's support, and Barbara Spinelli is exploring FRAP funding and plans to apply for a SOLER grant for 2027–2028. These efforts illustrate a broader pattern: the fellowship does not just produce prototypes—it produces faculty who are equipped and motivated to advocate for their own projects, write proposals, seek institutional partnerships, and connect their pedagogical innovations to the broader scholarly conversation. This capacity building may be among the program's most durable contributions.

The most consistently expressed need was continued intern support—but nearly all faculty who want their interns to continue reported having no confirmed funds. Other resource needs include API and software credits for continued prototyping, research assistant support for data analysis, financial support for student computation costs and pilot study incentives, and institutional guidance on AI tool policies.

6. Recommendations

The following recommendations are drawn directly from faculty and intern program reviews, weekly reflections, bootcamp feedback, and patterns observed across the semester. They are offered not as criticism of what was accomplished, but as the natural next iteration of a program that has already proven its value.

6.1 Program Structure

Prioritize a one-to-one intern-to-faculty ratio. This was the single most consistent recommendation from both faculty and interns. Faculty with dedicated interns consistently reported stronger outcomes and smoother collaboration.

Introduce a phased project structure. A three-phase model—Conceptualization (weeks 1–4), Development (weeks 5–10), and Completion/Deployment (weeks 11–14)—with rough milestone dates would help teams manage momentum and identify when pivots are needed.

Add a formal mid-semester check-in. A structured review around week 7 would create an opportunity to formally reset scope before the final stretch.

6.2 Onboarding and Training

Provide a foundational AI literacy primer. A brief resource explaining tokens, APIs, models, costs, and common platforms would reduce early friction for less technical faculty and ensure all participants start from a shared baseline.

Offer structured technical onboarding for interns. A practical session covering commonly used tools—prompt engineering, API access, key frameworks—would allow interns to focus on project work sooner.

6.3 Institutional Support

Advocate for course-load credit or formal recognition for participating faculty. Multiple faculty noted that participation required significant time investment not recognized in their formal workload. Teaching relief or equivalent credit would signal institutional commitment and make the program sustainable.

Clarify and communicate university AI policy guidance. Faculty need clear guidance on what AI tools they may use and how—including whether copyrighted content can be used to configure AI tools. This guidance should be issued before the next cohort begins.

Establish a clear API and computation funding pipeline. A streamlined process for requesting API keys, cloud credits, and reimbursements—ideally with a small discretionary fund available from day one—would reduce technical friction significantly.

6.4 Ethics and Privacy

Provide proactive IRB guidance for pilot studies. As projects move into student-facing pilots, a brief orientation on when IRB review applies would ensure compliance before evaluations launch.

Develop a shared data governance template. A template covering AI tools, hosting, data retention, access controls, consent mechanisms, and applicable regulations would improve consistency and institutional accountability across all projects.

7. Looking Forward

The inaugural semester of the aiX Faculty Fellows Program established a powerful proof-of-concept for fostering meaningful pedagogical advancement. By uniting 20 faculty fellows and 15 student interns across disparate academic landscapes, the initiative translated institutional urgency into tangible outcomes. While faculty emerged with refined skills and a robust peer network, interns experienced significant professional and technical growth, surpassing their initial expectations for the collaboration.

The most meaningful insight gained is the efficacy of a discipline-centered approach. When faculty engage with artificial intelligence on their own terms—anchored in deep subject-matter expertise and supported by dedicated technical partners—they cultivate solutions that are far more innovative and durable than those produced by prescriptive

training models. This model treats the integration of technology not as a technical hurdle, but as a sophisticated form of *disciplinary renewal*.

Artificial intelligence represents a fundamental shift in the production and communication of knowledge, posing a critical challenge to traditional educational frameworks. The aiX Fellowship addresses this transformation by positioning faculty as the intellectual owners of change. By anchoring the university's AI strategy in the cultivation of professional judgment and creativity, the program ensures that pedagogical inquiry remains at the heart of institutional evolution.

The program generates impressive momentum in one semester but its work has just started. With most initiatives currently in beta testing, there remains to be a sustained need for individualized mentoring and technical resources. The program's potential to scale effectively depends on establishing a stable institutional infrastructure, including policy clarity and a dedicated funding pipeline for API and computation costs.

The faculty-led, discipline-grounded model developed in the aiX Faculty Fellowship program here has already emerged as a replicable approach for peer institutions. By committing to the expansion of this work, Columbia does more than support its own community; it strengthens a broader movement to ensure that the future of AI in higher education is governed by those who understand the nuances of teaching and learning most deeply: the faculty.