

Leveraging Large Language Models to Create Learner Personas for Training Design

Lovett, N., Yang, M., Herman, K., Li, B., Sosonkina, M., Purwanto, W., Jiang, P., Wu, M. H.

This case examines the innovative use of Large Language Models (LLMs) to generate learner personas for developing learner-centered cybersecurity training materials when direct access to initial learner data is not available. The team developed a nine-stage iterative process for creating and refining AI-generated personas to address this constraint, integrating ethical review, stakeholder feedback, and action research principles. The process expanded upon Kouprie and Visser's (2009) empathic design framework to ensure cultural responsiveness and mitigate potential biases in LLM outputs. Through multiple refinement cycles, initial generic personas evolved into detailed, context-rich archetypes which informed the development of effective and context-responsive training materials. The case demonstrates how AI-generated personas can serve as valuable tools for instructional design in emerging technical domains when developed through rigorous iterative processes and ethical oversight. This approach offers insights for educational initiatives facing similar challenges in understanding diverse learner needs before direct audience data becomes available.

Introduction

The proliferation of sophisticated cyber threats, combined with the critical role of high performance computing (HPC) and artificial intelligence (AI) in cybersecurity, highlights the urgent need for a skilled workforce capable of leveraging advanced capabilities (ISC2, 2024). HPC and AI are key components of cyberinfrastructure (CI), which is defined as the collection of hardware, software, networks, data and people which are linked together to enable transformative scientific research and discovery (Atkins et al., 2003; Stewart et al., 2010). The current state of cybersecurity education has not prepared a workforce skilled in leveraging HPC and AI in their professions. The Train-the-Trainer Approach to Fostering CI- and Data-Enabled Research in CyberSecurity (T3-CIDERS) project directly addresses this challenge. T3-CIDERS aims to equip educators, researchers, and practitioners to integrate cutting-edge CI techniques into cybersecurity research and education. By doing so, it seeks to accelerate the adoption of HPC-enabled methods, foster advanced problem-solving skills, and ultimately bolster global cyber-resilience.

Designing effective educational interventions in a rapidly evolving domain requires carefully understanding learners' diverse backgrounds, technical proficiencies, and learning contexts. Learner-centered design is a practical necessity in a field as varied and dynamic as cybersecurity. Learners differ widely in linguistic backgrounds, cultural contexts, resource constraints (e.g., limited access to HPC clusters), and prior exposure to advanced computational tools (Qawasmeh et al., 2024). Without considering this heterogeneity, curricula risk misaligning with learners' contexts, inadvertently excluding particular groups of learners, and failing to foster the broad and diverse pool of talent needed to combat cyber threats effectively.

Persona-based approaches in instructional design have long been employed in user experience (UX) design and increasingly appear in educational design contexts (Salminen et al., 2022). By creating vivid, context-specific learner archetypes, designers can empathize with prospective

users, anticipate their needs, and tailor instruction to align with diverse motivators, constraints, and skill sets. In an ideal scenario, persona creation would rely heavily on direct learner data, such as interviews, surveys, and observational studies. However, the T3-CIDERS project faced a key constraint early on: marketing and recruitment were still underway, and direct access to prospective learners remained limited.

This situation led to an innovative design decision: employing generative Large Language Models (LLMs), such as GPT-4 and Claude, to generate provisional learner personas. Rather than waiting for empirical data, the team prompted LLMs to synthesize plausible learner archetypes from textual cues, domain knowledge, and inferential logic. This approach is supported by recent research demonstrating LLMs' capability to predict human responses in experimental settings accurately. Hewitt et al. (2024) found LLM-derived predictions strongly correlate with actual experimental effects ($r = 0.85$) across diverse demographic groups, and maintain high accuracy even for scenarios absent from training data. These AI-generated personas offered a starting point for instructional designers. They provided a means to conceptualize different learners who might engage with cybersecurity-focused HPC training and to begin shaping effective, context-aware, and relevant instruction before actual learners were fully identified. To ensure ethical and responsible use of AI-generated personas, the team developed a novel nine-stage methodological process incorporating prompt refinement, ethical review, stakeholder feedback, and iterative improvement.

Utilizing AI-generated personas as a foundational design element presents challenges and ethical considerations. While LLMs offer efficiency in persona generation, they risk reproducing biases present in their training data (Holmes et al., 2022). In security-relevant domains, these biases might manifest as personas who disproportionately associate advanced technical capabilities with learners from well-resourced institutions, underestimate the proficiency of educators from regions with historically limited representation, or frame multilingualism as an obstacle rather than an asset in global cybersecurity collaboration. Although Hewitt et al. (2024) reported encouraging evidence regarding LLMs' ability to provide relatively consistent predictions across demographic groups, potential biases still require careful attention. Addressing these concerns necessitates a rigorous, iterative design process involving prompt refinement, ethical review, stakeholder feedback, and careful alignment with project goals.

This article documents how T3-CIDERS designed and integrated AI-generated personas into the instructional design of an HPC training program for the cybersecurity domain, responding to calls for rigorous evaluations of persona creation methodologies (Chapman & Milham, 2006; Vestergaard et al., 2016). The narrative details a nine-stage persona creation process, each carefully vetted through naturalistic inquiry (Lincoln & Guba, 1985) and action research principles (Hinchey, 2008). This process creates tension between learner-centered instructional ideals and the practical challenges of teaching in a volatile, high-stakes setting. In doing so, the team offers insights into how AI tools can be leveraged responsibly to inform instructional design decisions, even when direct audience data is incomplete or evolving.

The following sections detail the team's nine-stage methodology for generating, refining, and implementing AI-created personas. They then discuss how these personas informed the development of effective, engaging training materials. Through this design case, we will contribute to the conversation about the thoughtful integration of AI into instructional design and highlight the potential of persona-based methods to foster adaptive and effective cybersecurity education.

Project Context

The T3-CIDERS project, supported by the U.S. National Science Foundation, was conceived to address a pressing challenge in contemporary cybersecurity education: the need to rapidly cultivate teaching expertise that integrates advanced CI skills (including HPC, big data analytics, and AI/machine learning), into -enabled threat detection. While many educational programs have begun to incorporate basic cybersecurity principles, fewer can equip instructors with the specialized computational and analytical proficiencies required to handle the complexity of modern cyber threats (Almoughem, 2023; Qawasmeh et al., 2024; Ricci et al., 2020).

Goals and Scope of T3-CIDERS

T3-CIDERS aims to establish a foundational layer of cybersecurity educators, both seasoned faculty members and emerging trainers, who can disseminate CI-driven pedagogical methods across various institutional contexts. By focusing on train-the-trainer, or the training of trainers (ToT) methods, T3-CIDERS extends its impact well beyond a single cohort of learners. In essence, it strives to create a multiplier effect: instructors trained through T3-CIDERS will adapt and integrate HPC-enabled cybersecurity modules into their curricula, reaching larger numbers of students over time.

The project's scope is intentionally broad, reflecting the landscape of cybersecurity education itself. Participating instructors may hail from research-intensive universities looking to infuse cutting-edge HPC tools into graduate-level courses, community colleges aiming to enhance employability skills in associate degree programs, or industry-based training centers focused on upskilling their workforce. Additionally, T3-CIDERS recognizes the significance of interdisciplinary collaboration, encouraging participants to integrate CI skills not only in computer science or information systems programs but also in fields such as bioinformatics, engineering, and data science where cybersecurity threats now permeate (Admass et al., 2024).

The Train-the-Trainer Approach

The ToT model is a well-established framework in professional development and capacity-building efforts, widely used across sectors such as healthcare (Mormina & Pinder, 2018). In the context of T3-CIDERS, the ToT model involves selecting instructors with some foundational knowledge of cybersecurity principles and providing them with intensive training to incorporate HPC-driven methodologies into their teaching. These “future trainers” then return to their home institutions equipped with technical competencies and pedagogical strategies, instructional materials, and formative assessment techniques tailored to HPC-oriented cybersecurity learning.

This approach aligns with a growing recognition that introducing new technical content is insufficient. Educators must understand how to scaffold complex concepts, use authentic simulations, integrate problem-based learning activities, and ensure that instruction remains accessible to a wide range of learners (McGettrick, 2013; Ricci et al., 2020). By focusing on the educators themselves, T3-CIDERS leverages their existing pedagogical expertise, enabling a more sustainable and widespread dissemination of advanced cybersecurity instruction.

The ToT model particularly amplified the need for well-developed personas. Unlike programs designed for a single audience, T3-CIDERS creates a cascading effect where each future trainer subsequently adapts materials across varied institutional contexts (Mormina & Pinder, 2018). This multiplier effect created the need for personas that could represent not just varied individual characteristics, but also the different educational environments, resource constraints, and student populations the program's trainers would encounter when implementing their training.

Challenges in Accessing Target Audiences

Despite the ToT model's potential, designing practical training modules for an unknown or partially defined audience presents significant challenges. Early in the T3-CIDERS project, recruitment and outreach efforts were still underway, and direct access to participants, potential trainers who would eventually shape HPC-enabled cybersecurity education, was limited. Without concrete information about their backgrounds, technical proficiencies, institutional contexts, or resource constraints, the design team faced the classic "designing in the dark" scenario. This condition risked creating materials which were either too advanced or too rudimentary, too specialized in specific computational techniques, or not aligned with learners' actual motivations and time constraints.

While challenges in accessing target audiences exist in many educational contexts, T3-CIDERS faced particularly acute constraints when compared to other cybersecurity education programs. Most U.S. cybersecurity-educator programs serve relatively homogenous cohorts and evolve their materials over multiple cycles. For example, GenCyber camps recruit K-12 participants months ahead of delivery and run fixed curricula (Childers et al., 2023), while community college workshops address stable audiences with annual content updates (Burrell et al., 2019). In contrast, T3-CIDERS issued a late-stage national call to faculty with widely varying experience levels and institutional resources, from low-tech rural campuses to HPC-equipped research universities. The grant required a cyberinfrastructure-rich syllabus launched within a single funding year, far faster than typical program timelines and exemplifying the "just-in-time" curriculum pressure increasingly common in cybersecurity education (Murphy et al., 2023). This combination of unknown learners, extreme resource disparity, and compressed development timeline made traditional audience analysis methods impractical.

Conventional methods of gathering participant data, such as surveys or interviews, often require established communication channels and a known pool of respondents. In the fast-moving and globally dispersed domains of HPC training and cybersecurity upskilling, instructors may be scattered across multiple institutions, time zones, and resource environments, making synchronous data collection approaches impractical. As a result, T3-CIDERS' designers had to make informed guesses about the audiences they aimed to serve. Recent scholarship identifies specific competency gaps among cybersecurity instructors that need to be addressed, including: technical fluency with advanced cyber-infrastructure such as a secure cloud or HPC systems (Almoughem, 2023; Murphy et al., 2023), conceptual-pedagogical expertise for scaffolding complex content and running authentic simulations (Ricci et al., 2020), and applied teaching practice, including inquiry-driven lesson design and classroom management for hands-on labs (Burrell et al., 2019; Childers et al., 2023).

Based on these identified gaps and the preliminary research, the team anticipated some participants might be highly experienced in cybersecurity theory but less familiar with HPC tools and parallel programming environments. Others might be adept at handling big data sets but uncertain about integrating machine learning approaches for threat detection. Some future trainers might work in resource-limited institutions without access to HPC clusters or advanced visualization tools. Still, others might be multilingual K-12 facilitators needing culturally responsive pedagogical strategies to engage multilingual or underrepresented student populations. These varied profiles directly informed the team's persona development process, ensuring that each institute module addressed the specific competency needs represented across the persona set.

Without direct learner data to guide the balancing of these variables, the design team risked producing generic or misaligned modules. This tension set the stage for an innovative solution: leveraging Large Language Models (LLMs) to generate provisional learner personas. By creating synthetic yet plausible archetypes of future participants, the designers could begin anticipating

needs and constraints, thereby moving from a vacuum of information to a structured, if tentative, set of learner profiles.

In summary, the T3-CIDERS project's expansive goals, reliance on a train-the-trainer approach, and diverse potential audience created a complex environment for instructional design.

Limitations in early-stage data collection about future participants further complicate efforts to ensure that HPC-enabled cybersecurity training is technically robust and effective across differing learning contexts. Against this backdrop, the decision to explore AI-generated personas emerged, offering a pathway to inform design decisions before the entire learner landscape came into clear focus.

Design Challenge

The central design challenge for the T3-CIDERS project lay in crafting cybersecurity training modules which would be both effective and engaging despite limited access to direct audience data. Designing, engaging, and context-adaptive training modules for cybersecurity education requires careful consideration of multiple dimensions. "Effective" training in this context refers to materials which successfully develop the specific CI-enabled cybersecurity competencies identified as critical gaps in the current workforce (Almoughem, 2023). "Engaging" refers to instruction which employs evidence-based strategies to maintain learner motivation through relevance, challenge, and autonomy (Keller, 1987). This is particularly important given the technical complexity of HPC concepts. By "context-adaptive," we mean materials which can be meaningfully implemented across diverse institutional settings with varying resource availability, from research universities with dedicated HPC clusters to community colleges with limited computational infrastructure (Qawasmeh et al., 2024). These dimensions guided the team's design approach and informed the subsequent development of learner personas.

Early in T3-CIDERS's development, a direct needs assessment was not feasible through traditional methods such as interviews and surveys (Smith & Ragan, 2004). Marketing and recruitment efforts for the program were still underway, and the project team had only broad, generalized knowledge about the diverse pool of potential participants, who could range from seasoned cybersecurity instructors seeking to integrate HPC tools into their curricula, to educators with minimal experience in computational methods, to those operating in resource-constrained settings.

Creating Learner-Centered Cybersecurity Training

The goal of T3-CIDERS was not simply to impart technical knowledge in HPC-enabled cybersecurity; it was to do so in a manner that recognized and accommodated the heterogeneity of participant backgrounds. Cybersecurity education often involves intricate, domain-specific concepts that can be challenging to grasp without appropriate scaffolds (Qawasmeh et al., 2024). Adding HPC and data-intensive methodologies intensifies the complexity. Suppose educators are to effectively teach topics such as parallelized cryptographic computations, machine learning-based threat detection, and large-scale data analysis. In that case, the instructional materials must anticipate a wide range of learner scenarios. Without such foresight, modules risk being inaccessible, unengaging, or disconnected from the practical realities of instructors' teaching contexts.

Inclusivity here extends beyond addressing varying technical skill levels. It involves cultural responsiveness, linguistic considerations, and sensitivity to resource disparities. For instance, some future trainers might work at well-funded universities with readily available HPC clusters, while others may struggle with limited computational resources and bandwidth (Qawasmeh et al.,

2024). Some trainers may be English language learners or serve multilingual student populations, necessitating clear, concise, and easily localizable materials. Others may face institutional constraints such as limited instructional time or competing curricular demands.

Established models for developing existing personas exist, including the leveraging of Kouprie and Visser's (2009) Discover-Immerse-Connect-Detach empathic design process (see Baaki & Maddrell, 2020). However, to the knowledge of the design team, this model has not yet been used without designers having at least some access to end users. This design case aimed to understand the applicability and limitations of this design approach when utilizing AI-generated learner personas.

Rationale for Using AI-Generated Personas

To navigate these complexities, the project team generated provisional learner personas using LLMs. Personas have been recognized for helping designers empathize with users, identify key pain points, and tailor solutions accordingly (Cooper, 1999; Salminen et al., 2022). Typically, persona creation is grounded in empirical data. However, when such data are unavailable or incomplete, LLMs can offer an alternative.

The transition from traditional to AI-generated persona methodologies builds upon established frameworks in instructional design. Traditional persona development has primarily relied on qualitative approaches like interviews and observations (Huynh et al., 2021), with designers creating archetypal representations to enhance empathy in educational design. When direct learner access is limited, researchers have employed "proto-personas" as provisional starting points (Salminen et al., 2020) and developed clustered representations from small educational datasets (Zagallo et al., 2019). More recently, Automatic Persona Generation systems have emerged that use computational methods to create data-driven personas (Jansen et al., 2021), and since 2022, LLMs have introduced new capabilities for persona creation (Salminen et al., 2024).

LLMs synthesize plausible learner archetypes from textual cues, constraints, and domain knowledge provided in carefully crafted prompts. The design team produced personas which highlighted distinct sets of needs and challenges through an iterative process. This approach positioned LLMs as sounding boards for idea generation rather than autonomous creators, aligning with evidence that such refinement processes yield more effective outcomes (Yang et al., 2025). AI-generated personas acted as a heuristic device, shifting the design process from guesswork to a more structured approximation of learner diversity. LLMs rapidly proposed multiple plausible scenarios, enabling the design team to consider varied learner conditions.

Building upon these emerging AI-driven persona methodologies, the team developed a nine-stage process for generating, refining, and implementing AI-created personas. This process was adapted from Kouprie and Visser's (2009) Discover-Immerse-Connect-Detach framework, rather than developed inductively, though the team significantly expanded it to address the unique challenges of creating personas without direct user access. The following section details this methodological approach, beginning with data collection and progressing through multiple iterations of persona refinement and validation.

Initial Design Considerations and Constraints

Despite the potential of AI-generated personas, the team recognized several constraints and considerations upfront. First, LLM outputs inherently reflect the biases and limitations of their training data (Holmes et al., 2022). The design team needed to implement ethical review

processes and multiple iterations to ensure that personas did not reinforce stereotypes, ignore cultural nuances, or present unrealistic expectations of available resources.

Second, without direct validation from actual participants, these personas were, at best, provisional. They offered a starting point for context-responsive design, but the team understood that future refinements would depend on feedback from actual T3-CIDERS trainers once they were recruited. The initial personas were intended to guide early decision-making, informing choices about content complexity, multimedia inclusion, pacing, and support materials. Over time, as empirical data became available, the personas would be revisited and updated to align with the real composition of the learner audience more closely.

Third, the complexity of cybersecurity education, intertwined with HPC and data-intensive methods, demanded careful balancing. On the one hand, the training needed to reflect cutting-edge techniques and anticipate future workforce demands. On the other hand, training had to remain accessible to instructors with limited computational backgrounds or who teach students new to the concept of cyberinfrastructure. The design team recognized persona-based planning would help maintain this balance, ensuring early prototype modules accounted for varying levels of prior knowledge, technical infrastructure, and motivational factors.

The design challenge was creating effective, engaging, HPC-driven cybersecurity training without direct access to future learners. Embracing AI-generated personas as a guiding framework allowed the team to envision multiple participant profiles and consider their distinct constraints, thus laying a foundational approach for the subsequent iterative development and refinement of the training materials.

Persona Design Process

The persona design process implemented within the T3-CIDERS project was carefully structured to ensure rigor, transparency, and iterative refinement. Drawing on action research principles (Hinchey, 2008), empathic design-based research methodologies (Kouprie & Visser, 2009), and frameworks for trustworthy qualitative inquiry (Lincoln & Guba, 1985), the team developed a nine-stage procedure. Each stage contributed to progressively shaping AI-generated learner personas from initial conjectures into more contextually grounded and ethically vetted representations of future trainees.

Although the team's workflow adapts Kouprie and Visser's (2009) empathic design cycle, it is also grounded in instructional design literature. Stages 1 through 3 reflect Stefaniak and Baaki's (2013) layered audience analysis and encouraged designers to move beyond surface demographics when direct learner contact is limited (Stefaniak, 2015). The iterative LLM refinements in Stages 2 and 5 through 7 follow the prototype-revision logic of design-based research (Wang & Hannafin, 2005). The team's attention to multiple ways trainees might approach cyber-infrastructure implements principles from variation theory (Marton & Booth, 1997), and the development of provisional personas aligns with phenomenographic approaches to creating learner archetypes in educational contexts (Huynh et al., 2021).

The nine-stage process also incorporates elements from established instructional design frameworks. The early stages align with the Analysis phase of the ADDIE model (Analysis, Design, Development, Implementation, Evaluation), where thorough learner analysis traditionally occurs (Branch, 2009). However, the team's approach extends typical ADDIE analysis by using AI-driven personas to address gaps when direct learner data is unavailable. Similarly, the team's methodology reflects Backward Design principles (Wiggins & McTighe, 2005) by beginning with the desired training outcomes for cybersecurity educators and working backward to develop personas which would inform appropriate learning experiences and instructional materials. In this way, the AI-generated personas extend, rather than replace, existing learner-centered, design-

based research and variation-theory practices, demonstrating that the nine-stage model builds on a recognized lineage in educational design.

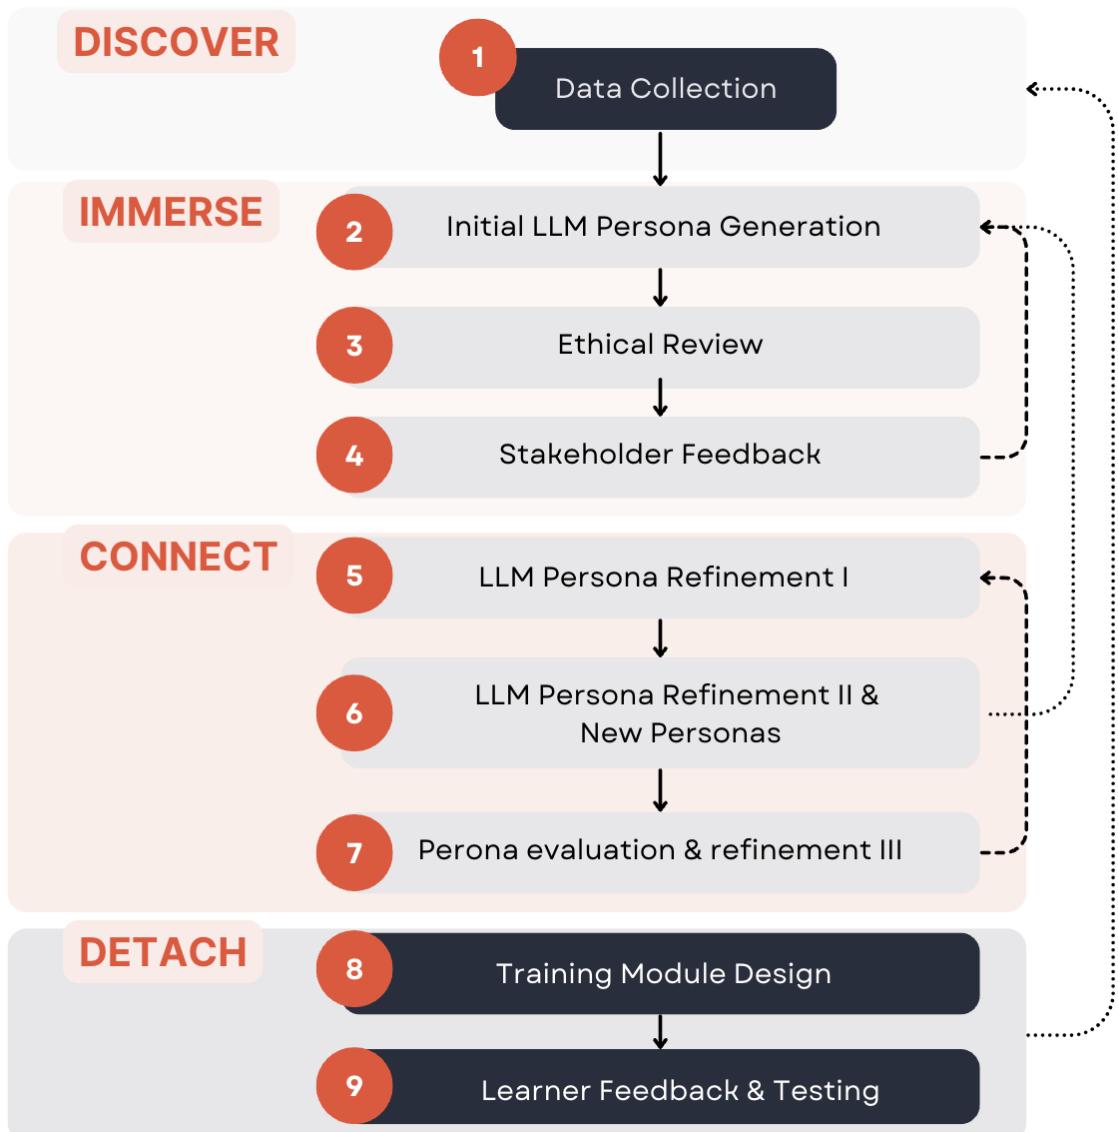
Nine-Stage Process Overview

As previously mentioned, the design team began with Kouprie and Visser's (2009) Discover-Immerse-Connect-Detach framework for empathic persona design. Traditionally, this model breaks the design process into four distinct phases, each with an emphasis on empathy. However, a more iterative interpretation of this framework proved necessary when using AI to generate learner personas.

Figure 1 visually represents the nine-stage process, illustrating how each stage informed the next and highlighting the feedback loops which permitted continual improvement of the personas:

Figure 1

Visual representation of the 9-Step Empathic Design Process



Note. The process flows through Discover, Immerse, Connect, and Detach (DICD) phases, with arrows indicating iterative feedback loops between stages.

This nine-stage process expands the original DICD framework as follows: (1) in the discover phase, designers name learners and pique curiosity by beginning to explore their unique characteristics; (2) in the immerse phase, designers dig deeper, leaving the design world to uncover and even participate in the lived experiences of identified learners; (3) in the connect phase, designers reflect on their own lived experiences to find and refine understanding; (4) and in the detach phase, designers step back into their role as designers, reflect on the design as is, consider alternate views, and put forth a nuanced design. However, when using AI to generate learner personas, a more iterative interpretation of this framework proved necessary, with the design team expanding on the existing model via the following nine-stage process: DISCOVER: (1) Data collection; IMMERSE: (2) Initial LLM persona generation, (3) Ethical review, (4) Stakeholder feedback; CONNECT: (5) LLM persona refinement I, (6) LLM persona refinement II and new persona generation, (7) Persona evaluation and refinement III; DETACH: (8) Training Design; and a return to DISCOVER: (9) Learner feedback and pilot testing. Although linear in presentation, the process was cyclical in practice. Insights from later stages often informed revisions to earlier assumptions, embodying the iterative nature of design-based research and suggesting that the reflection typically exhibited in the latter phases of empathic persona design was necessary throughout the design process.

The rationale for this expanded multi-stage process lay in the complexity of the design challenge. Rather than relying solely on AI outputs or designer intuition, this approach integrated multiple sources of feedback and oversight. Ethical review steps ensured cultural responsiveness and bias mitigation (Holmes et al., 2022), stakeholder feedback connected personas to real-world educational conditions, and pilot testing provided preliminary validation. By layering these methods, the team sought to balance the creative potential of LLMs with grounded, context-driven scrutiny.

DISCOVER: Data Collection

Lacking direct access to learners, the team triangulated data from three critical sources: (1) cybersecurity workforce reports identifying specific HPC skill gaps; (2) internal project documentation, including evaluations from previous related NSF projects; and (3) targeted consultations with cybersecurity faculty about teaching constraints. A key methodological decision was maintaining an explicit audit trail, distinguishing literature-based assumptions from speculative elements, enabling later revision when participant data became available. This foundation provided essential domain constraints for initial LLM prompt construction and subsequent evaluation of persona plausibility.

IMMERSE: Initial LLM Persona Generation, Ethical Review, and Stakeholder Feedback

The initial prompt design represented a critical methodological decision: the team specified four deliberately diverse roles (community college professor, HPC researcher, PhD student, and K-12 educator) to ensure varied institutional contexts and resource environments. The team primed the LLMs with T3-CIDERS project documentation to ground outputs in project goals before requesting personas with specific attributes (job role, experience level, and primary motivation).

The resulting sketches provided foundational archetypes but required significant refinement. For example, the initial "Cybersecurity Professor at a Community College" lacked contextual specificity about teaching environment, technological constraints, and cultural context—elements crucial for instructional design decisions.

The ethical review revealed systematic issues requiring intervention. Most significantly, early personas exhibited three problematic patterns: (1) portraying non-native English speakers as disadvantaged rather than linguistically resourceful; (2) framing resource constraints as deficits rather than contexts for innovation; and (3) implicit gender biases in technical roles. The team implemented a structured bias-mitigation protocol with explicit prompt revisions such as: "Present multilingualism as an asset that enables access to broader research communities" and "Frame resource limitations as contexts that foster creative pedagogical approaches."

A stakeholder review by cybersecurity faculty and HPC education experts identified critical gaps absent from the AI-generated personas. Three key omissions emerged: none addressed asynchronous learning needs for instructors with administrative responsibilities; none reflected the reality of teaching in multilingual contexts; and none operated in settings with intermittent internet connectivity. These expert insights directly shaped subsequent refinements, demonstrating how domain expertise complements AI generation.

CONNECT: Strategic Persona Refinement and Challenge Alignment

Initial version: "A mid-career community college instructor with general IT experience, curious about HPC but uncertain how to integrate parallel programming tasks into a basic cybersecurity course."

After expert-informed refinement: "A mid-career, multilingual community college instructor with over 10 years of experience teaching IT and introductory cybersecurity courses. They are fluent in Spanish and English and teach a diverse student body, many of whom are non-native English speakers. While they have a strong foundation in traditional cybersecurity topics, they are new to HPC and eager to learn how to incorporate simple, asynchronous HPC examples into their courses. They are particularly interested in using cloud-based HPC resources due to their institution's limited on-site computing infrastructure and their students' intermittent internet connectivity."

This evolution demonstrates how critical instructional design expertise transformed generic descriptors into contextually rich profiles directly informing design decisions.

A methodological breakthrough occurred when participant survey data became available, allowing validation and expansion of the synthetic personas. After analyzing this data with LLM assistance, the team identified gaps in representation that led to the development of five additional personas, ensuring more comprehensive coverage of learner diversity.

A pivotal design decision was aligning each persona with specific cybersecurity challenges they would encounter in their teaching contexts. This decision transformed personas from demographic descriptions into instructional design tools that directly informed content complexity, resource requirements, and pedagogical approaches. For example, the team matched teaching-focused personas in resource-limited environments with challenges requiring simplified simulations and offline resources. In contrast, research-oriented personas were aligned with advanced threat detection scenarios requiring hands-on code development.

DETACH and DISCOVER 2: Implementation, Validation, and Reflection

The refined personas directly informed three critical instructional design decisions: (1) content format adaptation—developing offline-compatible PDFs, chunked videos, and self-paced assignments for personas facing bandwidth constraints; (2) platform selection—choosing the K-12 educator version of Canvas to ensure accessibility across diverse institutional settings; and (3) pedagogical approach—incorporating simulated teaching sessions where trainers could practice teaching one fully developed activity, addressing the varied teaching experience levels identified across personas.

After the summer institute, the team implemented a structured feedback process where participants evaluated the personas. Among the feedback received, one volunteer instructor unit (one professor and an undergraduate student) who matched several attributes of the personas provided particularly valuable insights. This team conducted a local training session and then walked through their implementation with the research team. This review process provided an initial opportunity to evaluate the persona-informed design decisions. By analyzing the pilot users' interactions with the training materials, the team gathered insights on whether the personas' predicted challenges and needs aligned with actual participant experiences.

The team also used LLM assistance for meta-analysis of persona evolution, having it summarize how iterative refinement enhanced cultural sensitivity and scenario relevance. While treated with appropriate caution, this reflective process helped articulate the transformation from abstract AI outputs to practical design tools that effectively supported diverse learner needs in cybersecurity contexts.

Detailed Persona Example: Evolution of One Persona

While the preceding sections have outlined a broad, iterative methodology for generating and refining AI-driven learner personas, a more granular, person-centered illustration can reveal how these processes unfold in practice. This section provides a detailed case tracing the evolution of a single persona, Dr. Aisha Nguyen, from her initial, AI-generated inception to her final, context-rich form. By examining how Dr. Nguyen's profile was iteratively adjusted, refined, and validated, the team gained insight into how theoretical principles, ethical considerations, and stakeholder input converge to produce a persona capable of guiding real instructional design decisions.

Initial Generation

During the initial Large Language Model (LLM) persona generation phase, Dr. Nguyen first emerged as a generalized character sketch. The process began by providing the LLM with the complete publicly available T3-CIDERS grant abstract to establish necessary context about the program's goals, audience, and structure. After ensuring the model had this foundational information, the team presented a straightforward prompt which specified the key elements the team wanted in each persona, as shown in Figure 2.

Figure 2

Initial Persona Generation Prompt

"Based on the grant abstract, outline three distinct learner personas for the T3-CIDERS program. Include these details for each persona:
Job Title/Role: (e.g., Cybersecurity Professor, HPC Systems Researcher, Ph.D. Student in Computational Science)
Experience Level: (e.g., Novice with CI but expert in the field, Intermediate CI knowledge, etc.)
Primary Motivation: Why are they interested in this program? What do they hope to gain?"



Based on the T3-CIDERS program described in the grant abstract, here are **three** distinct learner personas that might be interested in participating in the program:

Persona 1: Cybersecurity Professor

- **Job Title/Role:** Cybersecurity Professor at a Community College
- **Experience Level:** Expert in cybersecurity principles and practices but novice with advanced cyberinfrastructure (CI) tools and techniques.
- **Primary Motivation:** The professor aims to integrate advanced CI concepts and tools into their curriculum to enhance the learning experience of their students and to ensure that graduates are well-prepared for the demands of the cybersecurity industry. They are particularly interested in the Train-the-Trainer model to learn and then teach CI-enabled cybersecurity research methods, hoping to bridge the gap between traditional cybersecurity education and the emerging demands for CI proficiency.

Persona 2: HPC Systems Researcher

- **Job Title/Role:** HPC (High-Performance Computing) Systems Researcher in a National Laboratory
- **Experience Level:** Advanced knowledge of HPC systems and some experience with CI, but



As illustrated in the screenshot, the initial prompt built upon the grant context by requesting three distinct learner personas with specific characteristics: Job Title/Role, Experience Level, and Primary Motivation. This approach allowed the LLM to generate personas with clear professional contexts and motivations aligned with the T3-CIDERS program objectives.

In response to this prompt, Dr. Nguyen initially appeared as a “cybersecurity professor at a community college,” holding a Ph.D. in Computer Science and interested in integrating advanced cyberinfrastructure (CI) tools, such as high-performance computing (HPC) and machine learning, into her courses.

At this early stage, her persona lacked detail on the nature of her teaching environment, the challenges she faced, or the cultural and infrastructural contexts influencing her pedagogical strategies. For example, while the persona mentioned her interest in HPC-enabled cybersecurity, the persona did not specify whether she had access to HPC resources, what languages her students spoke, or how rapidly changing threat landscapes affected her professional development. As a result, Dr. Nguyen served primarily as a placeholder, offering a glimpse into the potential complexity of instructors who might join T3-CIDERS training but had not yet conveyed the authenticity or granularity needed to inform design decisions.

Though the initial prompting approach was simple compared to more advanced prompt engineering techniques, it proved effective in establishing foundational personas that could be iteratively refined through subsequent prompts and stakeholder feedback. This example demonstrates how even straightforward prompts can yield useful results when part of a systematic, multi-stage refinement process, especially when the model has been provided with relevant contextual information first.

Ethical Considerations and Adjustments

Subsequent ethical reviews and prompt refinements targeted common pitfalls associated with AI-generated personas (Holmes et al., 2022). Early outputs sometimes defaulted to simplified stereotypes. To counter this tendency, the team introduced explicit instructions in prompts, encouraging the LLMs to highlight Dr. Nguyen's multilingual abilities, emphasize her unique perspective as a professor balancing administrative duties and teaching responsibilities, and portray resource constraints without reducing her agency or innovation.

For Dr. Nguyen, this meant reimagining her background to include an international dimension, such as noting that she was born in Vietnam and later immigrated to the United States, thereby contributing a multicultural lens to her teaching. Ethical refinement also ensured she was not depicted as "struggling" due to language differences, but leveraging her bilingual capabilities to access a broader range of HPC and cybersecurity materials. Additionally, refinements addressed technological challenges, e.g., intermittent internet connectivity or limited HPC clusters, framed not as insurmountable barriers but as conditions prompting creative solutions.

Stakeholder Feedback and Refinement.

After initial ethical adjustments, Dr. Nguyen's persona was shared with T3-CIDERS leadership and subject matter experts, whose input guided further refinements. Stakeholders noted that while Dr. Nguyen now appeared as a culturally nuanced figure, additional specificity could strengthen her persona's connection to the T3-CIDERS program objectives. For instance, they suggested incorporating references to widely used cybersecurity tools that would challenge her HPC integration efforts, such as Wireshark for analyzing network traffic or Metasploit for penetration testing. By doing so, Dr. Nguyen's persona illustrated a realistic trajectory: a professor who must keep pace with industry-standard tools, incorporate these into her curriculum, and ensure that her students gain hands-on experiences aligned with current workforce demands.

Stakeholders also recommended highlighting Dr. Nguyen's bridging role between academia and industry partners. This enhancement was achieved by adding details about her involvement in local workshops and training sessions, bringing academic knowledge into practical alignment with industry needs. Moreover, the persona now articulated how Dr. Nguyen navigated rapid technological changes, e.g., the emergence of AI-driven threat detection methods, and translated these developments into updated course materials for her students.

Final Version and Its Application in Training Design

The final refined version of Dr. Aisha Nguyen's persona portrayed a cybersecurity professor who captured more nuances. Figure 3 shows a screenshot of one section of this persona.

Figure 3

Dr. Aisha Nguyen's persona



Dr. Aisha Nguyen

Cybersecurity Professor

#Education

#Cybersecurity

#HPC

Education

Ph.D. in Computer Science (Network Security and AI Applications)

Teaching Environment

Teaching-focused institution with limited HPC resources

Technical Expertise

- Open-source cybersecurity tools
- HPC-enabled threat modeling

Cultural Competencies

- International background
- Multilingual capabilities

Professional Development Needs

- Ongoing training in HPC techniques and ML frameworks
- Flexible, asynchronous learning opportunities
- Industry collaboration for curriculum relevance

This final persona influenced training design in several tangible ways. Recognizing Dr. Nguyen's resource limitations prompted the development of offline-compatible training materials and modular, chunked content that could be reviewed in short intervals. Her interest in bridging academia and industry partnerships led the design team to incorporate scenario-based learning activities centered around real-world cybersecurity challenges. Her international background encouraged the inclusion of culturally responsive pedagogical strategies during the summer institute, ensuring that learners in her context could better connect with the material.

Equally importantly, the process of refining Dr. Nguyen's persona demonstrated the iterative, reflective nature of the design approach. Initial LLM outputs sparked creative thinking, ethical oversight ensured fairness and cultural sensitivity, and stakeholder feedback provided domain-specific granularity. This interplay resulted in a persona who was not merely a fictional character but a tool informing course development, material selection, and balancing complexity and accessibility. In other words, Dr. Nguyen's persona became a lens through which the T3-CIDERS team could anticipate and address the multifaceted realities of teaching HPC-driven cybersecurity content.

Examining the evolution of Dr. Aisha Nguyen's persona underscores the value of treating AI-generated personas as evolving constructs refined through ethical considerations, expert input, and iterative testing. The final persona stands as a testament to how theory, technology, and human judgment can converge to create more technically adaptable, contextually grounded, and impactful instructional design.

Challenges and Solutions

Generating and refining AI-driven learner personas within the T3-CIDERS project was neither linear nor without obstacles. Various challenges emerged as the team iterated through data

collection, persona generation, ethical reviews, stakeholder feedback, and subsequent refinements. Identifying these hurdles and devising corresponding solutions was critical to maintaining the trustworthiness and utility of the personas. This section discusses some of the major issues encountered and the strategies implemented to overcome them, illustrating how action research principles and design-based research methodologies guided responsive, reflective practice (Hinchey, 2008; Wang & Hannafin, 2005).

Challenge 1: Limited Baseline Data

Initially, T3-CIDERS lacked direct empirical information about future participants' backgrounds, skill levels, and contextual realities. Without such data, persona creation ran the risk of being too generic or grounded in assumptions that might not accurately reflect the target audience. This scarcity complicated efforts to ensure that modules would address the correct level of HPC complexity or respond to cultural and linguistic diversity.

Solution

To mitigate this challenge, the design team turned to multiple indirect data sources, including domain literature, project documentation, and expert consultations. These sources informed the initial prompting of LLMs and guided iterative refinements. By maintaining an audit trail of assumptions and documenting their origins, the team ensured that future empirical data from actual learners could be used to validate or adjust personas.

Challenge 2: Risk of Bias in LLM Outputs

Advanced AI models like GPT-4 and Claude can inadvertently incorporate biases present in their training data (Holmes et al., 2022). Early persona drafts sometimes defaulted to stereotypes, such as implying that non-native English speakers might struggle more with technical content or that resource-limited contexts inherently indicated lesser motivation. If left unaddressed, these biases could distort the design team's understanding and lead to materials that were insensitive or misaligned with learner strengths.

Solution

The explicit integration of an ethical review stage helped identify and correct biased narratives. Consider this comparative example from the work with Dr. Nguyen's Persona:

Initial LLM output:

"As a non-native English speaker, Dr. Nguyen sometimes struggles to keep up with rapidly changing technical terminology in cybersecurity documentation."

Revised prompt approach:

"Present multilingualism as an asset rather than a deficit. Highlight how diverse linguistic and cultural backgrounds enrich cybersecurity education and research."

Result after revision:

"Dr. Nguyen's multilingual capabilities allow her to access a broader range of international cybersecurity resources and research, enriching her curriculum with diverse perspectives that monolingual instructors might miss."

This iterative approach, guided by culturally responsive pedagogy, ensured that personas emerged as strength-based profiles, treating learners as capable and adaptive agents rather than reducing them to simplistic labels.

Challenge 3: Ensuring Cultural and Contextual Authenticity

Even with ethical considerations addressed, early personas lacked the granularity and cultural responsiveness needed to inform effective instructional design. Issues such as resource availability, institutional support, and time constraints remained too vaguely defined, limiting the personas' practical value for tailoring training modules.

Solution

Stakeholder feedback, including input from T3-CIDERS leadership and subject matter experts, was instrumental in enriching personas with realistic scenarios. These experts suggested referencing commonly used cybersecurity tools, acknowledging asynchronous learning needs for instructors with administrative duties, and reflecting global linguistic diversity. Incorporating these suggestions produced more authentic and contextually resonant personas. Additionally, iterative refinements emphasized scenario-based details, such as regional data protection laws or HPC cluster availability, bridging abstract persona attributes with actionable design decisions.

Challenge 4: Balancing Technical Complexity and Accessibility

The complexity of cybersecurity education and HPC-driven workflows threatened to produce modules too advanced for some users or too basic for others. Without direct learner feedback, determining the appropriate balance of technical rigor, scaffolding, and pacing posed a significant challenge.

Solution

Aligning personas with specific cybersecurity challenges helped the design team gauge technical complexity more accurately. For example, a persona who was experienced in machine learning but new to HPC job schedulers indicated a need for intermediate-level explanations and hands-on code snippets. Another persona operating in a bandwidth-limited setting inspired simpler, offline-compatible materials and asynchronous support. These persona-driven insights ensured that the resulting modules would offer multiple entry points, scaffolds, and resource formats, aligning with universal design for learning principles (Gordon et al., 2014).

Challenge 5: Validation and Iterative Improvement

Because personas were initially developed without direct engagement with actual learners, validating their accuracy remained a challenge. Although the team anticipated that future participant feedback would refine the personas further, immediate opportunities for validation were limited to pilot testing with volunteers who approximated some persona attributes.

Solution

To address this, the team employed preliminary pilot tests and LLM-based meta-analyses. While these measures were not substitutes for real participant data, they allowed the team to gather early signals regarding the plausibility and usefulness of the personas. Volunteer feedback confirmed the relevance of asynchronous materials, cultural responsiveness, and differentiated resource levels, lending credibility to the persona-informed design choices. The audit trail and transparent documentation ensured that the team could realign personas and training materials as soon as actual participant data became available.

Reflections on the Effectiveness of These Strategies

The team's challenges demonstrate the complexity of designing effective, adaptable HPC-enabled cybersecurity training without direct learner information. The solutions highlight the importance of methodological rigor, ethical vigilance, and responsive adaptation. Employing action research cycles and design-based research principles fostered a dynamic environment where each identified challenge became an opportunity for reflective improvement (Hinchey, 2008).

By treating personas as evolving constructs, rather than static artifacts, the T3-CIDERS project demonstrated that AI-generated personas could serve as provisional but valuable aids to instructional design. The integration of ethical review, stakeholder feedback, targeted prompt refinements, scenario-based alignment, and preliminary pilot testing formed a multi-tiered strategy for strengthening persona authenticity and utility.

In conclusion, navigating these challenges requires acknowledging the provisional nature of AI-generated personas, maintaining openness to iterative refinement, and employing diverse feedback mechanisms. This adaptive mindset ensured that personas remained aligned to produce equitable, engaging, and context-appropriate cybersecurity training modules, even in a data-scarce initial environment.

Lessons Learned

The iterative, ethically considered approach to creating AI-generated learner personas for the T3-CIDERS project yielded several key insights. First and foremost, the process underscored the value of treating personas not as static artifacts, but as evolving constructs refined through multiple lenses, ethical reviews, stakeholder input, and iterative prompt adjustments. By embracing continuous refinement, the team navigated initial shortcomings, including biased outputs and oversimplified characterizations, eventually producing more contextually authentic and pedagogically valid personas.

A second lesson emerged from the interplay between human judgment and AI capabilities. While LLMs such as GPT-4 and Claude provided rapid ideation and creative possibilities, human oversight was essential to ensure cultural responsiveness, accuracy, and relevance. This collaboration reflects a broader principle in AI-assisted instructional design: advanced technologies can inspire innovation and efficiency, but they do not replace the need for human educators and designers who bring ethical reasoning, domain expertise, and situational awareness to the design process (Holmes et al., 2022).

The persona generation process revealed three significant gaps in existing scholarship: (1) a lack of established methodologies for creating learner personas when direct audience access is unavailable, particularly in rapidly evolving technical fields, (2) insufficient frameworks for adapting empathic design processes to AI-assisted instructional design contexts, and (3) limited exploration of how multilingual educators and resource-constrained institutions integrate advanced computational tools into cybersecurity curricula. The findings particularly highlight the absence of research on educational approaches for contexts where instructors face both technical complexity and resource limitations simultaneously, a common scenario in cybersecurity education that remains underexplored in the literature on computational pedagogy.

Finally, the integration of AI into persona creation confirmed that strategic prompt engineering and ethical oversight can mitigate many common concerns associated with AI usage in education. While challenges remain, such as verifying persona authenticity without direct learner data, this experience demonstrates that ethical frameworks (Gay, 2018; Holmes et al., 2022) and culturally responsive practices can guide AI-assisted methods toward effective, learner-centered ends.

Limitations

Despite the promising results, this approach has several important limitations. First, these personas remain synthetic approximations until validated with actual participants. While Hewitt et al. (2024) demonstrated LLMs can predict human responses with high accuracy, the personas still require empirical verification. Second, the process is resource-intensive, requiring prompt engineering expertise and multiple review cycles which may not be feasible for all instructional design teams. Third, even with ethical oversight, subtle biases in LLM training data may have influenced the personas in undetected ways, as Holmes et al. (2022) caution regarding AI systems reflecting training data biases. Fourth, the approach privileges textual representations of learners, potentially overlooking dimensions of experience that emerge through interaction rather than description. Finally, while justified for T3-CIDERS given its multiplicative impact, this nine-stage process may require streamlining for contexts with tighter development timelines or fewer resources. These limitations suggest AI-generated personas complement rather than replace direct learner engagement when possible.

Future Directions

Looking ahead, the T3-CIDERS team envisions multiple avenues for deepening and extending the insights gained from this design case. The immediate next step involves rigorous validation through mixed-methods research including: (1) comparative surveys measuring alignment between persona predictions and actual participant characteristics, (2) semi-structured interviews with T3-CIDERS participants to assess perceived accuracy and usefulness of persona-based design decisions, (3) quantitative analysis of learning outcomes in courses taught by trainers, and (4) A/B testing of instructional materials designed with and without persona guidance. The validation will assess demographic matching and functional accuracy (whether personas correctly predicted learning needs, technical constraints, and pedagogical preferences). These steps will provide a crucial check on whether the personas' anticipated needs, constraints, and motivations align with those of real educators.

Beyond T3-CIDERS, the persona-based methodology has potential applications in other contexts where learner profiles are uncertain or emerging, share critical characteristics with cybersecurity that make them ideal candidates for AI-generated personas. These fields face similar challenges: rapidly evolving technical content requiring frequent curriculum updates, diverse learner populations with varying technical backgrounds, institutional resource disparities affecting implementation, and often limited access to representative learners during initial curriculum development. For example, quantum computing education must address both physics and computer science backgrounds, making diverse learner personas particularly valuable. Similarly, data science programs often serve students transitioning from diverse disciplines, creating heterogeneous classrooms which benefit from persona-driven differentiation strategies. In these contexts, AI-generated personas could provide the same structured approximation of learner diversity that proved valuable in this cybersecurity case.

Additionally, the integration of domain-specific AI models, trained on corpora relevant to HPC and cybersecurity, may yield more technically accurate and context-specific personas. Exploring how fine-tuning LLMs can enhance persona authenticity represents a promising area for future development. Another line of inquiry is to investigate the relative contributions of different feedback mechanisms, ethical reviews, stakeholder consultations, or pilot tests, to persona quality and instructional design outcomes.

In the long term, the project team hopes to leverage longitudinal studies to examine how persona-driven adjustments to training materials influence participant performance over multiple cohorts.

Understanding the lasting effects on educator self-efficacy, learner outcomes in real cybersecurity classrooms, and alignment with industry workforce needs can solidify the role of AI-generated personas as a sustainable design strategy in complex educational domains.

Conclusion

This design case has charted a path from uncertainty to insight, demonstrating how AI-generated learner personas can guide the development of HPC-enabled cybersecurity training in the absence of direct learner data. The T3-CIDERS project team transformed rudimentary, biased persona sketches into nuanced, scenario-rich profiles by synthesizing action research cycles, design-based research principles, and ethical AI considerations. These personas, in turn, shaped technically adaptable and context-responsive instructional materials, addressing varied cultural, linguistic, and infrastructural conditions that future trainers might encounter.

The significance of this approach extends beyond cybersecurity training. In an era where educational demands rapidly outpace available data on learners, persona-based methods, enhanced by AI, offer a flexible, iterative means of anticipating and accommodating diversity. While challenges persist, and further validation is necessary, the lessons gleaned here offer a blueprint for responsibly integrating AI into instructional design processes.

Ultimately, this case demonstrates that AI-generated personas are not simply automated outputs but co-constructed artifacts emerging from human-AI collaboration. For instructional designers facing similar constraints of limited learner access, compressed timelines, or diverse audience needs, the nine-stage methodology offers a practical bridge for early-stage curriculum development. By positioning LLMs as reflective sounding boards rather than autonomous creators, and anchoring their outputs in rigorous ethical frameworks and iterative stakeholder refinement, designers can create more nuanced, contextually-grounded representations than either humans or AI could produce independently. This approach helps instructional designers anticipate diverse learner needs when direct access is limited, ultimately creating more equitable, accessible learning environments that can adapt to the evolving demands of rapidly changing technical fields.

Author Note

The T3-CIDERS training program is a collaborative project among multiple institutions, funded by the U.S. National Science Foundation CyberTraining grants #2320998 and #2320999.

References

Admass, W. S., Munaye, Y. Y., & Diro, A. A. (2024). Cyber security: State of the art, challenges and future directions. *Cyber Security and Applications*, 2, 100031. <https://doi.org/10.5213/ajrsp.en.2023.45.3>

Almoughem, K. (2023). The future of cybersecurity workforce development. *Academic Journal of Research and Scientific Publishing*, 4(45), 37-48. <https://doi.org/10.5213/ajrsp.en.2023.45.3>

Atkins, D. E., Drogemeier, K. K., Feldman, S. I., Garcia-Molina, H., Klein, M. L., Messerschmitt, D. G., Messina, P., Ostriker, J. P., & Wright, M. H. (2003). Revolutionizing science and engineering through cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure. <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=841a9bb02faa96203b5058252640b260a5538e3e>

Baaki, J., & Maddrell, J. (2020). Building empathy and developing instructional design experience and skills: A case study of using personas to design open education resources. *The Journal of Applied Instructional Design*, 9(3), 77-98. <https://doi.org/10.51869/93bjm>

Branch, R. M. (2009). *Instructional design: The ADDIE approach*. Springer.

Burrell, D. N., Dattola, A., Dawson, M. E., & Nobles, C. A. (2019). A practical exploration of cybersecurity faculty development with microteaching. *International Journal of Applied Management Theory and Research*, 1(1), 32-44. <https://doi.org/10.4018/ijamtr.2019010103>

Chapman, C. N., & Milham, R. P. (2006, October). The personas' new clothes: methodological and practical arguments against a popular method. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 50, No. 5, pp. 634-636). Sage.

Childers, G., Linsky, C. L., Payne, B., Byers, J., & Baker, D. (2023). K-12 educators' self-confidence in designing and implementing cybersecurity lessons. *Computers & Education Open*, 4, 100119. <https://doi.org/10.1016/j.caeo.2022.100119>

Cooper, A. (1999). *The inmates are running the asylum*. Springer.

Gay, G. (2018). *Culturally responsive teaching: Theory, research, and practice* (3rd ed.). Teachers College Press.

Gibbons, A. S. (2013). *An architectural approach to instructional design*. Routledge.

Gordon, D., Meyer, A., & Rose, D. (2016). *Universal Design for Learning*. CAST Professional Publishing Peabody.

Hinchey, P. H. (2008). *Action research primer* (Vol. 24). Lang.

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S. B., Santos, O. C., Rodrigo, M. T., Cukurova, M., Bittencourt, I. I., & Koedinger, K. R. (2022). Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, 32(3), 504-526. <https://doi.org/10.1007/s40593-021-00239-1>

Huynh, T., Madsen, A., McKagan, S., & Sayre, E. (2021). Building personas from phenomenography: a method for user-centered design in education. *Information and Learning Sciences*, 122(11/12), 689-708. <https://doi.org/10.1108/ILS-12-2020-0256>

ISC2. (2024). 2024 ISC2 Workforce Study. ISC2. <https://edge.sitecorecloud.io/internationf173-xmc4e73-prodbc0f-9660/media/Project/ISC2/Main/Media/documents/research/2024-ISC2-WFS.pdf>

Keller, J. M. (1987). Development and use of the ARCS model of instructional design. *Journal of Instructional Development*, 10(3), 2-10. <https://doi.org/10.1007/BF02905780>

Kouprie, M., & Visser, F. S. (2009). A framework for empathy in design: Stepping into and out of the user's life. *Journal of Engineering Design*, 20(5), 437-448. <https://doi.org/10.1080/09544820902875033>

Lincoln, Y. S., & Guba, E. (1985). *Naturalistic inquiry* (Vol. 75). Sage.

Marton, F., & Booth, S. (1997). *Learning and awareness*. Lawrence Erlbaum.

McGetrick, A. (2013). Toward Effective Cybersecurity Education. In *IEEE Security & Privacy* (Vol. 11, Number 6, pp. 66-68). IEEE. <https://doi.org/10.1109/MSP.2013.155>

Mormina, M., & Pinder, S. (2018). A conceptual framework for training of trainers (ToT) interventions in global health. *Globalization and Health*, 14. <https://api.semanticscholar.org/CorpusID:53038973>

Murphy, D., Tryfona, N., & Marshall, A. M. (2023). A targeted study on the match between cybersecurity higher-education offerings and workforce needs. *Virginia Journal of Science*, 74(1-2), 1-17. <https://doi.org/10.25778/jx3e-6785>

Qawasmeh, S. A.-D., AlQahtani, A. A. S., & Khan, M. K. (2024). Navigating Cybersecurity Training: A Comprehensive Review. <https://doi.org/10.48550/ARXIV.2401.11326>

Ricci, S., Hajny, J., Piesarskas, E., Parker, S., & Janout, V. (2020). Challenges in Cyber Security Education. *International Journal of Information Security and Cybercrime*, 9(2), 7-11. <https://doi.org/10.19107/ijisc.2020.02.01>

Salminen, J., Wenyun Guan, K., Jung, S.-G., & Jansen, B. (2022). Use cases for design personas: A systematic review and new frontiers. *CHI Conference on Human Factors in Computing Systems*, 1-21. <https://doi.org/10.1145/3491102.3517589>

Smith, P. L., & Ragan, T. J. (2004). *Instructional Design*. John Wiley & Sons.

Stefaniak, J. E. (2015). The implementation of service-learning in graduate instructional design coursework. *Journal of Computing in Higher Education*, 27(1), 2-9. <https://doi.org/10.1007/s12528-015-9092-7>

Stefaniak, J. & Baaki, J. (2013). A layered approach to understanding your audience. *Performance Improvement*, 52 (6), 5-10. <https://doi.org/10.1002/pfi>

Stewart, C. A., Simms, S., Plale, B., Link, M., Hancock, D. Y., & Fox, G. C. (2010). What is cyberinfrastructure? In *SIGUCCS '10: Proceedings of the 38th Annual ACM SIGUCCS Fall Conference: Navigation and Discovery*, pp. 37-44. <https://doi.org/10.1145/1878335.1878347>

Vestergaard, L., Hauge, B., & Hansen, C. T. (2016). Almost like being there: The power of personas when designing for foreign cultures. *CoDesign*, 12(4), 257-274, <https://doi.org/10.1080/15710882.2015.1127385>

Wang, F., & Hannafin, M. J. (2005). Design-based research and technology-enhanced learning environments. *Educational Technology Research and Development*, 53(4), 5-23. <https://doi.org/10.1007/bf02504682>

Wiggins, G. P., & McTighe, J. (2005). *Understanding by design* (2nd ed.). ASCD.

Yang, M., Jiang, S., Li, B., Herman, K., Luo, T., Moots, S. C., & Lovett, N. (2025). Analysing nontraditional students' ChatGPT interaction, engagement, self-efficacy and performance: A mixed-methods approach. *British Journal of Educational Technology*. <https://doi.org/10.1111/bjet.13588>

Zagallo, P., McCourt, J., Idsardi, R., Smith, M. K., Urban-Lurain, M., Andrews, T. C., Haudek, K., Knight, J.K., Merrill, J., Nehm, R., Prevost, L.B., & Lemons, P.P. (2019). Through the eyes of faculty: Using personas as a tool for learner-centered professional development. *CBE—Life Sci. Educ.*, 18(4), ar62. <https://doi.org/10.1187/cbe.19-06-0114>

Read this online at https://jaid.edtechbooks.org/jaid_14_3/fkuiqkevdr