

Key Considerations for Domain Expert Involvement in LLM Design and Evaluation: An Ethnographic Study

Annalisa Szymanski
Computer Science and Engineering
University of Notre Dame
Notre Dame, IN, USA
aszzyman2@nd.edu

Toby Jia-Jun Li
toby.j.li@nd.edu
Computer Science and Engineering
University of Notre Dame
Notre Dame, IN, USA

Oghenemaro Anuyah
Microsoft Corporation
Redmond, WA, USA
maroanuyah@microsoft.com

Ronald A. Metoyer
Computer Science and Engineering
University of Notre Dame
Notre Dame, IN, USA
rmetoyer@nd.edu

Abstract

Large Language Models (LLMs) are increasingly developed for use in complex professional domains, yet little is known about how teams design and evaluate these systems in practice. This paper examines the challenges and trade-offs in LLM development through a 12-week ethnographic study of a team building a pedagogical chatbot. The researcher observed design and evaluation activities and conducted interviews with both developers and domain experts. Analysis revealed four key practices: creating workarounds for data collection, turning to augmentation when expert input was limited, co-developing evaluation criteria with experts, and adopting hybrid expert-developer-LLM evaluation strategies. These practices show how teams made strategic decisions under constraints and demonstrate the central role of domain expertise in shaping the system. Challenges included expert motivation and trust, difficulties structuring participatory design, and questions around ownership and integration of expert knowledge. We propose design opportunities for future LLM development workflows that emphasize AI literacy, transparent consent, and frameworks recognizing evolving expert roles.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI; Interactive systems and tools.**

Keywords

Large Language Models, Ethnography, Evaluation Methods, LLM-as-a-Judge, Human-AI Interaction

ACM Reference Format:

Annalisa Szymanski, Oghenemaro Anuyah, Toby Jia-Jun Li, and Ronald A. Metoyer. 2026. Key Considerations for Domain Expert Involvement in LLM Design and Evaluation: An Ethnographic Study. In *31st International Conference on Intelligent User Interfaces (IUI '26)*, March 23–26, 2026, Paphos, Cyprus.



This work is licensed under a Creative Commons Attribution 4.0 International License. *IUI '26, Paphos, Cyprus*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1984-4/2026/03
<https://doi.org/10.1145/3742413.3789105>

Cyprus. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3742413.3789105>

1 Introduction

Large Language Models (LLMs) are increasingly used in high-stakes domains such as education, healthcare, and law, where the accuracy of outputs carries important ethical and practical implications [1, 5, 15, 63, 64]. These models can produce fluent, contextually relevant responses that expand access to expertise and assist in complex decision-making [8, 76]. However, these same models may also reproduce bias, generate misleading outputs, or hallucinate information [33, 37]. As LLMs become more accessible through open-source and commercial platforms, questions of validation, reliability, and appropriate deployment have become increasingly urgent [32, 47, 82]. The reliability of LLM-based systems depends not only on technical performance but also on how development teams curate data, define fine-tuning objectives, and establish evaluation criteria [6, 60], requiring teams to balance scalability, rigor, and contextual relevance across iterative development cycles [43].

Several strategies for evaluating LLMs have emerged, including performance benchmarks [40, 54], LLM-as-a-Judge methods [25, 39, 83], and direct review by domain experts [53, 72]. Each approach offers trade-offs: benchmarks and automated metrics may fail to capture domain-specific requirements; LLM-as-a-Judge methods, though scalable, can introduce bias or overlook subtle errors [16, 74]; and domain expert evaluation can provide rich, contextualized feedback [48, 49, 77], but their involvement is resource-intensive, time-consuming, and often limited [28]. While these techniques are often studied in isolation, little is known about how they are combined in real-world development workflows, how teams involve experts, or how organizational factors influence these decisions.

Unlike traditional software, LLM-based systems are non-deterministic and require human judgment and iterative refinement [7, 11, 71]. Development involves coordination across technical and domain expertise, often under constraints of time, interpretability, and institutional expectations. Yet, current research offers limited insight into how these trade-offs are negotiated *in situ* and how reliability is prioritized in development work. To address this gap, we argue for empirical approaches that examine how evaluation and design decisions are made in practice. Ethnographic

methods allow researchers to observe not only formal workflows but also informal practices, negotiations, and constraints that shape design and evaluation [27, 65]. By embedding in the team’s day-to-day work, we aim to understand how decisions about expert involvement and evaluation are made, and what trade-offs emerge. This work addresses the following research questions:

- **RQ1:** How do members of an LLM development team carry out design and evaluation activities over the course of a development cycle?
- **RQ2:** What challenges and trade-offs emerge in decision-making about evaluation and expert involvement?
- **RQ3:** What lessons can inform better workflows, tools, and practices for supporting domain experts in LLM development?

We conducted a 12-week participatory ethnographic study of a university-based team developing an LLM-powered pedagogical chatbot. Through a thematic analysis of observations and interviews with both developers and pedagogy experts, we found that the team approached design and evaluation through four interrelated practices: (1) creating workarounds to collect gold-standard conversations from experts, (2) turning to data augmentation when expert data was limited, (3) utilizing collaboration with experts to co-develop evaluation criteria, and (4) adopting hybrid evaluation strategies that combined expert input, developer judgment, and LLM-as-a-Judge scoring. These practices illustrate how the team balanced rigor with time and resource constraints. However, expert involvement also introduced challenges. Experts reported fatigue with repetitive tasks, struggled with limited AI literacy when asked to design evaluation criteria, and expressed concerns about embedding their knowledge in training data. Despite these challenges, expert involvement was motivated by institutional trust, alignment with the project’s pedagogical mission, and the opportunity to shape a tool for their community.

We contribute insights into the design and evaluation practices that evolve during LLM development, the constraints that shape expert involvement, and the requirements to sustain meaningful expert engagement over time. We summarize the following research contributions:

- An in-depth understanding of the design and evaluation decisions made by an LLM development team during the creation of a domain-specific chatbot.
- An analysis of the challenges and trade-offs involved in integrating domain expertise, drawn from both team observations and interviews.
- Key considerations for promoting expert–developer collaboration, including ethical implications and best practices for involving domain experts in LLM development.

2 Background and Related Work

2.1 Existing Approaches to Studying LLM Development

To gain insight into LLM development, previous literature has used surveys, interviews, and large-scale data analyses. Studies of public developer communities have characterized challenges in prompt design, API use, and plugin integration [17], as well as fine-tuning,

dataset management, and deployment [2]. Analyses of open-source software have identified model, parameter, and integration challenges in generative AI applications [12, 68]. Survey-based studies have mapped methods, metrics, and benchmarks for LLM tasks such as model compression [84] and evaluation [13, 14, 30, 58], while qualitative work has examined issues of bias, privacy, and usability [36, 56, 62]. Although these studies provide valuable overviews, most rely on retrospective or aggregated data and give limited visibility into how development decisions are made in context [50, 55].

As a result, we know little about the day-to-day work through which teams interpret model behavior, negotiate evaluation strategies, and adapt to constraints. Ethnographic methods, by contrast, capture these practices as they unfold to reveal what practitioners do and why they make particular decisions [57, 69]. This approach provides a critical complement to existing survey- and interview-based studies by grounding observations of development work.

2.2 Integrating Domain Experts in LLM Design

Prior research in HCI has emphasized the value of incorporating multiple stakeholders in AI development, yet less attention has been given to how development teams decide when and how to involve them within the constraints of organizational context. In applied machine learning (ML), domain expertise often enters the pipeline through the creation of gold-standard datasets, where experts label or provide examples that serve as benchmarks for training and evaluation [31]. This approach is common in fields such as healthcare and law, where expert judgment defines correctness [77]. However, relying on static datasets limits adaptability when models are deployed in dynamic professional settings and requires significant time and resources from experts [9, 28].

Beyond dataset construction, participatory and co-design methods have emerged to embed expertise more directly into the design process. Human-in-the-loop approaches enable experts to guide requirement gathering, prompt design, and model evaluation [61, 66, 73]. Participatory workshops have also been used to co-create evaluation rubrics, elicit tacit knowledge, and ground model outputs in professional standards [53, 72]. However, while participatory methods capture context-specific expertise that benchmarks often miss, they are difficult to scale and require substantial coordination and time commitment from participants [28].

Despite these advances, little is known about how expertise is integrated into day-to-day development. Most studies focus on discrete activities such as annotation or evaluation, without examining how teams balance expert input, automated methods, and shifting project goals.

2.3 Strategies for Evaluating LLMs

Evaluating LLMs remains an active area of research, especially for complex tasks where accuracy and contextual understanding are difficult to quantify [14]. Traditional automated metrics such as BLEU and ROUGE [40, 54] offer scalability but often miss the domain-specific qualities that matter most in applied settings. As a result, human evaluation has become the preferred standard [53, 72].

Human evaluators provide grounded judgments that capture reasoning, context, and appropriateness [79], but their involvement

can be costly and time-consuming [28]. Expert input is especially valuable in specialized domains to align model outputs with professional standards [18, 41]. However, these evaluations are difficult to scale and sustain and motivates interest in automated alternatives. To address these limitations, researchers have explored the LLM-as-a-Judge approach, where a separate model acts as an evaluator to approximate human judgment [25, 39, 83]. Under certain conditions, these automated judges show promising correlation with human preferences [24] and offer a scalable complement to human review. Typically, LLM evaluators apply pre-defined natural-language evaluation criteria, or statements reflecting correctness, completeness, clarity, or tone, to rank model responses [67, 81].

In real-world contexts, these LLM-mediated evaluations offer several practical benefits. They help developers identify strengths and weaknesses in model performance, detect emerging errors, and decide when to refine prompts or retrain models [3, 67]. They can also serve as early indicators of when human or expert input may be necessary to support more efficient allocation of limited human review. Continuous, criterion-based evaluation [23, 67] may help developers and end-users build confidence in LLM outputs to offer guidance on when systems can be trusted or require further oversight [42]. However, LLM evaluators inherit the same biases and limitations as the models they assess [16, 70, 83], which further complicates decisions about when their judgments can be trusted and when expert oversight is needed.

Understanding how development teams navigate these questions, how they decide when to rely on automated evaluation, when to seek expert input, and how to reconcile the two, is critical for building responsible, reliable LLM systems. Through an ethnographic lens, we trace how these methods were integrated, adapted, and negotiated throughout the creation of a pedagogical chatbot to reveal how evaluation workflows evolve under real-world constraints.

3 Method

3.1 Ethnographic Approach and Rationale

To address our research questions (RQ1-RQ3), a 12-week ethnographic study was conducted. We adopted a participatory ethnographic approach to obtain a situated understanding of day-to-day practices, interactions, and meanings that shaped design and evaluation decisions within the development team. Ethnography enables the production of thick description [27, 65], revealing how values, constraints, and professional expertise are negotiated through practice. The first author, referred to as the researcher, served as a member of the development team, engaging in meetings and informal discussions while maintaining a reflective and analytical stance. This approach allowed for direct access to how design and evaluation unfolded in real-time and revealed both explicit decisions and the reasoning behind them.

The ethnographic method was chosen for two main reasons. First, LLM development differs from traditional software engineering in its iterative and interpretive nature: model behavior is unpredictable, evaluation criteria evolve dynamically, and expert input is often negotiated under resource and organizational constraints. Ethnography captures how these indeterminate processes unfold *in situ* and reveals the tacit judgments that shape the technology. Second, by participating as a recognized member of the team while

maintaining an analytic stance, the researcher could access both formal meetings and informal exchanges central to sense-making. This positioning allowed the study to trace how developers and pedagogical experts defined “good” performance, justified evaluation decisions, and established boundaries between human and LLM responsibility.

3.2 Study Context and Participants

This study was conducted between June and August 2025 with a development team at a private university developing an LLM-powered pedagogical chatbot. The project took place within the university’s teaching and learning center, an academic unit providing pedagogical resources, consultations, and programming for faculty and instructors. The center’s mission is to promote effective teaching practices and enhance learning by offering instructors individualized consulting sessions, workshops, and ongoing professional development opportunities. The consultants, or domain experts, are experienced educators who support instructors by providing expertise in pedagogy, instructional design, and faculty development. The chatbot project, referred to as *InstructAI*, originated from a collaboration between the center and a development team within the institution. The goal was to create an AI-powered system that provides real-time, pedagogically grounded teaching suggestions to faculty. Specifically, *InstructAI* was envisioned as a tool that could respond to instructors’ questions about classroom challenges, teaching strategies, or student engagement by drawing upon a knowledge base of best practices and pedagogical guidance.

When this study began, the project was entering a critical phase of transitioning from initial requirements gathering to iterative model fine-tuning, evaluation, and system testing. Prior to the researcher joining, the development team had interviewed several pedagogical consultants and facilitated an early design workshop to storyboard potential chatbot features and was ready to integrate domain experts into the evaluation loop and begin systematic testing of the model’s responses. The development team operated semi-independently from the center but maintained close collaboration through recurring joint sessions. It consisted of three members (UX project manager, technical lead, and software engineer) that operated within a U.S.-based academic institution, all with prior experience in AI and LLM research or development. The UX project manager had prior training in AI in education and human-computer interaction, the technical lead brought prior industry research experience and expertise in developing and optimizing large language models, and the software engineer contributed a background in full-stack web development and machine learning, including experience integrating generative AI components into interactive systems. The development team worked closely with a primary stakeholder, a pedagogical scholar for the center, responsible for setting the project vision and coordinating requirements. The domain experts who were affiliated with the center, participated in design and evaluation activities at key points, such as co-developing evaluation rubrics, reviewing chatbot outputs, and advising on pedagogical framing. Figure 1 depicts this multi-stakeholder structure, with both technical contributors and pedagogical specialists.

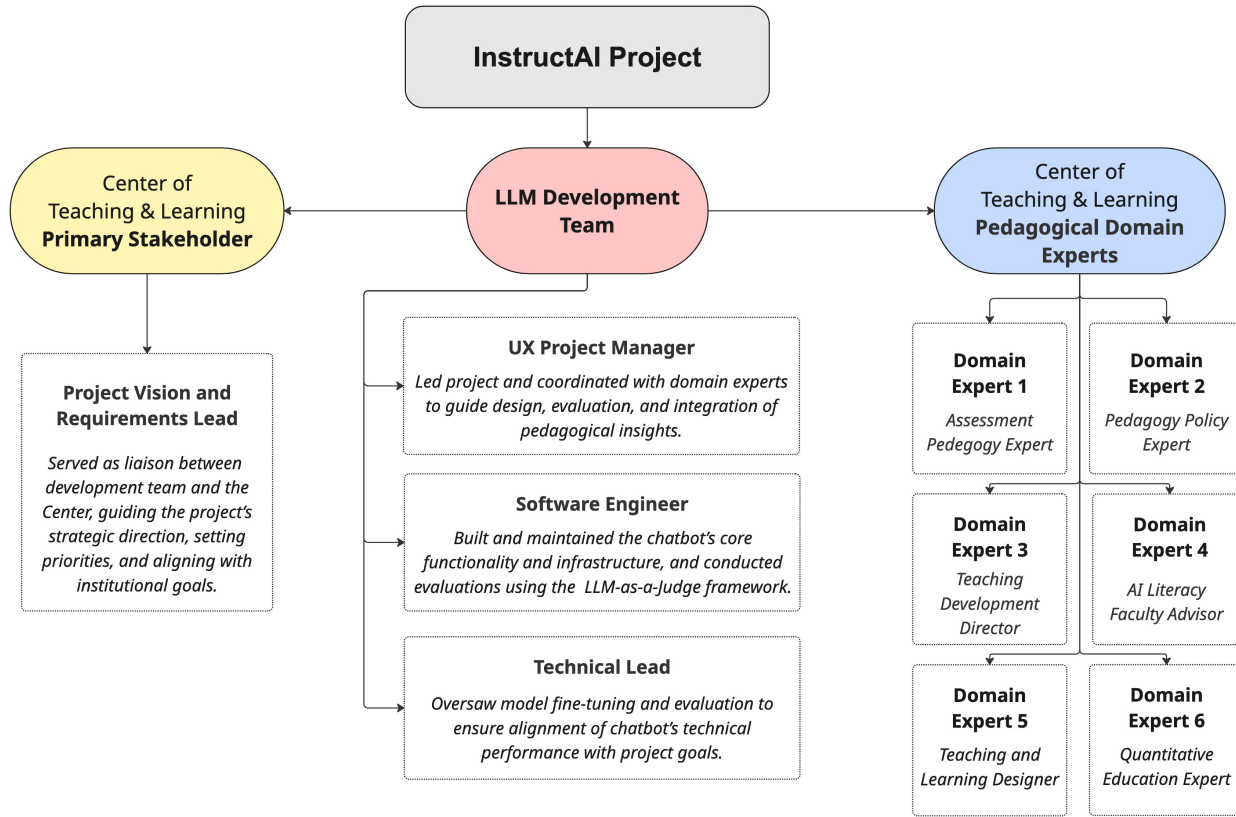


Figure 1: Stakeholder structure and roles in the InstructAI project. The LLM development team coordinated design, implemented the system, and oversaw model development and evaluation. The team worked in collaboration with a primary stakeholder from the center of teaching and learning, who set the project vision and requirements, and with pedagogical domain experts, who contributed expertise during design and evaluation activities.

3.3 Researcher Role and Participatory Ethnographic Approach

This study employed a participatory ethnographic approach, positioning the first author as an embedded researcher within the LLM development team. Unlike a detached observer, participatory ethnography involves active engagement in the social setting while maintaining reflexive awareness of the researcher’s dual role as a participant and analyst [45]. As a recognized and trusted member of the team, the researcher participated in daily stand-ups, weekly planning meetings, and evaluation sessions, contributing observations and perspectives when relevant but intentionally refraining from taking ownership of project deliverables or influencing final design decisions. This positioning provided access not only to formal design and evaluation practices but also to the reasoning, trade-offs, and interpersonal dynamics that shaped them.

The researcher had access to all team communications, including Slack channels, shared documents, and email threads, which enabled observation of ongoing interactions where design and evaluation decisions were negotiated. The team’s awareness of the

researcher’s dual role encouraged reflective discussions, which allowed observation of collective sensemaking and self-critique. This embedded perspective also captured developer-expert interactions that shaped how expert knowledge was incorporated into LLM evaluation. Observing and occasionally facilitating these engagements revealed how expert contributions were framed, interpreted, and operationalized across iterative development cycles. This participatory lens captured not only what the team did, but how technical and pedagogical members understood, justified, and adapted their practices in response to emerging challenges.

3.4 Data Collected

To capture a holistic view of the development and evaluation process, multiple qualitative data sources were drawn upon. These included daily and weekly observations of team practices, interviews with both developers and domain experts, and a collection of artifacts that documented design decisions, communications, and evaluation outcomes.

3.4.1 Daily Observations. Over a 12-week period, the researcher conducted 70.8 hours of observations embedded within the team’s day-to-day workflow. The team held daily 30-minute stand-ups (focusing on blockers, evaluation tasks, and model updates) and two recurring weekly meetings (30 minutes and 75 minutes) with the primary stakeholder to align evaluation plans, rubric revisions, and milestone targets. The researcher also observed ad-hoc activities central to the project’s design and evaluation loop: one-on-one strategy check-ins, small-group discussions on emergent issues, planning sessions for upcoming design/evaluation workshops, and both virtual and in-person team checkpoints. Integration into Slack and email channels, allowed observation of informal negotiations where decisions formed, issues emerged, trade-offs were accepted, and priorities shifted with stakeholder input or evaluation results. Observations also extended to sessions with domain experts, including design workshops, rubric co-development, review of AI-generated outputs, and debriefs after both evaluation and data contribution activities (e.g., collaborative sessions to develop evaluation criteria or provide exemplar data for model training).

The researcher balanced participation with analytic distance by contributing context when invited without owning deliverables or decisions. During meetings, the researcher produced real-time jottings and same-day expanded field notes. Notes were organized under consistent headings: Actors, Activities, Artifacts (documents, prompts, rubrics), Decision Points, Rationales, Tensions/Trade-offs, Follow-ups. Furthermore, the notes were accompanied by weekly analytical memos that captured emerging interpretations. These memos were revisited to check for interpretive drift and to refine questions for subsequent observations or interviews.

Analytically, observation time concentrated on key episodes where design and evaluation intertwined:

- Gold-standard data creation (how exemplar responses were elicited from domain experts, prepared for fine-tuning, and positioned as reference material for evaluation).
- Evaluation-method design (introducing, revising, or retiring criteria; deciding thresholds; debating inter-rater alignment between developers and pedagogical experts).
- Human–automation balance (when to invoke LLM-as-a-Judge vs. expert review; how to resolve disagreements; how automated scores influenced prioritization).
- Stakeholder alignment (how pedagogical priorities, timeline constraints, and technical feasibility were reconciled).

This structured and flexible observation strategy yielded a thick description of everyday development and evaluation practices to capture both formal processes and the informal workarounds through which the team progressed.

3.4.2 Interviews. We combined informal, *in-situ* conversations with semi-structured interviews to investigate situated practices and reflective accounts of the design and evaluation. Informal interviews (contextual inquiries) were conducted continuously after meetings or key episodes (e.g., rubric changes, conflicting scores), while semi-structured sessions were scheduled once the team had completed several evaluation cycles, allowing participants to reflect across iterations. We interviewed all members of the development team ($n = 3$) to capture complementary perspectives across technical leadership, engineering, and research roles. We

also interviewed the the primary stakeholder and domain experts from the center who had directly participated in design and/or evaluation activities (e.g., rubric co-development, output review, data collection tasks). Experts were selected via criterion sampling to ensure variation in pedagogical specialization and degree of involvement across the study period.

Each interview session lasted approximately 45 to 60 minutes and was conducted toward the end of the 12-week study, once participants had substantial experience with design and evaluation activities. We used insights from our field notes and early observations to develop two tailored semi-structured interview protocols: one for developers and one for domain experts. The interviews were tailored to deepen understanding of how design and evaluation unfolded in practice and to complement the ongoing participant observation data.

Development Team interviews. Interviews with the three development team members focused on their roles, responsibilities, and reflections on integrating evaluation into the system’s lifecycle. Questions probed how evaluation was built into the development process, whether it was planned from the beginning or added iteratively, and how collaboration with pedagogical experts shaped decision-making. For instance, developers were asked: “*Can you walk me through how evaluation was built into the development process? Were evaluation and feedback considered from the beginning, or integrated later?*” “*In your view, where is human expert input most necessary, and where might an LLM be sufficient?*” and “*What challenges did you encounter when trying to incorporate domain expertise into design or evaluation?*”

These questions were grounded in themes that emerged during daily observations, such as the negotiation between automated evaluation (LLM-as-a-Judge) and human expert review, as well as the trade-offs between model performance goals and pedagogical fidelity. Developers were also asked to reflect on tensions between the team’s technical priorities and the center’s educational mission, and to suggest improvements to future workflows or protocols.

Domain Expert Interviews. Interviews with pedagogical consultants (domain experts) examined their experiences participating in the design and evaluation of the chatbot. Experts were asked to describe how they were recruited, what kinds of evaluation tasks they performed, and how clearly those tasks were framed. For example, questions included: “*Can you describe how you were involved in the evaluation or data collection process for InstructAI? What kind of tasks were you asked to complete—reviewing outputs, designing rubrics, or participating in the teacher simulator?*” “*How clear were the instructions or expectations for your role in this process? Did you feel your contributions were taken into account?*” “*Based on your experience, in what areas do you feel your evaluation as an expert added the most value, and where might an LLM do just as well or even better?*” Additional questions invited reflection on comfort and motivation: “*How did you feel about the time and effort the process required?*” and “*What would you want developers of AI systems to understand about the importance of domain expertise?*”

By tailoring questions to participants’ actual engagement in the project, such as their involvement in designing rubrics, reviewing AI outputs, or calibrating evaluation criteria, we encouraged concrete, example-driven accounts rather than abstract opinions. This helped reveal how participants made sense of their contributions, what

they found meaningful or challenging, and how they viewed the balance between expert judgment and automated evaluation.

All interviews were audio-recorded and transcribed verbatim, and identifiers were removed to ensure confidentiality. Informal follow-up conversations were also documented in field notes to clarify statements or revisit emerging interpretations. Together, these interviews provided insight into both how design and evaluation practices were enacted and how participants understood and rationalized their roles in shaping the LLM's development.

3.4.3 Artifacts. In addition to observations and interviews, a range of artifacts was collected that documented the team's iterative design and evaluation process. These included draft evaluation rubrics, participatory design session materials, chatbot outputs, annotated model responses, expert-generated data, and LLM-as-a-Judge results and performance summaries. These materials showed how evaluation criteria evolved, how expert input was incorporated, and how model performance was interpreted over time, which complemented what was observed in meetings and interviews.

3.5 Ethical Considerations

This study was approved by the university's Institutional Review Board (IRB), and all participants provided informed consent. The researcher's role as an embedded team member was disclosed to all collaborators, and care was taken to maintain voluntary participation and confidentiality. Identifiers were removed during transcription, and pseudonyms are used throughout. Because the researcher was involved in team activities, reflective memos and regular debriefs were used to mitigate bias and ensure analytic transparency. No student or instructional data were used, and all materials were stored on secure, access-restricted servers in accordance with institutional guidelines.

3.6 Data Analysis

We conducted a thematic analysis [10] across observation field notes and interview transcripts, complemented by an analytic review of collected artifacts. The first author led the analysis and interpretation of the data.

For RQ1, analysis focused on longitudinal observation field notes, interviews with developers, and artifacts documenting day-to-day development work. Through repeated readings, the researcher embedded initial codes capturing observable practices (e.g., refining evaluation criteria, coordinating with experts). These codes were iteratively grouped into subthemes (e.g., ethical barriers to real data, expert input improves benchmarks) and then synthesized into higher-level themes describing interrelated design and evaluation practices (e.g., designing workarounds to collect gold-standard conversations, turning to augmentation strategies due to limited data). Artifacts such as evaluation criteria rubrics, annotated model outputs, and expert evaluation documents were examined to trace the evolution of criteria and to support insights emerging from observations and interviews.

For RQ2, analysis drew primarily on interviews with developers and pedagogical experts. Interview transcripts were coded for

expressed concerns and constraints (e.g., confusion about task focus, loss of ownership), which were iteratively grouped into subthemes (e.g. difficulty understanding evaluation abstractions, concerns about giving away expertise) and higher-level themes (e.g. challenges in structuring participatory design processes, concerns about how expert feedback would be used) capturing key challenges and trade-offs in expert involvement.

Across both analyses, themes were refined through repeated engagement with the data, reflexive memo writing, and regular discussions with co-authors. Although formal inter-coder reliability was not employed, these discussions supported critical reflection on analytic decisions, challenged interpretations, and guided the iterative refinement of emerging themes.

3.7 Positionality Statement

All authors are computer scientists with strong backgrounds in human-computer interaction and human-centered computing. The first author brought prior experience in human-AI interaction research, including extensive work on LLM evaluation and methods for incorporating domain experts into LLM evaluation processes. This background supported informed interpretation of evaluation and design practices during analysis and guided clarifying questions in discussions with developers and domain experts when needed. The remaining authors contributed complementary expertise in community-centered technology research, qualitative and ethnographic methods, and technical development and evaluation of LLM systems. Collectively, these perspectives shaped the study's attention to situated practices, trade-offs, and organizational constraints influencing LLM evaluation. Analytic rigor was supported through reflective memo writing, iterative review of data, and regular collaborative discussions among the research team.

4 Findings

4.1 Development Team Design and Evaluation Practices (RQ1)

Analysis revealed that the development team carried out design and evaluation activities through four interrelated practices: (1) creating workarounds to collect gold standard conversations from pedagogical experts, (2) turning to augmentation strategies due to limited data collected, (3) utilizing collaboration with pedagogical experts to co-develop evaluation criteria, and (4) adopting hybrid evaluation approaches to address time constraints and inconsistencies. These practices illustrate the adaptive strategies the team employed in the development cycle.

4.1.1 Developers Designed Workarounds to Collect Gold Standard Conversations from Pedagogical Experts. The development team experimented with multiple strategies to collect what they referred to as "gold standard" conversational data from pedagogical experts. Their goal was to obtain authentic examples of pedagogical consultation that would serve two purposes: (1) as training data for fine-tuning the model, and (2) as reference data for evaluating model outputs. The dataset was intended to capture authentic conversations between instructors and pedagogical consultants that the chatbot was designed to model and to illustrate how a pedagogical expert would respond.

The team's initial strategy was to secure recordings of real consultations between pedagogical experts and faculty clients that occurred at the university's teaching and learning center, where consultations were frequent and took place daily. However, a challenge arose regarding the ethical concerns of recording these consultations as part of the data collection and design process. Developers learned that the process of recording could alter the nature of the conversations, discourage faculty participation, and compromise the trust central to these consultations. One developer reflected on the difficulty of securing access:

"Recording would significantly change the [center's] current services...and probably wouldn't work...and I wasn't going to push...they weren't seeing the direct benefit [of recording]. It would require extra work like getting consent...and it might even negatively impact [the consultation] service." (UX Project Manager)

As a workaround for collecting consultation data, the team recorded a small set of staged conversations in which members of the development team, together with participants who had classroom teaching experience, met with a pedagogical expert for consultation. Developers were limited to recruiting four participants due to time and scheduling constraints. Although these recordings generated some useful material, the developers quickly realized that the dataset was too small and too labor-intensive to clean, since the transcripts contained disfluencies, short filler responses, and language patterns that required extensive editing. They also noted that the conversational style differed from how a user may interact with an AI chatbot, making it less suitable for training or evaluation purposes.

To take a different approach, the team made a pivotal decision to move away from relying on recordings or small staged conversations and instead collected data in text form. For this purpose, developers created a simulation that generated consultation prompts and captured the responses of pedagogical experts in a chatbot-style format. They quickly prototyped the system, training it to adopt an instructor persona, pose an initial question about instruction, and provide follow-up responses. The development team then recruited 18 domain experts and invited them to interact with the chatbot, contributing asynchronously at their own pace and with hourly compensation. Through this strategy, the team collected over 130 simulated dialogues to create a more scalable form of "gold standard" data.

In comparing the two approaches, one developer highlighted why the simulator data felt more practical than transcripts from real conversations:

"[Real conversations] need extra effort to clean by removing those short answers or other fillers that aren't useful. But chatting [through text] is another good way to collect data, because people may [be inclined to] type more when chatting with [the simulator]." (Software Engineer)

To evaluate the simulator, the development team applied the LLM-as-a-Judge framework to assess how well the simulated instructor generated questions that reflected those an instructor might naturally ask. A team member developed tailored evaluation criteria

to guide these automated assessments, which focused on dimensions such as pedagogical relevance, cognitive depth, contextualization, and breadth of instructor concerns. These automated evaluations suggested that the instructor dialogues ranged from good to excellent in quality, giving the developers confidence that the simulator could be well-suited to ask questions.

In summary, each shift in data collection reflected a decision shaped by practical constraints. Recordings were abandoned because privacy concerns outweighed potential benefits, staged conversations proved too limited in size and format, and the simulator was adopted as a feasible way to involve experts more widely and build a usable dataset. The developers had to acknowledge potential issues and amend their approach to work around the issue of collecting expert feedback.

4.1.2 Developers Turned to Augmentation Strategies Due to Limited Data Collected. Upon reviewing the gold-standard dialogues collected through expert interactions with the simulator, the team determined that while the data were useful for evaluation, the sample was too limited in size and diversity to support fine-tuning. This recognition motivated a shift in strategy. Rather than continuing to seek additional expert dialogues through the simulator, they decided to adopt data augmentation methods to artificially expand the dataset and approximate the kinds of expert-like dialogues needed for training. One participant who was part of the development team summarized this reasoning:

"I had the instinct that I wouldn't have enough data to train a model...[after group discussion]...I should think of new ways [to expand data], like augmentation." (UX Project Manager)

To implement this approach, the team generated synthetic multi-turn dialogues that modeled the expert data that had been collected from the simulator. The augmented conversations were crafted to mirror the structure and tone of authentic consultations while expanding the dataset to 1,000 synthesized dialogues. This step reflected the team's focus on efficiency, as augmentation allowed them to quickly scale beyond what could be collected from experts in real time. In addition, the team did not conduct a direct evaluation of the augmented data itself due to time constraints and assessed the model's augmented outputs only after fine-tuning (see Section 4.1.4).

In summary, augmentation enabled the developers to move forward despite limited data from experts to produce a larger and more consistent dataset for training. This shift, however, also marked a turning point in the design process: rather than grounding the model primarily in expert-provided conversations, the team chose to use synthetic data to approximate expertise at scale.

4.1.3 Developers Collaborated with Pedagogical Experts to Co-Develop Evaluation Criteria. After generating data for fine-tuning the model, the team needed a way to evaluate whether the chatbot's responses reflected sound pedagogical practice. While domain experts offered the highest standard of evaluation, the team recognized that relying exclusively on them would be prohibitively time-consuming and costly. Instead, they began exploring strategies that could balance the domain experts' oversight with scalable evaluation methods.

Within this process, the idea of using an LLM-as-a-Judge again emerged as a central strategy to capture the essence of expert judgment without requiring experts to rate every output. Their solution was to translate expert feedback into evaluation criteria that the LLM-as-a-Judge could apply at scale. This reasoning reflected a compromise: experts were still essential for setting standards, but once those standards were formalized, the automated judge could take over routine evaluation. In this way, the team preserved the authority of expert input while creating a workflow that was more efficient, sustainable, and repeatable. As one developer explained, the team’s limited pedagogical expertise meant they could only achieve a baseline score with the LLM-as-a-Judge alone, but expert guidance could significantly raise the benchmark:

“We are not experts in this domain...we have a baseline [set of criteria] that will be rated [by the LLM-as-a-Judge] six out of 10, but we can get to probably nine out of 10 [with expert input into defining criteria]” (UX Project Manager)

This criteria was used for two main purposes: first, to quickly evaluate the model and identify outputs that did not align with expert expectations, and second, to serve as metrics for comparing the fine-tuned model to a standard GPT model and testing performance differences.

Initially, to create criteria, the development team drew on existing standard practices for pedagogical consultation without experts to draft an initial set of criteria grouped into categories. These criteria were created from utilizing online resources on “best practices” for assessing both the quality of pedagogical advice and the overall coherence of the dialogue. However, the developers realized that the criteria that they developed were generic and lacked the domain specificity that they could not supply themselves. As one developer explained:

“We didn’t want to use standard evaluation criteria, such as [measures of] fluency or typical checks of how correct the language is... those are criteria that can be used for every LLM task. It is important that we trust human involvement more.” (UX Project Manager)

With this in mind, the team concluded that the most effective way to elicit high-quality input from domain experts was to organize a series of participatory design workshops, rather than simply asking individuals to generate evaluation criteria independently. They reasoned that a collaborative process would make it easier for experts to refine their ideas, align on shared standards, and better connect their knowledge to the system’s needs. As one developer explained, when thinking about the time and cost constraints:

“I feel like the only way that we can get experts [to actively participate] is to do a workshop so that they can focus and get the thing done.” (UX Project Manager)

The team intentionally recruited pedagogical experts with diverse backgrounds, such as consultants who regularly advised faculty, researchers who studied pedagogy, and professors with practical classroom experience to surface different perspectives on what counts as “good” teaching advice. The development team conducted three small-group sessions, each lasting 1.5 hours, with two hourly compensated experts to facilitate in-depth discussions. In planning

these sessions, the developers also debated how best to capture feedback, such as through whiteboard brainstorming, paper annotations, or digital criteria editing. Ultimately, they distributed the draft criteria in a shared document, allowing experts to directly revise and suggest edits on their laptops.

The workshops followed a structured progression, with each step intentionally designed to strike a balance between efficiency and expert understanding. First, experts were briefed on the project’s goals, the purpose of the evaluation criteria, and the concept of using LLM-as-a-Judge to scale assessment. The team reasoned that without this orientation, experts might view the task as a traditional human assessment exercise, rather than understanding how their criteria would be operationalized by an automated judge. Second, experts reviewed the draft criteria created by the developers individually and in discussion, and suggested amendments or additions. Third, experts were shown chatbot outputs alongside the LLM-as-a-Judge’s automated ratings and asked whether the criteria were capturing meaningful qualities. This comparison was included to make the abstract concept of automated evaluation more concrete, giving experts a chance to see how the judge “interpreted” the initial criteria in practice. Finally, experts interacted directly with the chatbot to test whether their refinements better reflected pedagogical standards in real-world use. This final step created an opportunity for iterative reflection and ensured that the criteria were not only theoretically valid but also practical in guiding model evaluation.

During the initial rubric review, the experts’ criteria evolved to be more descriptive, particularly by clarifying the guiding statements or questions used by the LLM-as-a-Judge. As experts compared chatbot outputs with LLM-as-a-Judge scores, they identified disagreements between their own assessments and the judge’s rankings. These mismatches prompted attempts to clarify wording, refine numerical scales, and add contextual guidance. However, because changes to the criteria could not be tested interactively to observe immediate updates in LLM-as-a-Judge scoring, refinements remained limited to minor edits rather than substantive restructuring. After experts interacted directly with the chatbot, no major structural changes to the criteria were introduced, though several usability issues were identified and documented by the development team.

In summary, the developers recognized their own limitations in defining robust evaluation criteria and sought expert involvement to ensure that the chatbot was assessed against pedagogical standards of quality and accuracy. The team also invested considerable effort in designing strategies for introducing criteria to present them in ways that would allow experts to readily engage with, critique, and refine them.

4.1.4 Developers Used a Hybrid Evaluation Approach to Address Time Constraints and Inconsistencies. After fine-tuning the model with new augmented training data and establishing evaluation criteria, the development team adopted a hybrid evaluation approach that combined domain expert input, developer judgment, and automated scoring from the LLM Judge.

Initially, the team asked pedagogical experts to evaluate chatbot responses using the criteria developed in earlier sessions, with the

goal of comparing expert ratings to those produced by the LLM-as-a-Judge. To do this, experts participated in informal interviews in which they reviewed chatbot responses and ranked the outputs using the criteria. However, this task quickly proved too tedious, as experts had to read each full conversation, decide which criteria applied, and provide both a rating and justification.

The team then shifted responsibility to members of the development group, who were more readily available to conduct evaluations. Developers reasoned that, after working closely with experts in participatory design sessions, they had developed a reasonable sense of how experts would apply the criteria. This gave them confidence to take over the evaluations in order to save time and flag any obviously erroneous responses.

In addition, the team used the LLM-as-a-Judge to rank the provided criteria against both their chatbot and a different GPT model in order to compare performance. To do this, they tested different LLM-as-a-Judge prompts until the rankings aligned more closely with expert and developer manual assessments. Yet development team members noted that automated scores were not always consistent:

“If we use the LLM-as-a-Judge, sometimes the scores are not consistent. Without [the LLM] giving a justification, it’s kind of a black box. For example, if I set the temperature to zero, the scores may be consistent, but if I add something else to the prompt, the scores can change.” (Technical Lead)

Despite these limitations, developers would discuss which criteria were more impactful, meeting regularly to review results and decide whether metrics truly differentiated model performance. As one developer described, the process was less about finding perfect scores than about identifying which criteria meaningfully discriminated between the models that they were comparing:

“We look at how the scores turn out. If most of the scores are perfect, maybe that criterion is not important. Then we try another criterion to see if one model scores differently from another, which can help show meaningful distinctions between models.” (Technical Lead)

These findings suggest that while the development team recognized expert input as ideal, they ultimately relied on developer involvement and automated LLM-as-a-Judge methods to conduct evaluations. Importantly, the criteria not only shaped how evaluations were conducted but also informed model refinement. Developers used patterns in the scores, such as criteria that consistently exposed weaknesses, to guide adjustments in fine-tuning and prompt engineering and effectively looped evaluation results back into the development process.

4.2 Challenges and Trade-offs (RQ2)

Five overarching themes were identified concerning the challenges and trade-offs the team faced as they navigated the design and evaluation of the LLM chatbot.

4.2.1 Motivations and Trust Featuring Expert Participation. Interviews with pedagogical experts revealed that their willingness to participate was shaped by several motivating factors, including professional curiosity about AI, institutional trust, available time,

and compensation. For some, the experience was described as intellectually and professionally rewarding. Others emphasized a desire to be involved so that they could better understand the tool and eventually promote it to faculty instructors:

“It gives a lot of credibility if I’m going to be working with a faculty member, and say, hey, there’s this tool it’s pretty helpful, I know the designers, and here’s what goes into it, as opposed to some company somewhere made something.” (Domain Expert 3)

Another expert described initial hesitation toward participating, stemming from uncertainty about AI and its implications for writing instruction. This reluctance evolved into motivation as participation became a chance to learn, reflect, and contribute meaningfully to the project:

“I was a little reluctant at first because I don’t use AI...but I always walked away from sessions feeling like I’d learned something new...It was definitely more work than I expected, but the time felt well spent.” (Primary Stakeholder)

Another domain expert framed skepticism as a productive motivation for engagement, emphasizing that diverse viewpoints strengthened the design and evaluation process:

“I appreciate being included because I am an AI skeptic. I support the use of AI, and I think we should use and embrace it, but I feel like I’m often the most skeptical person in the room in a lot of meetings...it’s good to have people with a range of opinions about AI.” (Domain Expert 2)

Trust in the institution emerged as a consistent theme, as experts felt comfortable contributing because the project was led by colleagues within their university rather than driven by profit motives. This sense of institutional alignment made participation feel safer and more meaningful. Experts stated that they would have felt less comfortable contributing their expertise to commercial or for-profit efforts:

“In an organization that has thoughtful leaders, that has strong consensus around institutional values and goals and adequate resources to support all of that...I’m not worried.” (Domain Expert 3)

In summary, the pedagogical experts emphasized that their motivation to participate was not only practical, as driven by their curiosity, trust, or compensation, but also ethical. Experts carried with them a sense of responsibility for how their knowledge would be used and expressed a desire for their contributions to benefit the university community rather than external commercial interests.

4.2.2 Challenges in Structuring Participatory Design Processes. Observations of the participatory design sessions revealed several challenges that shaped how domain experts engaged with creating evaluation criteria for the LLM-as-a-Judge. While the workshops provided valuable input, experts frequently struggled to distinguish between evaluating the chatbot itself and refining the criteria. These challenges stemmed from three main factors: limited prior experience with design sessions, difficulties in understanding the LLM-as-a-Judge concept, and insufficient time for reflection.

Limited Experience with Design Sessions. For many experts, this was their first time participating in a structured design session to create evaluation criteria. The process involved reviewing model outputs and comparing them to criteria, which included tasks that felt unfamiliar and cognitively demanding. As one expert reflected:

“I’ve never done anything like that before...so it was challenging for me to remember the steps that the LLM was evaluating the chatbot, and we were evaluating the criteria. It took me a minute to figure out what we were actually doing or what we were actually evaluating.” (Domain Expert 5)

Several experts described the process as “intense” and suggested that receiving the criteria beforehand might have made the task more manageable:

“I think the process was a bit intense...For me, it takes time to connect the evaluation criteria to my work experience...so I need more time to process... maybe it will be better if I can access these evaluation criteria beforehand to make more connections.” (Domain Expert 4)

Difficulties Understanding the LLM-as-a-Judge Concept. Although the development team introduced the idea of the LLM-as-a-Judge, some experts found it difficult to keep this framework in mind during the sessions. They often defaulted to thinking about how they would want the chatbot to behave, rather than focusing on whether the evaluation criteria captured meaningful standards. As one participant explained:

“It was tricky, because I was thinking to myself about what I would want out of [the chatbot], but then I also had to make that extra step of thinking that I’m looking at the evaluation criteria...just like trying to remind myself to stick with the LLM Judge piece of it. I think we got there eventually.” (Domain Expert 5)

Insufficient Time for Reflection. Experts valued the collaborative environment of the sessions but noted that more time to reflect on the criteria would have improved the process. They suggested follow-up opportunities or real-time iteration to better connect the criteria with practice. One expert described the benefit of live feedback loops:

“Instead of doing [the criteria edits] on paper...do it in real time...make changes live. [Such that] we suggest some language or a few descriptors to the criteria, then we look at the output...and see what does it rate...to see how that’s changing...[whether] it’s not interpreting the information the way an expert would.” (Domain Expert 3)

These findings show that while participatory design workshops offered an opportunity to align evaluation criteria with expert standards, they also opened discussion around the use of expertise, the understanding of AI, and time. Experts expressed the need for more guidance and opportunities for reflection to fully engage and best contribute to the process.

4.2.3 Concerns About How Expert Feedback Would Be Used. While many experts described their participation as rewarding, they also voiced concerns about how their contributions would ultimately

be used in the development of the LLM chatbot. This was mostly about sharing data as well as the ability of the chatbot to mimic the job of an expert.

Concerns About Giving Away Knowledge. Experts had some reservations about contributing their expertise to the project and how it might be repurposed. While they recognized that providing data to the simulator could help train novice consultants by modeling realistic interactions, imputing their tacit “tips and tricks” into a dataset raised concerns about ownership and loss of control once the knowledge was embedded into the system. As one expert reflected:

“I was working with the simulator...even though you’re trying to take my job physically...I’m giving it all my tips and tricks on how I go about my consultation.” (Domain Expert 3)

Concerns About Replicating Expert Work. Most domain experts were confident that the LLM could not replicate the full scope of their work. They emphasized that pedagogical consultations often involve faculty sharing personal or vulnerable information, which requires confidentiality, trust, and professional judgment to handle responsibly. These dimensions go beyond providing advice and reflect the ethical and relational aspects of expert practice. As one expert explained:

“We try to protect the confidentiality of our clients...there are times when faculty come to us in a place of vulnerability and we would like to be able to help that person without it putting their career in jeopardy.” (Domain Expert 1)

These reflections suggest an important trade-off in expert involvement. Expert knowledge can either be essential to creating meaningful evaluation criteria and training data, or it could surface anxieties about ownership, control, and the irreplaceable aspects of human reasoning and care in pedagogical practice.

4.2.4 Fatigue and Engagement in Data Contribution. The simulator enabled the team to collect dialogues at scale; however, experts reported that the process was repetitive and, over time, exhausting. Experts developed their own strategies for staying engaged, such as selectively choosing more interesting or challenging scenarios. As one expert explained:

“I would try to do 15 to 20 minutes [of the conversations]...I could only get through two or three at a time before some of the prompts or problems that it presents are kind of repeated throughout...so I try to find one that’s more interesting or different to keep myself engaged right in the process.” (Domain Expert 3)

Beyond repetition, experts also highlighted that the typing-based format of the simulator created additional strain. They suggested that alternative modes of contribution to the simulator, such as using voice-to-text, could help sustain longer engagement. As one expert noted:

“If there’s an opportunity for a voice interaction in conversation...that would certainly sustain longer attention and engagement with it, because it’s similar to a conversation with the person.” (Domain Expert 3)

These findings reveal that while the simulator made expert involvement more scalable, the format could introduce trade-offs in terms of fatigue and depth of contributions. Designing future workflows may require diversifying interaction formats to reduce repetition and better mirror the natural dynamics of consultation practices.

5 Discussion

This study provides an in-depth examination of how expert-informed evaluation unfolds within an academic LLM development team. Our findings reveal both the potential and complexity of involving domain experts in the design and evaluation of AI systems. While the development team valued expert input and experts were motivated to participate, resource constraints, mismatched participation formats, and limited AI literacy support introduced challenges. These frictions point to the need for intentional design of workflows, tools, and ethical practices that can better support meaningful and sustainable expert engagement. In this section, we synthesize our findings into three key lessons for HCI researchers and practitioners seeking to integrate domain expertise into LLM development (RQ3).

5.1 Designing Tools That Capture Rich Expert Knowledge

Recent work in ML has increasingly emphasized data-centric approaches to improving model quality [35]. Our findings extend this perspective by showing how the design of data collection tools fundamentally shapes the kind of expert knowledge that can be captured [52]. In our study, the simulator served as a flexible data collection tool and gave experts the autonomy to contribute at their own pace. However, the repetitive questions and decontextualized prompts constrained their ability to express the deeper reasoning that reflects real consultations. This raises questions about whether the tool adequately supported the diversity of consulting scenarios and instructional queries relevant to expert practice.

Prior research highlights the persistent difficulty of encoding professional expertise for computational systems [51]. In our study, experts noted that short written responses in the simulator did not fully express deeper pedagogical reasoning, which risked reducing their expertise to surface-level data. This aligns with critiques in the ML community that model development often prioritizes algorithmic performance over the quality and representativeness of training data [34, 78]. When AI systems aim to supplement or replicate expertise, the risks of shallow data become especially concerning, since insufficient training data may harm users [74].

Lesson 1: Experts should be provided with tools that are both flexible and data-centric for collecting insights. Future workflows should consider how to elicit richer forms of expert knowledge. Approaches such as interactive demonstrations, think-aloud protocols, or focus sessions that support reflection may help capture not only explicit “answers” but also the strategies and reasoning processes underlying expert practice. In doing so, data-centric design tools can better balance the needs of developers for structured datasets with the needs of experts for autonomy in representing their work.

5.2 Evolving Roles of Experts in the Era of AI

Domain experts often struggled with a limited understanding of AI concepts, which made it challenging to fully grasp the task of co-designing evaluation criteria. Although they possessed the necessary knowledge and experience to contribute to the system evaluation, the experts struggled to fully grasp how their knowledge was being applied in the LLM-as-a-Judge, even after being introduced to the concepts and the co-design task. This raised two key risks: first, that the evaluation criteria might result in reduced quality or misalignment with the system’s needs; and second, that experts could unknowingly contribute knowledge without fully approving of its usage and application within the evaluation.

This finding is aligned with other literature that supports incorporating AI literacy support into participatory design [19]. Previous studies indicate that stakeholders are commonly engaged in critiquing system interfaces even though they are not properly prepared to assess or understand the underlying model itself [21]. Research has also shown that all participants need certain design skills and conceptual understandings to be most successful in these settings [59, 85]. Our study reinforces this gap and provides further evidence that structured AI literacy support is essential for domain experts to make meaningful contributions to the development and evaluation of AI systems.

Beyond completing the task, experts should be able to critically question the goals of the system, the risks of automation, such as LLM Judge approaches, and the boundaries of their own contributions. This requires creating reflective spaces in the design process, rather than just providing task instructions. Long et al. have discussed specific core competencies required for AI literacy [44], including high-level overviews of how algorithms operate, explainable AI features that clarify outcomes, and interactive demonstrations that enable experts to test or simulate system behavior, all of which contribute to AI literacy [44]. In our study, for instance, domain experts only began to fully grasp the role of the LLM-as-a-Judge once they were shown sample scores accompanied by the judge’s explanations. However, what was missing from the team’s participatory design process was a mechanism for experts themselves to test and refine the criteria in real-time. One expert explicitly discussed the need for interactive tools that support the iterative development and refinement of criteria within the design session. Recent HCI work has explored interactive evaluation support tools (e.g., [22, 26, 38]) which can be used to support all stakeholders in these evaluation contexts.

Our findings also point to broader shifts in the nature of expert labor. As domain experts in this study were increasingly asked to train and evaluate AI systems, their consulting roles expanded from advising instructors to shaping algorithmic decision-makers. This shift signals a potential redefinition of professional identity, from a consultant to a data provider or even a co-designer.

Lesson 2: Expert Involvement Requires Both AI Literacy Support and Recognition of Shifting Roles. HCI researchers and practitioners should take seriously the implications of this evolution. Sustaining expert engagement requires not only technical support for AI literacy but also structural recognition of experts’ changing roles and contributions in the era of AI. The involvement of experts should be considered a two-way exchange, as they need sufficient training

and support to understand how their knowledge is being used, and at the same time, their labor must be valued as more than just data collection. As LLM Developers recruit domain experts for their tasks, they should be explicit about the roles needed and provide proper training.

5.3 Ethical Implications of Expert Contributions

Involving domain experts in LLM development introduces not only technical advantages but also broader ethical considerations related to how expert knowledge is incorporated and sustained over time. Pedagogical experts were asked to share their expertise by contributing “tips and tricks” for consultations through the simulator as well as refining evaluation criteria to reflect their best practices. Most of these experts had spent years developing their skills through training and professional experience. As the project progressed, expert-generated materials were incorporated into multiple stages of the development pipeline, including data augmentation and model fine-tuning, in ways that supported evolving technical and resource needs. While these practices were consistent with approved data use and project goals, the rapid integration of expert knowledge into automated workflows meant that contributions became increasingly abstracted from the individuals who produced them. This observation points to the importance of planning and governance mechanisms for data reuse that support transparency and help preserve expert agency as development priorities and system capabilities evolve.

This raises an important question: *who owns a domain expert’s data once it is embedded into an AI system?* This concern is not unique to pedagogy [46]. Broader debates are emerging across creative and professional domains, with artists discovering their work in image-generation datasets [80], authors noting that their books have been scraped to train large models [4], and musicians raising concerns about voice synthesis and imitation [75]. In each of these cases, once expertise or creative labor is transformed into data, it becomes difficult to trace, attribute, or govern. Our findings suggest that the same problem extends to professional expertise. Pedagogical knowledge, developed through years of practice, risks being reduced to invisible “training material” for LLMs. This echoes concerns raised in Ghost Work about the invisibility of labor that sustains AI systems [29].

This invisibility raises critical questions about ownership and recognition. When inventors patent a process, they receive recognition, legal protection, and sometimes royalties. By contrast, domain experts who provide strategies, evaluation heuristics, or consultation practices may receive short-term compensation (e.g., hourly pay) but often lack recognition or control over how their contributions are utilized. If their knowledge becomes central to a system that is later scaled, commercialized, or deployed, they have little agency over its usage. Moreover, when real-world data is constrained by privacy, access, or scale, teams often rely on augmented or synthetic data, as observed in our study, further complicating questions of ownership. Prior work shows that such data can introduce bias or misrepresentation if not carefully governed [20].

Lesson 3: Experts Should Be Explicitly Informed About How Their Knowledge Will Be Used. While research teams cannot reshape

global data policy, they can implement more transparent data collection practices. Recruitment and consent forms should state how expert contributions will be utilized, whether for training, evaluation, or deployment. Teams should also consider adopting licensing-style agreements that define boundaries of use (e.g., restricted to research, restricted to this organization’s use) and clarify ownership. In addition, compensation structures could be expanded to recognize the long-term value of knowledge contributions, particularly when these shape evaluation frameworks or fine-tuned models. This might include authorship, acknowledgment, or shared rights to future system outputs. When synthetic data or simulators are used in place of real expert-authored data, experts should be involved in reviewing and refining generated examples over time to mitigate risks and retain influence as systems evolve. Above all, systems and workflows should be designed to sustain expert agency through recurring expert review checkpoints and ensure that experts remain collaborators and co-designers rather than invisible data providers.

6 Limitations and Future Work

Our work has limitations that also open directions for future research. First, this study focused on a single LLM development team observed over a 12-week period. Therefore, the findings reflect practices shaped by a specific institutional and organizational context and may not generalize to other domains or industry settings.

Second, although multiple pedagogical specialists were included, all participants came from a single institution and shared similar professional backgrounds. Future work should broaden its scope to include experts from diverse disciplines, institutions, or cultural contexts, thereby enhancing understanding of how different forms of expertise influence LLM design and evaluation practices.

Third, the 12-week observation window captured only one phase of model refinement and evaluation and did not extend into deployment or long-term system use. As such, the study cannot address how expert–developer relationships or evaluation practices evolve over time. Future longitudinal research could investigate how collaboration practices, evaluation workflows, and perceptions of trust and ownership evolve as systems mature and are implemented in real-world contexts.

7 Conclusion

Through a twelve-week ethnographic study and interviews with developers and pedagogical experts, we investigated how an LLM development team approached the design and evaluation of a pedagogical chatbot, while navigating the challenges of integrating domain expertise. Our findings show how the team navigated organizational constraints and limited expert availability through strategies, including data-collection workarounds, co-developing evaluation criteria, and hybrid expert–LLM assessment methods. These findings demonstrate the practical and ethical complexities of expert involvement, including issues of trust, motivation, and ownership. For the HCI and IUI communities, this work offers design implications for workflows and tools that better support expert participation in LLM development. We argue that future intelligent systems should incorporate mechanisms for transparent

consent, enhanced AI literacy, and recognition of expert contributions that allow domain specialists to participate as co-designers rather than solely as data providers. Such approaches can support more responsible, participatory, and sustainable LLM development practices.

8 GenAI Usage Disclosure

LLMs played a central role in the development activities observed in this study. These tools were employed by study participants as part of their normal development workflow. Generative AI tools (e.g., OpenAI's GPT models) were also used to improve grammar and clarity in writing. All text was reviewed and verified by the human authors.

Acknowledgments

This work was supported in part by the Notre Dame–IBM Technology Ethics Lab. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

References

- [1] Muhammad Aurangzeb Ahmad, Ilker Yaramis, and Taposh Dutta Roy. 2023. Creating trustworthy llms: Dealing with hallucinations in healthcare ai. *arXiv preprint arXiv:2311.01463* (2023).
- [2] Khairul Alam, Kartik Mittal, Banani Roy, and Chanchal Roy. 2024. Developer Challenges on Large Language Models: A Study of Stack Overflow and OpenAI Developer Forum Posts. *arXiv preprint arXiv:2411.10873* (2024).
- [3] Ian Arawjo, Priyan Vaithilingam, Martin Wattenberg, and Elena Glassman. 2023. ChainForge: An open-source visual programming environment for prompt engineering. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–3.
- [4] Leah Asmelash. 2023. These books are being used to train AI. no one told the authors. <https://www.cnn.com/2023/10/08/style/ai-books3-authors-nora-roberts-ccc>
- [5] John W Ayers, Adam Poliak, Mark Dredze, Eric C Leas, Zechariah Zhu, Jessica B Kelley, Dennis J Faix, Aaron M Goodman, Christopher A Longhurst, Michael Hogarth, et al. 2023. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA internal medicine* 183, 6 (2023), 589–596. doi:10.1001/jamainternmed.2023.1838
- [6] Scott Barnett, Zac Brannely, Stefanus Kurniawan, and Sheng Wong. 2024. Fine-tuning or fine-failing? debunking performance myths in large language models. *arXiv preprint arXiv:2406.11201* (2024).
- [7] Earl T Barr, Mark Harman, Phil McMinn, Muzammil Shahbaz, and Shin Yoo. 2014. The oracle problem in software testing: A survey. *IEEE transactions on software engineering* 41, 5 (2014), 507–525.
- [8] Manuela Benary, Xing David Wang, Max Schmidt, Dominik Soll, Georg Hilfenhaus, Mani Nassir, Christian Sigler, Maren Knödler, Ulrich Keller, Dieter Beule, et al. 2023. Leveraging large language models for decision support in personalized oncology. *JAMA Network Open* 6, 11 (2023), e2343689–e2343689.
- [9] Jessica G Borger, Ashley P Ng, Holly Anderton, George W Ashdown, Megan Auld, Marnie E Blewitt, Daniel V Brown, Melissa J Call, Peter Collins, Saskia Freytag, et al. 2023. Artificial intelligence takes center stage: exploring the capabilities and implications of ChatGPT and other AI-assisted technologies in scientific research and education. *Immunology and cell biology* 101, 10 (2023), 923–935.
- [10] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [11] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [12] Yangxiao Cai, Peng Liang, Yifei Wang, Zengyang Li, and Mojtaba Shahin. 2025. Demystifying issues, causes and solutions in LLM open-source projects. *Journal of Systems and Software* (2025), 112452.
- [13] Yixin Cao, Shibo Hong, Xinze Li, Jiahao Ying, Yubo Ma, Haiyuan Liang, Yantao Liu, Zijun Yao, Xiaozhi Wang, Dan Huang, et al. 2025. Toward generalizable evaluation in the llm era: A survey beyond benchmarks. *arXiv preprint arXiv:2504.18838* (2025).
- [14] Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2024. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology* 15, 3 (2024), 1–45.
- [15] Angeline Chatelan, Aurélien Clerc, and Pierre-Alexandre Fonta. 2023. ChatGPT and Future Artificial Intelligence Chatbots: What may be the Influence on Credentialed Nutrition and Dietetics Practitioners? *Journal of the Academy of Nutrition and Dietetics* 123 (2023), 1525–2529. Issue 11. doi:10.1016/j.jand.2023.08.001
- [16] Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. 2024. Humans or llms as the judge? a study on judgement biases. *arXiv preprint arXiv:2402.10669* (2024).
- [17] Xiang Chen, Chaoyang Gao, Chunyang Chen, Guangbei Zhang, and Yong Liu. 2025. An Empirical Study on Challenges for LLM Application Developers. *ACM Transactions on Software Engineering and Methodology* (2025).
- [18] Szu-Wei Cheng, Chung-Wen Chang, Wan-Jung Chang, Hao-Wei Wang, Chih-Sung Liang, Taishiro Kishimoto, Jane Pei-Chen Chang, John S Kuo, and Kuan-Pin Su. 2023. The now and future of ChatGPT and GPT in psychiatry. *Psychiatry and clinical neurosciences* 77, 11 (2023), 592–596.
- [19] Aayushi Dangol, Michele Newman, Robert Wolfe, Jin Ha Lee, Julie A Kientz, Jason Yip, and Caroline Pitt. 2024. Mediating Culture: Cultivating Socio-cultural Understanding of AI in Children through Participatory Design. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 1805–1822.
- [20] Philippe De Wilde, Payal Arora, Fernando Buarque de Lima Neto, Yik Chin, Mamello Thinyane, Serge Stinckwich, Eleonore Fournier-Tombs, and Tshilidzi Marwala. 2024. Recommendations on the use of synthetic data to train ai models. (2024).
- [21] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2021. Stakeholder Participation in AI: Beyond" Add Diverse Stakeholders and Stir". *arXiv preprint arXiv:2111.01122* (2021).
- [22] Michael Desmond, Zahra Ashktorab, Werner Geyer, Elizabeth M Daly, Martin Santillan Cooper, Qian Pan, Rahul Nair, Nico Wagner, and Tejaswini Pedapati. 2025. Evalassist: Llm-as-a-judge simplified. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 29637–29639.
- [23] Michael Desmond, Zahra Ashktorab, Qian Pan, Casey Dugan, and James M Johnson. 2024. EvalLLM: LLM assisted evaluation of generative outputs. In *Companion proceedings of the 29th international conference on intelligent user interfaces*. 30–32.
- [24] Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. 2024. Length-Controlled AlpacaEval: A Simple Way to Debias Automatic Evaluators. *arXiv preprint arXiv:2404.04475* (2024).
- [25] Yann Dubois, Chen Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy S Liang, and Tatsunori B Hashimoto. 2024. AlpacaFarm: A simulation framework for methods that learn from human feedback. *Advances in Neural Information Processing Systems* 36 (2024).
- [26] Simret Araya Gebreegziabher, Charles Chiang, Zichu Wang, Zahra Ashktorab, Michelle Brachman, Werner Geyer, Toby Jia-Jun Li, and Diego Gómez-Zarà. 2025. MetricMate: An Interactive Tool for Generating Evaluation Criteria for LLM-as-a-Judge Workflow. In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work*. 1–18.
- [27] Clifford Geertz. 2017. *The interpretation of cultures*. Basic books.
- [28] Sebastian Gehrmann, Elizabeth Clark, and Thibault Sellam. 2023. Repairing the cracked foundation: A survey of obstacles in evaluation practices for generated text. *Journal of Artificial Intelligence Research* 77 (2023), 103–166.
- [29] Mary L Gray and Siddharth Suri. 2019. *Ghost work: How to stop Silicon Valley from building a new global underclass*. Harper Business.
- [30] Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, et al. 2024. A survey on llm-as-a-judge. *arXiv preprint arXiv:2411.15594* (2024).
- [31] Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964* (2020).
- [32] Xinyi Hou, Yanjie Zhao, Yue Liu, Zhou Yang, Kailong Wang, Li Li, Xiapu Luo, David Lo, John Grundy, and Haoyu Wang. 2024. Large language models for software engineering: A systematic literature review. *ACM Transactions on Software Engineering and Methodology* 33, 8 (2024), 1–79.
- [33] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2023. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems* (2023).
- [34] Oana Inel, Tim Draws, and Lora Aroyo. 2023. Collect, measure, repeat: Reliability factors for responsible AI data collection. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 11. 51–64.
- [35] Johannes Jakubik, Michael Vössing, Niklas Kühl, Jannis Walk, and Gerhard Satzger. 2024. Data-centric artificial intelligence. *Business & Information Systems Engineering* 66, 4 (2024), 507–515.
- [36] Sajed Jalil. 2025. The Transformative Influence of LLMs on Software Development & Developer Productivity. In *2025 International Conference on Artificial Intelligence, Computer, Data Sciences and Applications (ACDSA)*. IEEE, 1–10.
- [37] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in

- natural language generation. *Comput. Surveys* 55, 12 (2023), 1–38.
- [38] Tae Soo Kim, Yoonjoon Lee, Jamin Shin, Young-Ho Kim, and Juho Kim. 2024. Evallm: Interactive evaluation of large language model prompts on user-defined criteria. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [39] Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. AlpacaEval: An Automatic Evaluator of Instruction-following Models. https://github.com/tatsubal/alpaca_eval.
- [40] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, 74–81. <https://aclanthology.org/W04-1013>
- [41] June M Liu, Donghao Li, He Cao, Tianhe Ren, Zeyi Liao, and Jiamin Wu. 2023. Chatcounselor: A large language models for mental health support. *arXiv preprint arXiv:2309.15461* (2023).
- [42] Yang Liu, Yuanshun Yao, Jean-Francois Ton, Xiaoying Zhang, Ruocheng Guo Hao Cheng, Yegor Klochkov, Muhammad Faaiz Taufiq, and Hang Li. 2023. Trustworthy LLMs: A survey and guideline for evaluating large language models' alignment. *arXiv preprint arXiv:2308.05374* (2023).
- [43] Adian Liusie, Vatsal Raina, Yassir Fathullah, and Mark Gales. 2024. Efficient LLM Comparative Assessment: a Product of Experts Framework for Pairwise Comparisons. *arXiv preprint arXiv:2405.05894* (2024).
- [44] Duri Long and Brian Magerko. 2020. What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–16.
- [45] Raymond Madden. 2022. *Being Ethnographic: A Guide to the Theory and Practice of Ethnography*. SAGE.
- [46] Armaan K Malhotra, Aayush R Malhotra, and Mark Bernstein. 2024. Who Owns the Data? In *Ethical Challenges for the Future of Neurosurgery*. Springer, 57–63.
- [47] Jiya Manchanda, Laura Boettcher, Mathews Westphalen, and Jasser Jasser. 2024. The open source advantage in large language models (llms). *arXiv preprint arXiv:2412.12004* (2024).
- [48] Peter J Martin. 1999. Influences on clinical judgement in mental health nursing. *NT Research* 4, 4 (1999), 273–281.
- [49] Benjamin Paul Michael. 2019. *Clinical Judgement: an investigation of clinical decision-making*. Ph. D. Dissertation. University of Sheffield.
- [50] Nadia Nahar, Christian Kästner, Jenna Butler, Chris Parmin, Thomas Zimmermann, and Christian Bird. 2025. Beyond the comfort zone: Emerging solutions to overcome challenges in integrating llms into software products. In *2025 IEEE/ACM 47th International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 516–527.
- [51] Pauli Pakarinen and Ruthanne Huising. 2025. Relational expertise: What machines can't know. *Journal of Management Studies* 62, 5 (2025), 2053–2082.
- [52] Indranil Pan, Lachlan R Mason, and Omar K Matar. 2022. Data-centric Engineering: integrating simulation, machine learning and statistics. Challenges and opportunities. *Chemical Engineering Science* 249 (2022), 117271.
- [53] Qian Pan, Zahra Ashktorab, Michael Desmond, Martin Santillan Cooper, James Johnson, Rahul Nair, Elizabeth Daly, and Werner Geyer. 2024. Human-Centered Design Recommendations for LLM-as-a-Judge. *arXiv preprint arXiv:2407.03479* (2024).
- [54] Kishore Papineni, Salim Roukos, Todd Ward, and Wei jing Zhu. 2002. BLEU: a Method for Automatic Evaluation of Machine Translation. 311–318.
- [55] Chanjun Park and Hyeonwoo Kim. 2024. Understanding llm development through longitudinal study: Insights from the open ko-llm leaderboard. *arXiv preprint arXiv:2409.03257* (2024).
- [56] David KH Pasaribu, Agung Dewandaru, and GA Putri Saptawati. 2024. Development of LLM-Based System for IT Talent Interview. In *2024 IEEE International Conference on Data and Software Engineering (ICoDSE)*. IEEE, 108–113.
- [57] Carol Passos, Daniela S Cruzes, Tore Dybå, and Manoel Mendonça. 2012. Challenges of applying ethnography to study software practices. In *Proceedings of the ACM-IEEE international symposium on Empirical software engineering and measurement*. 9–18.
- [58] Ji-Lun Peng, Sijia Cheng, Egil Diau, Yung-Yu Shih, Po-Heng Chen, Yen-Ting Lin, and Yun-Nung Chen. 2024. A survey of useful llm evaluation. *arXiv preprint arXiv:2406.00936* (2024).
- [59] Xiang Qi and Junnan Yu. 2025. Participatory Design in Human-Computer Interaction: Cases, Characteristics, and Lessons. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–26.
- [60] Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. 2023. Fine-tuning aligned language models compromises safety, even when users do not intend to. *arXiv preprint arXiv:2310.03693* (2023).
- [61] Mohi Reza, Ioannis Anastasopoulos, Shreya Bhandari, and Zachary A Pardos. 2025. PromptHive: Bringing subject matter experts back to the forefront with collaborative prompt engineering for educational content creation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [62] Shaolun Ruan, Rui Sheng, Xiaolin Wen, Jiachen Wang, Tianyi Zhang, Yong Wang, Tim Dwyer, and Jiannan Li. 2025. Qualitative Study for LLM-assisted Design Study Process: Strategies, Challenges, and Roles. *arXiv preprint arXiv:2507.10024* (2025).
- [63] Malik Sallam. 2023. ChatGPT utility in healthcare education, research, and practice: systematic review on the promising perspectives and valid concerns. In *Healthcare*, Vol. 11. MDPI, 887.
- [64] Souvika Sarkar, Mohammad Fakhruddin Babar, Monowar Hasan, and Shubhra Kanti Karmaker. 2024. LLMs as On-demand Customizable Service. *arXiv preprint arXiv:2401.16577* (2024).
- [65] Devansh Saxena, Karla Badillo-Urquiola, Pamela J. Wisniewski, and Shion Guha. 2021. A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 348 (Oct. 2021), 41 pages. doi:10.1145/3476089
- [66] Chirag Shah. 2025. From Prompt Engineering to Prompt Science with Humans in the Loop. *Commun. ACM* 68, 6 (2025), 54–61.
- [67] Shreya Shankar, JD Zamfirescu-Pereira, Björn Hartmann, Aditya G Parameswaran, and Ian Arawjo. 2024. Who Validates the Validators? Aligning LLM-Assisted Evaluation of LLM Outputs with Human Preferences. *arXiv preprint arXiv:2404.12272* (2024).
- [68] Yuchen Shao, Yuheng Huang, Jiawei Shen, Lei Ma, Ting Su, and Chengcheng Wan. 2024. Are LLMs Correctly Integrated into Software Systems? *arXiv preprint arXiv:2407.05138* (2024).
- [69] Helen Sharp, Yvonne Dittrich, and Cleidson RB De Souza. 2016. The role of ethnographic studies in empirical software engineering. *IEEE Transactions on Software Engineering* 42, 8 (2016), 786–804.
- [70] Lin Shi, Chiyu Ma, Wenhua Liang, Xingjian Diau, Weicheng Ma, and Soroush Vosoughi. 2024. Judging the judges: A systematic study of position bias in llm-as-a-judge. *arXiv preprint arXiv:2406.07791* (2024).
- [71] Yifan Song, Guoyin Wang, Sujian Li, and Bill Yuchen Lin. 2024. The good, the bad, and the greedy: Evaluation of llms should not ignore non-determinism. *arXiv preprint arXiv:2407.10457* (2024).
- [72] Chongyan Sun, Ken Lin, Shiwei Wang, Hulong Wu, Chengfei Fu, and Zhen Wang. 2024. LalaEval: A Holistic Human Evaluation Framework for Domain-Specific Large Language Models. *arXiv preprint arXiv:2408.13338* (2024).
- [73] Annalisa Szymanski, Brianna L Wimer, Oghenemaro Anuyah, Heather A Eicher-Miller, and Ronald A Metoyer. 2024. Integrating Expertise in LLMs: Crafting a Customized Nutrition Assistant with Refined Template Instructions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–22.
- [74] Annalisa Szymanski, Noah Ziems, Heather A Eicher-Miller, Toby Jia-Jun Li, Meng Jiang, and Ronald A Metoyer. 2025. Limitations of the llm-as-a-judge approach for evaluating llm outputs in expert knowledge tasks. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*. 952–966.
- [75] Kits AI Team. 2025. Friend or faux: The ethics of Ai Voice Training and why it matters. <https://www.kits.ai/blog/ai-voice-cloning-ethics>
- [76] Shubo Tian, Qiao Jin, Lana Yeganova, Po-Ting Lai, Qingqing Zhu, Xiuying Chen, Yifan Yang, Qingyu Chen, Won Kim, Donald C Comeau, et al. 2023. Opportunities and challenges for ChatGPT and large language models in biomedicine and health. *Briefings in Bioinformatics* 25, 1 (2023).
- [77] Ruth Vo, M Smith, and N Patton. 2021. *Journal of Human Nutrition and Dietetics* 34, 1 (2021), 124–133.
- [78] Steven Euijong Whang, Yuji Roh, Hwanjun Song, and Jae-Gil Lee. 2023. Data collection and quality challenges in deep learning: A data-centric ai perspective. *The VLDB Journal* 32, 4 (2023), 791–813.
- [79] Xingjiao Wu, Luwei Xiao, Yixuan Sun, Junhang Zhang, Tianlong Ma, and Liang He. 2022. A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems* 135 (2022), 364–381.
- [80] Lixuan Zhang, Katelyn Wilson, and Clinton Amos. 2025. The rise of AI art: A look through digital artists' eyes. *First Monday* (2025).
- [81] Yue Zhang, Ming Zhang, Haipeng Yuan, Shichun Liu, Yongyao Shi, Tao Gui, Qi Zhang, and Xuanjing Huang. 2024. LlmEval: A preliminary study on how to evaluate large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 19615–19622.
- [82] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).
- [83] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems* 36 (2024).
- [84] Hongyin Zhu. 2024. Challenges and responses in the practice of large language models. *arXiv preprint arXiv:2408.09416* (2024).
- [85] Douglas Zytco, Pamela J. Wisniewski, Shion Guha, Eric PS Baumer, and Min Kyung Lee. 2022. Participatory design of AI systems: opportunities and challenges across diverse users, relationships, and application domains. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–4.