From Incident to Insight: Understanding AI Model Lifecycle Management Through Case Analysis

Yeqing Kong

Georgia Institute of Technology

Abstract Understanding the AI lifecycle—the series of tasks and decisions that shape an AI project from conception to deployment—is essential for ensuring responsible AI practices. This article presents a hands-on teaching case designed for technical and professional communication classrooms, where students examine lifecycle failures in actual AI incidents. This approach helps students gain a deeper understanding of AI model lifecycle management in real-world contexts.

Keywords AI lifecycle, generative AI, AI incidents, responsible AI, technical communication pedagogy, case analysis

Introduction and Rationale

As AI systems become integral to sectors ranging from healthcare to transportation, pressing questions arise: How can we ensure these powerful tools serve society safely, fairly, and ethically? A key to addressing this challenge lies in understanding the *AI lifecycle*—the sequence of tasks and decisions that guide an AI project from conception to deployment (De Silva & Damminda Alahakoon, 2022; Data Science PM, 2024; Gcore, 2024). AI model lifecycle management "manages the complicated AI pipeline and helps ensure the necessary results in enterprise" (Ishizaki, 2020). More than a series of technical steps, the AI lifecycle provides a framework for making informed choices that shape the behavior and societal impact of AI systems. De Silva and Alahakoon (2022) developed a comprehensive AI lifecycle model comprising three overarching phases (design, develop, and deploy) and 19 constituent stages that trace the journey of AI models from inception to production. These stages span from problem identification to reviewing data ethics, data preparation, risk assessment, and performance monitoring and evaluation.

Implementing effective AI lifecycle management strategies provides numerous advantages for organizations. These include maximizing model reliability and performance, optimizing resource use, improving scalability and explainability, and, most critically, promoting responsible practices that help prevent harm (Deepgram, 2024; Dupont, 2024; Gcore, 2024). Nevertheless, managing the AI model lifecycle presents significant challenges. These include growing model complexity, the need for rigorous data quality and governance, and the difficulty of interpreting sophisticated systems. Moreover, addressing model bias, often rooted in flawed or historical data, requires proactive strategies such as diverse data collection and bias auditing to foster ethical AI development (Dupont, 2024).

AI lifecycle management has drawn increasing interest across sectors. For instance, Troy Tazbaz and John Nicol (2024), specialists at the U.S. Food and Drug Administration's Digital Health Center of Excellence, published a blog post outlining how lifecycle management can support the delivery of safe, effective AI-enabled healthcare. They proposed leveraging Life Cycle Management (LCM), a business management methodology foundational in software development since the 1960s. They mapped traditional lifecycle phases onto specific AI development steps, resulting in a seven-stage model structured as a cycle: planning and design; data collection and management; model building and tuning; verification and validation; model deployment; operation and monitoring; and, finally, real-world performance evaluation. This final stage loops back to planning and design, reflecting the iterative nature of AI development. Their model also details the technical and procedural considerations associated with each phase, offering a structured approach to managing risks while ensuring effectiveness and ethical compliance.

Furthermore, the AI lifecycle approach aligns with the principles of *Responsible Artificial Intelligence (Responsible AI)*, which emphasizes developing, assessing, and deploying AI systems in safe, ethical, and trustworthy ways (What is Responsible AI, 2024). For example, Microsoft's Responsible AI Standard (2022) articulates six guiding principles—fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability—which provide a comprehensive framework to guide AI development. Integrating AI lifecycle management into technical communication courses offers students the opportunity to explore how Responsible AI principles apply at each phase of AI system development.

In addition to industry-led explorations, scholarship in Technical and Professional Communication (TPC) has developed pedagogical frameworks for teaching AI literacy and ethics. Many scholars emphasize the need for multi-level integration. Nupoor Ranade and Marly Saravia (2024) call for engagement at the institutional, curricular, and classroom levels, advocating both the adaptation of existing courses and the creation of new instructional modules. Similarly, Matthew Vetter et al. (2024) proposed a "local ethic" framework, stressing that ethical engagement with AI should be grounded in the specific values and contexts of individual classrooms. In parallel, Danielle DeVasto and Zsuzsanna Palmer (2024) demonstrated through classroom experiments that how AI is integrated into a curriculum can deeply influence students' ethical reasoning and conceptual understanding of AI's societal implications. Finally, Peter Cardon et al. (2023) offered an AI literacy framework for professional contexts, organized around four dimensions: application, authenticity, accountability, and agency. Taken together, these contributions reflect an emerging consensus in TPC: that pedagogy must move beyond technical skill-building to help students critically navigate the ethical, rhetorical, and institutional dynamics of AI systems.

Building on this growing body of pedagogical work, this article presents an instructional activity designed to introduce students to AI lifecycle management through a hands-on, analytical investigation. Students examine real-world AI projects that experienced significant failures, identify the lifecycle stage(s) where issues emerged, explore underlying causes, and propose corrective actions. Through this process, they not only deepen their technical understanding of lifecycle management but also enhance their critical AI literacy and ethical awareness.

The activity unfolds in four structured steps:

1. Step 1 introduces AI lifecycle concepts through targeted readings to prompt class discussion.

2. Steps 2 and 3 guide students through a case analysis, where they investigate failures in actual AI incidents.

3. Step 4 concludes with a debrief, in which students present their findings, reflect on lifecycle mismanagement, and consider how Responsible AI principles might prevent similar failures.

The remainder of this article is organized as follows. First, I describe the context of the activity, outline student learning outcomes, and offer step-by-step implementation guidance. Next, I provide an assessment plan for instructors. The article concludes with a discussion of the activity's strengths, limitations, and suggestions for future iterations.

Activity Design

Context

I implemented this activity in my upper-level undergraduate course, *Responsible AI in Communication*, in Fall 2024. Offered through the School of Literature, Media, and Communication at the Georgia Institute of Technology, the course fulfills a humanities requirement for students from a wide range of majors. The activity took place during a 75-minute session in the first unit of the course, which introduces foundational concepts in artificial intelligence. Prior sessions addressed questions such as What is (generative) AI? What is an algorithm? What is machine learning?, What is a large language model?, and What is Responsible AI? These topics provide students with the technical and ethical groundwork necessary to understand the AI lifecycle.

Although developed for this specific course, the activity is highly adaptable for other communication-focused classes, such as courses covering computers and writing, communication and emerging technologies, communicate science and technology to the public, contemporary issues in professional communication, human-centered design, AI ethics, and responsible AI. Its interdisciplinary relevance makes it a valuable addition to curricula that address the social, ethical, and technical dimensions of emerging technologies.

Student Learning Outcomes (SLOs)

By the end of this activity, students will be able to:

1. Define and describe key stages of the AI lifecycle.
2. Analyze real-world AI system failures to identify which AI lifecycle stage(s) contributed to the failure.
3. Evaluate the underlying technical, organizational, or ethical factors behind lifecycle mismanagement.
4. Apply Responsible AI principles to recommend improvements to flawed AI lifecycle practices.
5. Communicate findings through structured presentations and engage in reflective discussion on AI ethics and lifecycle governance.

Four Steps of the Activity

Step 1: Preparation Session (10 minutes)

The session begins with a warm-up designed to reinforce students' understanding of the AI lifecycle and its critical stages. To prepare, students are assigned two short readings prior to class. During the first 10 minutes, the instructor facilitates small-group discussions centered on the following prompts:

1. Thomas, Rob. (2019). ***The AI ladder: Demystifying AI challenges***
Authored by Rob Thomas, Senior Vice President of Software and Chief Commercial Officer at IBM, this booklet introduces AI in enterprise contexts and outlines four key stages of the AI development pipeline:

- **Collect:** Simplify and make data accessible
- **Organize:** Establish a foundational analytics framework
- **Analyze:** Build and scale AI systems with transparency and trust
- **Infuse:** Operationalize AI across the organization

The AI Ladder framework provides organizations with a structured approach to transforming raw data into actionable insights. It emphasizes streamlining data collection, organization, and analysis to create a foundation for a governed, efficient, and adaptable AI strategy.

Since this resource predates recent advancements in generative AI, instructors may wish to supplement it with a more current reading focused on AI model lifecycle management. Such materials can be sourced from professional organizations (e.g., the Association for the Advancement of Artificial Intelligence), academic research databases (e.g., ACM Digital Library), or industry-focused hubs like Hugging Face, which offer documentation and tools on model versioning, deployment, and lifecycle practices.

2. Dupont, Maxim. (2024, May). *Mastering AI model lifecycle management*
This trade publication, accessible to college-level readers, addresses fundamental questions about AI lifecycle management, including its definition, importance, components, challenges, best practices, and tools. To build on the broader phases outlined in Thomas' AI Ladder, Dupont (2024) breaks the lifecycle into five specific stages (refer to Figure 1):

- **Data Preparation and Preprocessing:** Collect and clean data to ensure it is accurate and suitable for model development
- **Model Development and Training:** Choose and refine algorithms using performance metrics to build a reliable predictive system
- **Model Evaluation and Validation:** Assess the model's accuracy, precision, and recall using validation datasets to confirm performance benchmarks
- **Model Deployment and Monitoring:** Deploy the model into production systems and actively monitor its performance to detect emerging issues
- **Model Governance and Maintenance:** Manage version controls, retraining, and updates to maintain security, compliance, and relevance

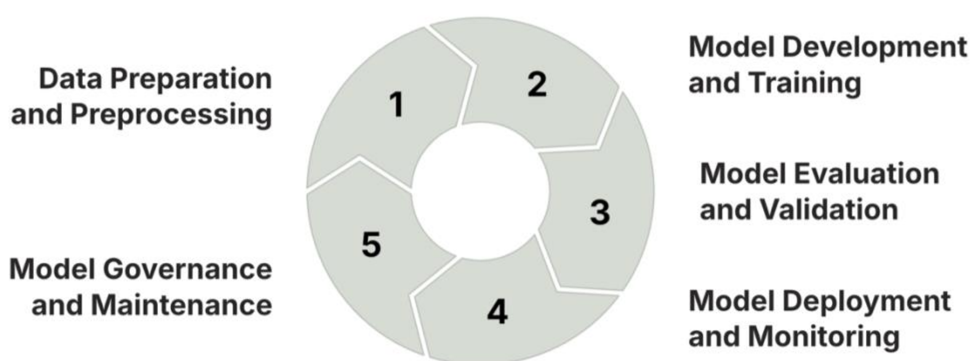

Figure 1. AI model lifecycle management process

Importantly, the AI lifecycle is not linear. Rather, it is an iterative process in which stages are revisited to incrementally refine and improve the model (see Figure 1; Data Science PM, 2024). Instructors can help students connect these AI lifecycle stages to familiar procedural frameworks in their disciplines, such as the writing process or the design thinking model, to scaffold the integration of new concepts with existing knowledge.

Step 2: AI Incidents Investigation (15 minutes)

In small groups of three to four, students select an AI incident from the [AI Incident Database](https://incidentdatabase.ai/) (https://incidentdatabase.ai/), a publicly accessible, crowdsourced platform launched in November 2020 to track documented harms caused by AI systems (McGregor, 2020). In 2022, the database became part of the Responsible AI Collaborative, an initiative dedicated to "identifying, defining, and cataloging artificial intelligence incidents." The platform is intended to make AI system failures more visible and analyzable for developers, researchers, educators, and policymakers, with the goal of preventing or mitigating future harms (Responsible AI Collaborative Founding Report, 2022).

The database includes over 5,000 AI incidents, defined as "an alleged harm or near harm event to people, property, or the environment where an AI system is implicated" (Editor's Guide, 2025). It is important to distinguish between an AI incident and an AI issue—the latter refers to potential harm that has not yet occurred or been detected. Each entry in the database includes a unique ID, date, summary, taxonomy classifications, timeline, and links to supporting reports (typically drawn from news and social media coverage). This structure is comparable to incident reporting systems in fields like aviation and cybersecurity. Students navigate to the "Discover Incidents" page (refer to Figure 2), where they can search for incidents using keywords or filter by criteria such as source, incident date, language, authorship, and tags. Students may also explore clusters of related incidents using the database's *Spatial Visualization* page, which visually maps patterns across AI failures.

According to a recent newsletter published by the AI Incident Database, more than 80 new incident IDs were added between April and May 2025 (Atherton, 2025). Many of these incidents reflect an intensification of previously identified trends, including deepfake scams, AI-enabled fraud, the misuse of generative tools to produce nonconsensual or harmful content, and the global spread of mis- and disinformation. These patterns underscore the evolving and increasingly complex landscape of AI-related harms.

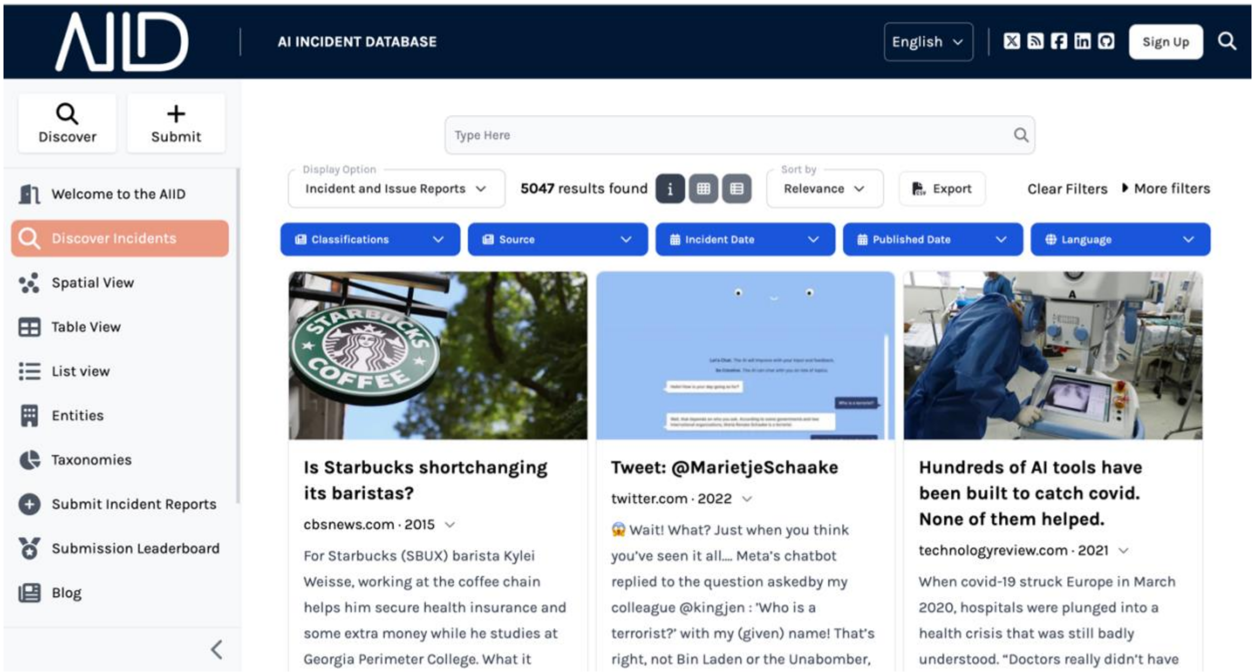

Figure 2. Screenshot of "Discover Incidents" page in the AI incident database

After selecting a case, students share initial impressions within their groups. Guiding questions include:

- What is your first impression of the incident?
- What do you assume went wrong?

Students then individually examine the incident reports to understand the context, stakeholders, and technical aspects of the failure.

Step 3: AI Lifecycle Failure Analysis (15 minutes)

Once familiar with their selected incident, students spend the next 15 minutes analyzing the case collaboratively. The goal is to identify where in the AI lifecycle the failure occurred and how it could have been prevented. Groups respond to the following guiding questions:

- At which stage(s) of the AI lifecycle did the failure occur?
- What were the root causes of the failure?
- What ethical considerations, if any, were overlooked?
- How could lifecycle management practices have been improved to prevent this failure?

Each group documents their responses in a shared Google Slides deck, with one slide designated per group. Each slide includes pre-filled prompts to ensure consistent analysis across teams.

Step 4: Group Presentations (35 minutes)

The activity concludes with group presentations, in which each team shares their case analysis using a designated slide. Each group is allotted approximately seven minutes, making this portion of the session about 35 minutes total for five groups. Presentations follow a structured format:

- A brief introduction of the selected AI incident
- Identification and analysis of the lifecycle stage(s) where failure occurred
- Discussion of ethical considerations that were overlooked or insufficiently addressed
- Summary of key takeaways and recommendations, with a reflection on how their understanding evolved through the analysis

Assessment of Student Learning Outcomes

This section presents a rubric for evaluating students' presentations on real-world AI lifecycle failure cases. The assessment criteria are aligned with the SLOs identified earlier in the activity. The total score for the presentation is 100 points, distributed across five key areas:

1. Understanding of AI Lifecycle Stages (SLOs 1 & 2)
2. Root Cause Analysis (SLO 3)
3. Application of Responsible AI Principles (SLO 4)
4. Presentation & Communication (SLO 5)
5. Critical Reflection & Engagement (SLO 5)

Table 1. AI Lifecycle Case Analysis Presentation Rubric

Criteria	Excellent (Full Marks)	Good (Minor Gaps)	Developing (Some Gaps)	Needs Improvement	Points
Understanding of AI Lifecycle Stages (SLOs 1 & 2)	Accurately identifies relevant lifecycle stages and clearly explains their role in the failure case. Demonstrates deep conceptual understanding.	Identifies relevant stages with minor gaps or unclear explanations.	Partial identification; explanations are vague or inaccurate.	Lifecycle stages are misidentified or missing.	20
Root Cause Analysis (SLO 3)	Thorough, well-reasoned diagnosis of technical, organizational, or ethical causes. Uses evidence from the case and theory.	Logical causes identified; some evidence or theory used.	Causes are too general, lack specificity or depth. Limited evidence.	Superficial or incorrect analysis with little to no justification.	20
Application of Responsible AI Principles (SLO 4)	Clear, actionable recommendations grounded in ethical frameworks (e.g., fairness, transparency). Demonstrates high-level ethical reasoning.	Recommendations align with ethical principles but are underdeveloped or general.	Limited or vague ethical application; lacks strong reasoning.	Recommendations are weak, unrealistic, or unrelated to Responsible AI.	20
Presentation & Communication (SLO 5)	Organized, clear, and professional delivery. Slides or visuals enhance understanding. Team or individual speaks with confidence.	Mostly clear and organized; minor delivery issues. Visuals are adequate.	Some clarity issues; inconsistent organization or delivery. Visuals may be lacking.	Presentation is disorganized, unclear, or unengaging. Poor visuals.	20
Critical Reflection & Engagement (SLO 5)	Offers insightful reflection on lifecycle management and Responsible AI. Engages thoughtfully with peer discussion/questions.	Reflection is relevant but lacks depth. Engages with discussion.	Reflection is brief or surface-level. Limited peer interaction.	Minimal or no reflection. Does not engage in discussion.	20

Conclusion: Autoethnographic Reflection on the Activity Design

As this article does not present student data, Institutional Review Board approval was not required. Instead, I offer an autoethnographic reflection on my instructional experience with this activity. I observed that while many students initially approached AI failures as purely technical problems, they ultimately came to recognize that these failures often stem from deeper systemic issues, such as inadequate data governance, lack of transparency, weak accountability mechanisms, or insufficient ethical oversight. Through the case analysis, students became skilled at identifying which stages of the AI lifecycle were implicated in the failures, most commonly data collection, model training, or inadequate validation during deployment. In doing so, they surfaced a wide range of ethical concerns, including bias, privacy violations, and threats to public safety. These realizations helped students better appreciate the broader societal responsibilities of AI developers.

This activity effectively introduces students to the complexities of AI lifecycle management by integrating real-world case analysis, fostering critical thinking, and enhancing their ethical awareness. A key strength of the activity lies in its ability to help students connect technical processes with ethical implications—an essential competency in the context of responsible AI development (Dupont, 2024). Moreover, the use of incident analysis as a pedagogical strategy reinforces a systems-level perspective. Students learn to see the AI lifecycle not as a linear sequence but as a dynamic, interdependent process, where failures at any stage can trigger cascading effects (De Silva & Alahakoon, 2022).

However, several limitations and challenges emerged in implementing this activity:

1. **Technical complexity**: Some AI incidents involved technical details or proprietary algorithms that exceeded students' existing knowledge. This challenge echoes Dupont's (2024) observation that model complexity and interpretability remain significant barriers to effective AI lifecycle management.

2. **Time constraints**: The limited class time poses a challenge for conducting thorough investigations. Given the breadth of AI lifecycle stages and the potential complexities of certain incidents, students may struggle to perform deep analyses within a single session.

3. **Limited access to tools and documentation**: Students relied primarily on publicly available incident reports and lacked access to internal documentation or technical platforms commonly used in enterprise AI lifecycle management, such as IBM Cloud Pak for Data (Ishizaki, 2020) and Microsoft's Responsible AI Toolbox (https://responsibleaitoolbox.ai/). These platforms, while powerful, often require institutional subscriptions, posing accessibility challenges in classroom settings.

To address these limitations, future iterations of the activity could be expanded in both scope and duration.

1. Instructors can select AI incidents with clear, accessible documentation, enabling students to focus on lifecycle and ethical analysis without being overwhelmed by technical complexity.

2. Dividing the activity across multiple sessions or developing it into a project-based module would provide students with more time to engage deeply with each stage of the AI lifecycle. For advanced courses, incorporating multiple

case studies for comparative analysis could further support pattern recognition and critical evaluation across diverse AI failures.

3. Instructors could integrate open-source tools such as TensorFlow, PyTorch, or Kubeflow to simulate aspects of AI lifecycle management, offering students hands-on experience with model debugging and risk assessment. For advanced courses, using multiple cases for comparative analysis could encourage students to identify patterns and variations across incidents.

In sum, these adaptations would enhance students' technical engagement and more effectively prepare them for ethical decision-making and responsible practice in AI-related fields.

References

Atherton, Daniel. (2025, June). AI Incident Roundup – April and May 2025. https://incidentdatabase.ai/blog/incident-report-2025-april-may/

Cardon, Peter; Fleischmann, Carolin; Aritz, Jolanta; Logemann, Minna; & Heidewald, Jeanette. (2023). The challenges and opportunities of AI-assisted writing: Developing AI literacy for the AI age. *Business and Professional Communication Quarterly*, *86*(3), 257-295.

Data Science PM (2024, November). *What is the AI life cycle?* https://www.datascience-pm.com/ai-lifecycle/

Deepgram. (2024, June). *AI lifecycle management*. https://deepgram.com/ai-glossary/ai-lifecycle-management

De Silva, Daswin De; & Alahakoon, Damminda. (2022). An artificial intelligence life cycle: From conception to production. *Patterns*, *3*(6), 100489, 1-13.

DeVasto, Danielle; & Palmer, Zsuzsanna. (2024). Building critical AI literacy in the business communication classroom. *Business and Professional Communication Quarterly*, *87*(4), 557-572.

Dupont, Maxim. (2024, May). *Mastering AI model lifecycle management*. https://www.labelvisor.com/mastering-ai-model-lifecycle-management/

Editor's Guide. (2025). https://incidentdatabase.ai/editors-guide/

Gcore. (2024, April). The complete guide to AI model cycle. https://gcore.com/learning/ai-model-lifecycle

Ishizaki, Kazuaki. (2020, November). *AI model lifecycle management: Overview*. https://www.ibm.com/think/topics/ai-lifecycle

Mcgregor, Sean. (2020, November). When AI systems fail: Introducing the AI incident database. https://partnershiponai.org/aiincidentdatabase/

Ranade, Nupoor; & Saravia, Marly. (2024). Teaching AI ethics in technical and professional communication: A systematic review. *IEEE Transactions on Professional Communication*, *67*(4), 422-436.

Responsible AI Standard (v2). (2022, June). https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4ZPmV

Responsible AI Collaborative Founding Report. (2022, March). https://asset.cloudinary.com/pai/cf01cce1af65f5fbb3d71fa092d001db

Tazbaz, Troy; & Nicol, John. (2024). Blog: A lifecycle management approach toward delivering safe, effective AI-enabled health care. https://www.fda.gov/medical-devices/digital-health-center-excellence/blog-lifecycle-management-approach-toward-delivering-safe-effective-ai-enabled-health-care?utm_medium=email&utm_source=govdelivery

Thomas, Rob. (2019). *The AI ladder: Demystifying AI challenges*. O'Reilly Media.

Vetter, Matthew; Lucia, Brent; Jiang, Jialei; & Othman, Mahmoud. (2024). Towards a framework for local interrogation of AI ethics: A case study on text generators, academic integrity, and composing with ChatGPT. *Computers and composition*, *71*, 102831.

What is Responsible AI? (2024, September). https://learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai?view=azureml-api-2

Author Information

Yeqing Kong is an Assistant Professor of Technical Communication in the School of Literature, Media, and Communication at Georgia Institute of Technology. Her research examines the communication of scientific, technological, health, and environmental risks in the global context. Her scholarly work has appeared in *Technical Communication Quarterly*, *IEEE Transactions on Professional Communication*, *Journal of Business and Technical Communication*, and Proceedings of the ACM International Conference on the Design of Communication.