Social-Emotional Learning in Neurodiversity-Centered Computing Education

Dr. D. Matthew Boyer, Clemson University

Dr. Boyer is a Research Associate Professor in the Department of Engineering and Science Education and an Educational Proposal Writer in the College of Engineering, Computing and Applied Sciences.

Ren Butler, Carnegie Mellon University

I am a Ph.D. student in Human-Computer Interaction at Carnegie Mellon University.

I research psychologically safe collaboration in neurodiverse AI engineering teams. I aim to discover design principles for software development tools that support psychological safety among engineers with social and emotional differences. This can foster team learning, well-being, and productivity.

I have worked with Deloitte, Google, and the Caribbean Science Foundation.

Andrew Begel, Carnegie Mellon University

Dr. Andrew Begel, PhD is an Associate Professor in the Software and Societal Systems Department at Carnegie Mellon University. Prior to this, he spent over 16 years as a Principal Researcher at Microsoft Research. Andrew's goal is to create inclusive workplaces where all people, especially those with disabilities and the neurodivergent, can be successful, without discrimination. His recent work has helped people on the autism spectrum gain employment and facilitate social interaction, helped blind software developers collaborate with their sighted colleagues, and used affective computing and biometrics to better understand how software developers do their work.

Somayeh Asadi, University of Virginia
Taniya Mishra
Rick Kubina, The Pennsylvania State University

Social-Emotional Learning in Neurodiversity-Centered Computing Education

Abstract

In this research paper, we present findings from the second year of a design-based research project preparing autistic community college students for AI and software careers. Building on evaluation cycles that documented persistent challenges with interview anxiety, instructor variability, and the tension between technical skill development and workplace readiness, we report how the program evolved from a "teach CS + teamwork" model to an integrated framework centered on adult-focused social-emotional learning (SEL) and ethical AI. Qualitative analysis of instructor and mentor interviews revealed that generic UDL-based "communication skills" instruction was insufficient. To make the best use of their new communication skills, students required additional coaching in measurable SEL microskills: self-advocacy under stress, emotional regulation during collaboration, and timing of disclosure decisions. Moreover, instructors struggled to consistently convey communication skills training, exposing a critical gap between pedagogical goals and classroom implementation. This gap motivated the development of a scenario-based train-the-trainer model with structured modules addressing SEL coaching, talk-through facilitation, and inclusive teamwork scaffolds.

The program faced another emerging challenge: supporting the ethical use of large language models (LLMs) as learning tools without compromising academic integrity. Our data suggest that for neurodivergent learners, LLMs function as nonjudgmental rehearsal partners for practicing interview responses and debugging explanations. However, instructors wished for clear frameworks for distinguishing productive AI support from problematic dependency. This tension reframes AI literacy not merely as a technical policy issue but as an emotional and cognitive support practice requiring explicit ethical decision-making guidance.

In this paper, we contribute an SEL-integrated design narrative transferable to other computing programs serving neurodivergent learners. Our work also speaks to the structural conditions shaping access to AI careers: the program's pedagogical redesign responds directly to normalized participation norms, e.g., rapid-turn conversation, on-demand self-presentation, and unsupported disclosure decisions, which function as invisible barriers for autistic learners in computing education [1].

We argue that SEL and ethical AI scaffolds support transfer across classroom, interview, and workplace contexts by making tacit participation norms visible, teachable, and assessable - a redesign that is necessary for autistic learners and potentially beneficial across engineering and computing education [2, 3].

Keywords: *Artificial Intelligence, Autism, Computer Science Education, Social-Emotional Learning, Training*

Introduction

Autistic adults face disproportionately high unemployment despite technical aptitudes that align well with computing and AI work [4, 5]. While many autistic adults possess strengths in pattern recognition, systematic thinking, and attention to detail [6, 7], their unemployment rate exceeds 85% [8]. PAS4AI: Preparing Autistic Students for the AI Workforce (PAS4AI) addresses this disparity through community college pathways that combine technical training, mentorship, and paid internships. Our Year 2 implementation (Spring-Summer, 2025) revealed that technical proficiency alone was insufficient to secure and sustain employment. Evaluation cycles documented persistent challenges that functioned as "invisible social access barriers" [1, 9]: overwhelming interview anxiety, significant variability in how instructors applied inclusive techniques, and ongoing tensions between technical skill development and workplace readiness. Despite Year 1 efforts emphasizing CS content alongside generic teamwork modules, instructors reported that students acquired conceptual knowledge of communication strategies but struggled to apply them during high-stress activities. The challenge was not teaching more content but designing learning conditions in which students could reliably enact workplace-relevant practices under performance pressure.

This challenge creates a hierarchy of educational risks. Concerns about AI overreliance, cognitive offloading, and academic integrity are real and require explicit boundaries. Yet, for learners who are routinely filtered out before they can demonstrate technical competence, the more immediate equity problem is access to the social performances that gate employment. Our design treats LLMs and SEL scaffolds not as replacements for independent competence, but as temporary supports that make entry into high-stakes practice possible and then fade through reflection, verification, peer rehearsal, and artifact-based demonstration.

We enact this work from a critical, justice-oriented view of the classroom as a terrain of struggle and a site of possibility, where seemingly "neutral" workplace-readiness practices can operate as social access barriers for disabled and neurodivergent learners [10].

Year 1 data suggested a predictable mismatch between strategy acquisition and situated enactment. Students could name and describe communication strategies in low-stakes contexts. Yet, those strategies often collapsed in the moments that function as career gateways: mock interviews, rapid-turn teamwork, and time-pressured explanation. By "generic communication pedagogy," we mean instruction that names and describes strategies, e.g., active listening, self-advocacy scripts, conflict resolution steps, without scaffolding for the stress and context conditions in which those strategies must ultimately be enacted [11, 12]. We interpret this not as a need for "more communication content" but as a design problem: the program had not yet built reliable conditions for enactment under stress, nor provided instructors with stable coaching routines to respond to breakdowns as they occurred.

We can also understand the enactment gap as a transfer-of-learning problem. Research on writing transfer has challenged the assumption that skills learned in one setting automatically carry over to another; instead, learners must recognize similarities and differences across contexts, adapt prior strategies, and reflect on how knowledge should be transformed for new rhetorical and social situations [2, 3]. Our use of repeated rehearsal, reflection, and scaffold fading follows this logic: the goal is not simply for students to possess communication scripts, but for them to adapt self-advocacy and explanation practices across classrooms, mock interviews, internships, and workplace interactions.

The Year 2 Pivot

In response, the program shifted to an integrated framework centered on adult-focused SEL and ethical AI literacy. Qualitative analysis of instructor and mentor interviews identified coachable, measurable SEL microskills that served as prerequisites for enacting communication strategies under stress: self-advocacy, emotional regulation during collaboration, and timing of disclosure decisions. By operationalizing abstract SEL competencies [13, 14] into discrete microskills, the program sought to provide foundations for both technical and teamwork success.

Simultaneously, instructors struggled to consistently convey communication skills training, exposing a gap between pedagogical goals and classroom implementation. This gap motivated the development of a scenario-based train-the-trainer model [15, 16, 17, 18] with structured modules addressing SEL coaching, talk-through facilitation, and inclusive teamwork scaffolding. A second design tension emerged: supporting students' use of large language models (LLMs) as learning tools while preventing academic integrity violations and avoiding unproductive dependence. Instructors observed that LLMs could function as nonjudgmental rehearsal partners, offering psychological safety during rehearsal [19, 20, 21, 22]. Yet instructors needed clearer boundary frameworks distinguishing productive support from problematic dependency, repositioning AI literacy as an ethical decision-making practice intertwined with cognitive and emotional support.

Research Questions and Contributions

Guided by instructor and mentor accounts from the Year 2 evaluation cycle, we address three research questions: (1) What SEL microskills did instructors identify as necessary for learners to enact communication strategies under stress? (2) How did instructors interpret the role of LLMs in learner rehearsal, and what ethical boundary frameworks did they need? (3) What instructor supports were needed to consistently implement SEL coaching and inclusive teamwork?

This work makes three contributions. First, we contribute an SEL-integrated design narrative that reframes workplace readiness as a context-of-enactment problem. This reframing also makes visible the structural conditions that create enactment failures: neurotypical participation norms embedded in technical hiring and collaborative computing culture that systematically disadvantage autistic learners even when they possess the requisite technical competencies. Second, we present a train-the-trainer approach designed to reduce instructor variability through scenario-based preparation. Third, we reframe ethical AI literacy for neurodivergent learners as an ethical decision-making practice closely tied to emotional and cognitive support.

Although we designed the intervention for autistic learners in AI and software career pathways, we do not intend for the implications to be limited to a single population or technical subfield. Engineering and computing programs often rely on tacit norms of rapid verbal response, confident self-presentation, and unsupported teamwork as proxies for professional readiness. Making these norms explicit, teachable, and assessable may be necessary for autistic learners. Still, it may also benefit students without disclosed disabilities, including multilingual students, first-generation students, students with anxiety, and students still learning the hidden curriculum of engineering professionalism.

Background and Conceptual Framework

Neurodiversity-Centered Computing Education

This work adopts a neurodiversity-affirming perspective, reframing cognitive differences such as autism, ADHD, and dyslexia as valuable human variations rather than deficits [23, 24]. In computing contexts, neurodivergent individuals often demonstrate exceptional capabilities in pattern recognition, systematic thinking, and precise attention to detail [6, 7]. However, traditional academic and workplace structures frequently obscure these strengths through rigid pedagogical approaches and standardized participation norms [25, 26].

Many instructional defaults assume uniform participation: fast conversational turn-taking, rapid inference under time constraints, and socially fluent self-presentation as a proxy for competence. When treated neutrally, these norms become invisible barriers that suppress the enactment of competence among learners whose strengths are expressed differently [1, 9]. By moving beyond simple accommodation toward a neurodiversity-affirming lens, we can reimagine collaboration as a social structure that supports diverse interaction styles [1, 27]. The classroom, in this frame, becomes a contested site where inherited participation norms can be named, questioned, and redesigned, as a space of both struggle and genuine possibility for learners whose capacities have been obscured by those norms.

Adult-Focused SEL

PAS4AI's redesign treated SEL not as a separate "life skills" strand but as an enabling condition for technical and collaborative success. Social-Emotional Learning encompasses the development of self-awareness, emotional regulation, and interpersonal skills [13, 14]. While often associated with K-12 education [28], SEL is a critical prerequisite for success in adult collaborative environments where team dynamics significantly influence project outcomes [27, 29]. For autistic learners, the challenge is often not a lack of social knowledge but difficulty enacting that knowledge under duress. SEL can remain too abstract to guide instruction unless operationalized into teachable, coachable units that provide systematic practice and support.

Transfer, Reflection, and Situated Enactment

The transfer scholarship in rhetoric and writing studies provides a useful vocabulary for the gap we observed between knowing a strategy and using it effectively in high-stakes contexts. Rather than treating transfer as the direct movement of a stable skill from one context to another, this literature emphasizes adaptation, metacognition, reflection, and recognition of contextual demands [2, 3]. This perspective helps explain why students could describe communication strategies in low-stakes settings but struggle to enact them during interviews, collaborative debugging, or disclosure decisions. In our design, reflection prompts, LLM-supported rehearsal, peer practice, instructor-facilitated scenarios, and portfolio talk-throughs function as transfer supports: they help learners abstract from one performance context, adapt to another, and gradually assume more independent control.

AI Literacy as Cognitive and Emotional Support

AI literacy frameworks have expanded rapidly in the generative AI era, moving well beyond early emphases on policy awareness and algorithmic fairness [30, 31] toward integrated competencies for understanding, evaluating, and using AI systems in context, including functional, ethical, rhetorical, and pedagogical literacies [32]. For neurodivergent learners, LLMs offer unique affordances as cognitive and emotional supports. This rehearsal partner model builds self-efficacy and aligns with strengths in solitary deep work while building collaborative readiness. Yet this framing also surfaces an unresolved boundary problem: distinguishing productive support from problematic dependency and preventing academic integrity violations without foreclosing beneficial uses. Consequently, we reframe AI literacy as an ethical decision-making practice requiring explicit instruction in prompt engineering and critical evaluation.

Program Context and Methods

Setting and Participants

PAS4AI is an NSF-funded community college-based bridge program preparing autistic students for entry-level AI and software careers, integrating technical training with structured internship experiences. During Spring-Summer 2025, we served nine autistic community college students recruited nationwide. Eligibility required completing two introductory programming courses and one mathematics course. The curriculum spanned 18–22 weeks: a 10-week virtual training phase (20 hours per week) followed by an 8-12-week paid internship. Technical instruction used Python, TensorFlow, Keras, and Jupyter Notebooks, emphasizing responsible AI design. The focal participants for this analysis are PAS4AI project team members interviewed during Year 2 data collection: five instructors/mentors (primarily PhD students and a postdoc) and two job coaches specializing in neurodivergent workplace support.

Design-Based Research Structure

We position PAS4AI as design-based research, in which successive implementations surface design tensions that drive targeted redesign [33, 34]. Year 1 focused on CS content and generic teamwork training; formative evaluation revealed that technical proficiency did not reliably translate to workplace readiness. Year 2 executed a strategic pivot to an integrated SEL and ethical AI framework. The Year 3 implementation in Spring 2026 is an unfunded sustainability extension, testing train-the-trainer modules and portfolio-based assessment with alumni peer mentorship and OER dissemination. We organized the Year 2 redesign around a central conjecture: If we operationalize SEL into coachable microskills and prepare instructors through scenario-based training, we will reduce implementation variability and improve learner readiness to enact workplace-relevant practices under performance stress.

We understand these three cycles as a cumulative design arc. Year 1 (Summer 2024) implemented a foundational “CS + teamwork” model, establishing the technical core and introducing generic communication training; formative evaluation documented the enactment gap described above: students who understood the strategies conceptually could not reliably deploy them under pressure. Year 2 (Spring-Summer 2025) addressed that gap by operationalizing SEL as coachable microskills and reframing ethical AI literacy as a boundary-making practice that requires explicit instruction from the first week of the program. Year 3 (Spring 2026) prioritizes train-the-trainer implementation, portfolio-based assessment, and alumni peer mentorship to determine whether we can maintain program quality beyond the original grant cycle. Each year's redesign was informed directly by qualitative evaluation data from the prior cycle, consistent with the iterative logic of design-based research [35].

Data Sources and Analysis

The primary dataset consists of semi-structured interviews (45–60 minutes each) with instructors and mentors collected through Year 2 external evaluation, supplemented by curriculum artifacts, mentor guides, and train-the-trainer materials. We use thematic analysis [36] in an iterative process: open coding to generate initial themes, followed by axial coding to explore relationships among themes and their connections to the research questions. As we used the interviews for external evaluation, the data received special consideration due to their dual use in both reporting and this research analysis. Separate thematic analysis from emergent and conceptual framings provided complementary themes, while this paper provides a more detailed analysis of the evidence-based themes and suggested revisions for Year 3.

Findings

Arc 1: From Communication Skills to SEL Microskills

Why Generic Communication Instruction Failed

While Universal Design for Learning informed the initial curriculum [25, 37], Year 2 data demonstrated that UDL-based communication skills instruction did not reliably translate into performance in high-stress contexts. Across accounts, a central pattern emerged: students understood communication strategies conceptually but struggled to enact them under pressure, uncertainty, and heightened social demands. Technical proficiency did not mitigate invisible social access barriers, particularly under performance pressure. Instructors noted that students grasped the "what" of communication but struggled with the "how" and "when" in high-stakes environments. One instructor described, "They know what to do but cannot do it in the moment," highlighting frustration when skills failed to transfer amid anxiety. Interview anxiety frequently caused capable learners to freeze despite technical mastery. Another instructor observed breakdowns during collaboration: "They were having issues with their team, and it was they needed somebody else to kind of like moderate." One project member noted that students were performing activities only because "they were told to," rather than developing the autonomy to apply their skills in real-world contexts.

Three Critical SEL Microskills

In response to the enactment gap, instructors and mentors identified granular SEL microskills that functioned as prerequisites for successful enactment (see Table 1). These were positioned not as peripheral "soft skills" but as enabling conditions for technical and collaborative success.

We interpret these scenarios as norm-breakdown pairings: moments when dominant expectations for synchronous participation, rapid verbal response, self-management, or unassisted professionalism made student competence less visible.

Domain	Breakdown	Observable Indicator	Coaching Response
Self-advocacy	Interview freeze	Student goes silent on unexpected question	LLM rehearsal partners; pre-planned clarification phrases
Self-advocacy	Overwhelm spiral	Student struggles silently rather than seeking help	*“Reel back” protocol: 5-minute instructor-initiated check-ins*
Self-advocacy	Anonymous routing	*“Some of them private message me...‘I want to ask this question anonymously’”*	Graduated disclosure practice; instructor models vulnerability
Emotional regulation	Camera shutdown	*“Their default was basically to clam up, turn off the camera, turn off their audio”*	Non-verbal Miro boards; text contribution as legitimate participation
Emotional regulation	Team conflict	*“They needed somebody else to kind of like moderate”*	Instructor-facilitated resolution; structured turn-taking
Emotional regulation	Processing mismatch	Student falls behind in rapid turn-taking; cannot formulate responses quickly	Extended wait time; asynchronous ideation before synchronous discussion
Emotional regulation	Defensive escalation	Student responds to code critique argumentatively or shuts down	Cooling-off period; written response before verbal
Disclosure	Premature disclosure	Student shares diagnosis without strategic context consideration	Disclosure decision framework; risk-benefit mapping
Disclosure	Missed accommodation	Student struggles during internship rather than requesting support	Early accommodation planning; instructor-initiated offers
Assessment	Knowledge-performance gap	*“They know what to do but cannot do it in the moment”*	Artifact talk-throughs replacing timed assessments
Ethical AI	Unreviewed acceptance	Students copying LLM outputs without verification; concealing AI use	Transparent use of logging; verification exercises
Ethical AI	Concealed dependency	Using LLMs to mask understanding gaps rather than address them	Reflection requirements: coaching on ethical boundaries

Table 1: Documented Breakdown Scenarios. Italicized entries are direct quotes from instructor interviews.

Self-advocacy under stress involves self-awareness and honesty about not understanding a concept or task. Observable indicators include requesting clarification calmly and proactively seeking support when overwhelmed. When this skill was missing, students froze during mock interviews rather than asking for clarification. Deliberate low-stakes practice with rehearsal partners helped. One instructor described a student who felt overwhelmed by a coding requirement and proactively sought a five-minute conversation to “reel back” the issue.

Interpreted structurally, these self-advocacy breakdowns show how engineering and computing settings often treat immediate clarification, public help-seeking, and confident verbal repair as neutral professionalism, even when those expectations can make access contingent on stress performance rather than technical understanding.

Emotional regulation during collaboration involves managing social dynamics and sensory overwhelm inherent in teamwork. Observable indicators include pausing before responding to conflict and maintaining engagement during disagreements. When missing, team members defaulted to shutting down, turning off cameras, and withdrawing from discussion. Instructors used nonverbal collaboration boards on Miro.com as scaffolding to help students regulate their thinking during team ideation without mandating a specific pace for interactive conversation. One instructor noted: “Their default was basically to clam up, turn off the camera, turn off their audio, and type in text only when spoken to.”

These emotional-regulation breakdowns similarly reflected a mismatch between student processing needs and dominant classroom norms that treat synchronous verbal participation as the primary evidence of engagement. The redesign legitimized text, visual boards, wait time, and asynchronous ideation as valid forms of participation rather than secondary accommodations.

Timing of disclosure decisions involves navigating the complex risks and benefits of self-disclosing autism status. Observable indicators include weighing risks before sharing a diagnosis and strategically requesting accommodations when appropriate. One instructor shared: “Some of them private message me and said, ‘You know, I want to ask this question anonymously...if I have such and such... pre-existing condition on top of being autistic-like, how do I deal with that in a workplace situation?’”

The disclosure examples also show that we cannot reduce workplace readiness to individual confidence. Students were being asked to make consequential disclosure decisions in environments where the risks, benefits, and norms for requesting support were often tacit rather than instructionally surfaced.

Method	**Description**	**Implementation**
“Reel back” protocol	5-minute instructor-initiated check-ins when students show overwhelm signs	Instructor notices behavioral indicators -> normalizing frame -> focusing question -> concrete next step
Non-verbal scaffolds	Miro boards for team ideation without verbal speed pressure	Asynchronous individual contribution -> synchronous structured synthesis
Graduated rehearsal	Progressive exposure from low to high social stakes	LLM partners (low threat) -> peer practice -> instructor-facilitated scenarios
Clarification phrases	Pre-planned language for requesting help under pressure	*“I want to make sure I understand - could you say more about [X]?”*

Method	Description	Implementation
Structured feedback	Written response protocols before verbal discussion	Student writes reaction -> cooling period -> verbal exchange

Table 2: SEL Coaching Methods Implemented

The SEL lead conducted scenario-based training sessions to prepare instructors on these techniques (see Table 2). However, the evaluation found instructors “lacked sufficient autonomy” to implement strategies without explicit direction; they executed activities “because we told them to” rather than from internalized understanding. One instructor adapted flexibly to student needs; another maintained “rigid expectations” and showed “less willingness to modify teaching tactics.” These fidelity challenges informed Year 3 train-the-trainer recommendations emphasizing sustained mentorship over initial workshops alone.

In this framing, instructor preparation is not only a fidelity mechanism. It is also a way of transforming the educational community by building shared interpretations of stress indicators, shared coaching routines, and shared standards for what counts as legitimate participation in technical learning spaces.

The Instructor Variability Challenge

A second finding concerned implementation consistency. Instructors reported difficulty conveying SEL and communication content in a stable, shared way, exposing a gap between pedagogical goals and classroom enactment. Qualitative data revealed competing teaching styles: one instructor preferred a nonlinear approach allowing for high-level ideation, while another maintained rigid expectations and was less willing to adapt to individual student capabilities. This rigidity often created conflict within the project’s hierarchical organizational structure. Instructors felt they lacked agency to implement SEL interventions independently, leading to inconsistent student experiences.

This variability motivated the development of scenario-based train-the-trainer modules to equip instructors with the pedagogical knowledge and autonomy to consistently coach SEL microskills. The proposed model centers structured modules focused on: (1) priming the instructional team on common student challenges through simulation scenarios, (2) developing facilitation protocols for “artifact talk-throughs” to replace high-stakes interviews, and (3) establishing consistent, inclusive teamwork scaffolds to prevent instructors from reverting to less flexible teaching habits when stressed.

A pattern in training drew directly from the Year 2 evaluation: when a student turned off their camera and went silent during a collaborative debugging session, one instructor paused, opened a brief private chat, and used the “reel back” protocol of naming the overwhelm, narrowing focus to one next step, and returning to the group with a defined task. A second instructor, encountering the same behavioral indicator, maintained the original pace and treated the silence as disengagement rather than a regulatory need. Simulating this contrast in training created a concrete occasion for the team to build shared language about what student stress indicators look like, what a coaching response requires, and how widely instructor interpretations can diverge without explicit shared frameworks [35].

This training example reframes instructor preparation as a collective interpretive practice: instructors learn not only what to do, but how to read classroom breakdowns as signals that existing norms may need redesign.

Arc 2: Ethical AI Literacy as SEL-Adjacent Practice

LLMs as Nonjudgmental Rehearsal Partners

A parallel set of findings concerned students' use of LLMs as learning tools. For neurodivergent learners, LLMs served dual purposes as both technical and social-emotional scaffolds. Instructors described that LLMs could serve as nonjudgmental rehearsal partners for practicing interview responses and debugging explanations. We consider this role instructionally meaningful because it expands access to iterative practice, reduces social threat during rehearsal, and supports preparation for high-stakes contexts.

However, LLM rehearsal's perceived safety is neither universal nor risk-free, and we treat it as a design hypothesis that requires boundaries rather than a settled benefit. Public and scholarly work has documented failure modes relevant to student well-being, including harmful outputs in mental-health-adjacent interactions and the risk of misplaced trust in AI guidance [38, 39, 40]. The program positions LLMs as one scaffold within a graduated rehearsal sequence: AI partners at low social threat, followed by peer practice and then instructor-facilitated scenarios, so that LLMs are intentionally not the endpoint of rehearsal.

A second concern is LLM sycophancy: the tendency to mirror or endorse a user's beliefs and framing rather than prioritize calibration and truth [41, 42]. For students practicing interview rehearsal and self-advocacy coaching, overly agreeable feedback can mask gaps, reinforce unhelpful narratives, or produce confident-sounding but miscalibrated responses. Crucially, sycophancy poses a specific risk for learners who are simultaneously building internal self-advocacy capacity: an always-affirming interlocutor may build rehearsal confidence without adequately preparing students for the reality of human feedback, which carries genuine stakes, genuine disagreement, and genuine uncertainty. The program treats "validation" from LLMs as something to be checked against external criteria, e.g., rubrics, exemplars, and instructor feedback, rather than accepted as truth.

Third, we recognize the risk of cognitive offloading or "cognitive debt" when learners delegate composing and reasoning to LLMs in ways that reduce later independent enactment [43]. For students learning to trust and act on their own judgment under pressure, this risk is particularly acute. Accordingly, the program treats LLM use as scaffolded rehearsal with intentional fading supports, such as reflection requirements, verification exercises, and periodic no-LLM demonstration tasks, rather than as a default productivity layer.

The Support Versus Dependency Tension

While LLMs offered safety, they also introduced a design tension between productive support and problematic dependency. Instructors reported needing clearer frameworks to distinguish productive AI-supported learning from problematic dependency and to address academic integrity risks without foreclosing beneficial uses. Students occasionally used LLMs uncritically, copying outputs or concealing their AI use rather than using the tools for learning. They did not frame tension management as a simple “allow versus prohibit” decision; rather, they framed it as a pedagogical need for consistent guidance on appropriate forms of AI support.

The resulting design response involves integrating ethical AI pedagogy into the first week of instruction, developing a simple decision rubric (rehearsal versus substitution, transparent versus concealed use, reflection requirements), explicit instruction in prompt engineering, and coaching students on “ethical moves,” treating decisions about AI use similarly to SEL microskills requiring self-awareness and self-advocacy.

Dimension	**Productive Support**	**Problematic Dependency**
Purpose	Rehearsal: practicing for performance student will own	Substitution: producing work to present as one's own
Transparency	Willing to acknowledge: “I practiced this with ChatGPT”	Concealing AI involvement from instructors/peers
Reflection	Can explain in own words; can defend approach if questioned	Adopts output without critical evaluation
Self-check	*“Am I using this to practice?”*	*“Am I using this to avoid learning?”*

Table 3: Ethical AI Use Framework

Because simple policy statements are insufficient for navigating the gray areas that arise in practice, evaluation suggests using explicit coaching language for three representative ambiguous cases. For example: Case 1: A student uses an LLM to generate an outline for their explanation, then writes the explanation in their own words, which is productive if the student can defend their reasoning and disclose the tool use, but warrants a coaching conversation if the outline is substituted for genuine planning. Case 2: A student practices their interview response with an LLM three times, building fluency, but then reads the LLM’s suggested phrasing verbatim during the mock interview, shifting from rehearsal to substitution at the moment of delivery. Case 3: A student asks an LLM to “make their code explanation more professional” without understanding what changed or why, which introduces a transparency and reflection deficit even without intent to deceive. Each case is designed as a coaching occasion rather than a disciplinary event, consistent with the program’s SEL-integrated approach.

Instructors observed LLMs functioning as valuable, nonjudgmental rehearsal partners, while also noting students “copying outputs uncritically” and concealing their use. This framework was emerging but incomplete by the end of Year 2. The evaluation recommended integrating ethical LLM pedagogy from Week 1 to establish norms before problematic patterns form. Planned Year 3 components include transparent use logging (framed as professional practice), reflection requirements after AI interactions, and coaching protocols for educationally responding to misuse.

Discussion and Implications

Workplace Readiness as a Stress-Context Performance Problem

Throughout this paper, we use "performance" in the sense developed by Schechner [44]: as social action taken within a specific context, shaped by the expectations of audiences and the demands of the situation. Using this construct distinguishes performance-as-enactment, i.e., what learners do in a technical interview or collaborative debugging session, from assessed proficiency or technical mastery, for which we use terms such as "demonstrated competency" or "assessed skill." The distinction matters because, in the Schechnerian sense, performance is structurally shaped: the context's conditions either enable or suppress what a learner can enact, regardless of their underlying knowledge or capability.

A central implication is that workplace readiness in neurodiversity-centered computing education is partly a stress-context performance problem: learners who can demonstrate capability in low-stakes settings may struggle to enact communication and collaboration practices in high-pressure moments that serve as career gateways. Traditional computing education often assumes that once students learn concepts or strategies, they will naturally transfer that knowledge. However, invisible social access barriers, such as rapid conversational turn-taking in pair programming or on-demand technical interviews, function as social performance chokepoints.

The findings suggest that communication skills are most instructionally useful when treated as context-dependent practices with coachable prerequisites. The design contribution is that SEL matters and that it becomes teachable in technical learning environments when operationalized as discrete behaviors that instructors can notice, prompt, and reinforce during authentic tasks. This responsibility makes SEL relevant to computing education as an enabling layer that directly shapes learners' ability to demonstrate competence.

For engineering education more broadly, this finding suggests that technical knowledge and social-emotional enactment should not be treated as competing curricular goals. Professional engineering education already requires students to explain their work, collaborate under uncertainty, request clarification, respond to critique, and make ethical decisions about tools. Making these practices explicit and assessable can support autistic learners while also strengthening professional preparation for students whose barriers are less visible or not formally disclosed.

This framing clarifies a key equity issue: common performance rituals such as technical interviews and rapid-turn collaborative debugging often function as invisible access barriers when they implicitly privilege particular interaction styles. The program's emerging response, i.e., portfolio-centered evidence and artifact talk-throughs, shifts evaluation away from high-stakes social performance toward structured demonstrations of technical work that can also elicit advocacy and communication practices. These assessment redesigns are not merely accommodations; they are structural interventions that challenge the assumption that on-demand self-presentation under time pressure is a valid proxy for technical competency.

This redesign changes the program's assessment ecology. Rather than helping students adapt to a single dominant performance ritual, portfolio talk-throughs redistribute what counts as evidence: completed artifacts, explanation, reflection, and self-advocacy become assessable together. The assessment intervenes in the norm that employability must be demonstrated through rapid, high-stakes verbal performance.

Instructor Preparation and Program Sustainability

Neurodiversity-centered interventions are unlikely to scale or sustain without explicit attention to instructor variability. The Year 2 pattern is not simply that instructors differed, but that pedagogical inconsistency can become a primary mechanism through which inclusive intent fails to translate into inclusive practice. Generic professional development cannot ensure inclusive goals translate into stable classroom practices; when the core instructional target includes SEL coaching and ethical AI boundary-setting, the consequences of inconsistency intensify.

The train-the-trainer response is therefore best understood as a sustainability mechanism. By standardizing core SEL microskills and ethical AI boundary protocols, programs can ensure students receive consistent coaching regardless of their specific instructor. The standardized core is the instructional theory of action: SEL microskills must be coached explicitly in contexts where enactment typically fails, and ethical AI literacy must be taught as boundary-making practice rather than left to ad hoc policy reminders. What can remain adaptable are specific scenarios used in training, delivery modalities, and the pacing of coaching routines.

Sustainability also depends on resource dependencies affecting long-term viability, including job coaches and specialized supports. PAS4AI's experience suggests that enabling conditions for enactment include not only pedagogy but sustained human supports that stabilize learners' experience across high-pressure moments.

Portfolio-Based Assessment

Stakeholders recommended replacing high-stakes interviews with portfolio-based assessment aligned with neurodivergent strengths. Students would build projects throughout the program, then deliver brief artifact talk-throughs demonstrating technical acumen through completed work: "this is the project that I'm proud of, this is why I'm proud of it... but it also shows the technical acumen that they have." This practice shifts assessment from on-demand problem-solving (where anxiety can cause capable students to "freeze up") to the structured demonstration of artifacts already produced, providing evidence of both technical capability and communication skills while reducing invisible social performance barriers.

Transferability

Three design elements appear most portable to other contexts. First, the SEL microskills framing and coaching approach provides a practical unit of design that programs can integrate into collaboration pedagogy and professional preparation. Second, the scenario-based instructor professional development structure targets enactment by preparing instructors to respond to predictable breakdowns with consistent coaching moves. Third, ethical AI boundary pedagogy treats AI literacy as a decision-making practice with explicit boundaries that can be taught from the beginning of instruction.

The likely transferability of these elements rests on their targeting of classroom and assessment structures rather than diagnostic categories alone. Programs outside autism-focused AI education can use the same design logic to ask where their own norms of speed, verbal fluency, unsupported teamwork, and high-stakes self-presentation may be suppressing otherwise valid demonstrations of competence.

Conversely, several features remain context-bound: specific industry internship partnerships, local support structures contingent on funding, and cohort composition. Recognizing these limits allows program developers to focus efforts on the transferable pedagogical core while navigating local resource constraints.

Resources Developed

SEL Microskills Map: Provides shared language for instructors with operational definitions for each microskill, observable indicators during authentic tasks, and coaching prompts for actionable facilitation moves.

Train-the-Trainer Module Overview: Includes scenario descriptions, facilitator decision points, and recommended coaching moves aligned to common failure modes, supporting consistent enactment of SEL coaching techniques and talk-through facilitation protocols.

*Ethical LLM Use Framework: I*ncludes a learner-facing decision guide that distinguishes rehearsal-oriented support from substitution-oriented dependency, an instructor-facing coaching protocol, and explicit integration with academic integrity expectations, designed for a Week 1 introduction.

Preliminary Student Outcomes

Formative data suggest students were highly receptive to explicit SEL instruction. Instructors observed that when they adopted facilitator roles aligned with SEL principles, students felt more comfortable sharing their challenges. Microskills coaching was associated with reduced withdrawal during team sessions and increased self-reported confidence during rehearsal activities. Artifact talk-throughs surfaced technical strengths that prior high-stakes interview formats had suppressed. Anecdotally, one student received a job offer almost immediately after training. We plan to track employment outcomes longitudinally for the Year 3 extension.

External evaluation of Year 2 corroborated these instructor-reported patterns [35]. The evaluation found that stakeholders described stronger alignment between instructional intent and student outcomes when SEL and self-efficacy were foregrounded as prerequisites for technical success, rather than treated as supplementary. The evaluation also confirmed that interview anxiety persisted as the primary career-gateway barrier, with otherwise technically capable students routinely failing to demonstrate their competencies under the on-demand, time-pressured conditions of conventional technical interviews. This convergence between instructor accounts and external evaluation strengthens our confidence in the intervention's formative effects, while acknowledging that the program's small cohort size and single-site context limit the conclusions that can be drawn from these early indicators.

Limitations and Future Work

Several limitations constrain interpretation. The small cohort size (n=9) limits the generalizability of the results. The single-program context means that findings are shaped by PAS4AI's specific staffing model and partnerships. Reliance on the instructor and mentor self-reports introduces perspective bias. The iterative, design-based nature of the intervention meant it was constantly evolving, complicating controlled comparisons. Limited longitudinal evidence on career outcomes means claims about workforce impact remain provisional.

The Spring 2026 unfunded extension will implement train-the-trainer modules and portfolio-centered assessment while testing sustainability mechanisms: alumni peer mentorship, OER dissemination via CMU repositories, and the formal integration of portfolio-based assessment to replace high-stress interviews. Additional evidence will include implementation fidelity indicators, learner outcomes linked to microskill enactment, documentation of ethical AI decision-making routines, and six-month employment-tracking data.

Conclusion

This paper documents how Year 2 implementation tensions in PAS4AI drove a shift from generic communication instruction toward enactment-oriented design. We conclude that workplace readiness for neurodivergent learners is not solely a matter of skill acquisition but a stress-context performance problem that requires redesigned learning conditions.

The Year 2 redesign does more than help autistic students adapt to existing computing and engineering norms by identifying how those norms produce breakdowns, then redesigns instruction, assessment, and instructor preparation so that they can demonstrate through multiple, structured, and ethically supported forms of participation.

Three core findings advance neurodiversity-centered computing education. First, SEL becomes teachable in technical environments only when operationalized into discrete, coachable microskills that instructors can notice, prompt, and reinforce during authentic tasks. Second, scenario-based instructor preparation is essential for maintaining implementation fidelity and reducing pedagogical inconsistency. Third, ethical AI literacy functions as an emotional and cognitive support practice that provides a space for scaffolded rehearsal while requiring explicit ethical boundaries, including attention to the sycophantic tendencies of current AI systems and the cognitive risks of unscaffolded reliance on LLMs.

Creating equitable pathways into AI work requires more than access to technical content; it requires building learning conditions in which learners can reliably demonstrate competence, advocate for needs, and navigate high-stakes social performance contexts without being filtered out by misaligned participation norms. Autistic learners face barriers in computing education, including the structure of computing classrooms, hiring processes, and collaborative workplaces, not in themselves. Redesigning those structures, even incrementally at the program level, is a pedagogical act and an equity intervention. Through these design refinements, we aim to empower neurodivergent learners to achieve career sovereignty in an increasingly AI-driven world.

Because these changes make tacit professional expectations explicit, they may also benefit learners beyond the focal autistic population. The strongest claim supported by this study is not that one intervention works for all students, but that equity-oriented pedagogy developed for a group facing acute access barriers can reveal design flaws in ordinary engineering and computing education.

Acknowledgments

The U.S. National Science Foundation supports this work under Award No. 2322555. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. National Science Foundation.

References

[1] D. E. M. Milton, "On the ontological status of autism: The 'double empathy problem'," Disability & Society, vol. 27, no. 6, pp. 883–887, 2012. doi: 10.1080/09687599.2012.710008

[2] R. L. Snead, K. P. Alexander, M. Davis, L. W. Mina, and R. P. Shepherd, Rethinking Transfer: The Rise of Writing Transfer Research Across Contexts, WPA-CompPile Research Bibliographies, no. 30. Fort Collins, CO: The WAC Clearinghouse, 2022. [Online]. Available: https://wacclearinghouse.org/docs/comppile/wpa/rethinking_transfer.pdf

[3] C. M. Anson and J. L. Moore, Eds., Critical Transitions: Writing and the Question of Transfer. Fort Collins, CO: The WAC Clearinghouse; Boulder, CO: University Press of Colorado, 2016. doi: 10.37514/PER-B.2016.0797.

[4] D. Hendricks, "Employment and adults with autism spectrum disorders: Challenges and strategies for success," Journal of Vocational Rehabilitation, vol. 32, no. 2, pp. 125–134, 2010. doi: 10.3233/JVR-2010-0502

[5] P. T. Shattuck et al., "Postsecondary education and employment among youth with an autism spectrum disorder," Pediatrics, vol. 129, no. 6, pp. 1042–1049, 2012. doi: 10.1542/peds.2011-2864

[6] S. Baron-Cohen, Autism and the Technical Mind. New York, NY: W.W. Norton & Company, 2009.

[7] L. Mottron, M. Dawson, I. Soulières, B. Hubert, and J. Burack, "Enhanced perceptual functioning in autism: An update, and eight principles of autistic perception," Journal of Autism and Developmental Disorders, vol. 36, no. 1, pp. 27–43, 2006. doi: 10.1007/s10803-005-0040-7

[8] A. J. Griffiths, C. M. Giannantonio, A. E. Hurley-Hanson, and D. N. Cardinal, "Autism in the workplace: Assessing the transition needs of young adults with autism spectrum disorder," Journal of Vocational Rehabilitation, vol. 52, no. 3, pp. 291–304, 2020. doi: 10.3233/JVR-201080

[9] B. Pfeiffer, A. Piller, A. Giagnacova, and C. Scali, "Social participation and autism: Invisible barriers to college success," Journal of Occupational Therapy, Schools, & Early Intervention, vol. 11, no. 3, pp. 290–305, 2018. doi: 10.1080/19411243.2018.1455552

[10] S. R. Harper, "An anti-deficit achievement framework for research on students of color in STEM," New Directions for Institutional Research, vol. 2010, no. 148, pp. 63–74, 2010. doi: 10.1002/ir.362

[11] D. P. Dannels, "Communication across the curriculum and in the disciplines: Speaking and writing in the disciplines," Communication Education, vol. 51, no. 1, pp. 148–158, 2001.

[12] D. P. Dannels and K. N. Martin, "Critique as culturally situated design education," Journal of Engineering Education, vol. 97, no. 3, pp. 339–351, 2008.

[13] CASEL, "CASEL's SEL Framework," Collaborative for Academic, Social, and Emotional Learning, 2020. [Online]. Available: https://casel.org/fundamentals-of-sel/what-is-the-casel-framework/

[14] C. E. Domitrovich, J. A. Durlak, K. C. Staley, and R. P. Weissberg, "Social-emotional competence: An essential factor for promoting positive adjustment and reducing risk in school children," Child Development, vol. 88, no. 2, pp. 408–416, 2017. doi: 10.1111/cdev.12739

[15] M. S. Garet, A. C. Porter, L. Desimone, B. F. Birman, and K. S. Yoon, "What makes professional development effective? Results from a national sample of teachers," American Educational Research Journal, vol. 38, no. 4, pp. 915–945, 2001. doi: 10.3102/00028312038004915

[16] L. M. Desimone, "Improving impact studies of teachers' professional development: Toward better conceptualizations and measures," Educational Researcher, vol. 38, no. 3, pp. 181–199, 2009. doi: 10.3102/0013189X08331140

[17] C. Carroll, M. Patterson, S. Wood, A. Booth, J. Rick, and S. Balain, "A conceptual framework for implementation fidelity," Implementation Science, vol. 2, p. 40, 2007. doi: 10.1186/1748-5908-2-40

[18] J. A. Durlak and E. P. DuPre, "Implementation matters: A review of research on the influence of implementation on program outcomes and the factors affecting implementation," American Journal of Community Psychology, vol. 41, no. 3–4, pp. 327–350, 2008. doi: 10.1007/s10464-008-9165-0

[19] J. Jang, S. Moharana, P. Carrington, and A. Begel, "'It's the only thing I can trust': Envisioning Large Language Model Use by Autistic Workers for Communication Assistance," in Proc. CHI '24, New York, NY: ACM, 2024, Art. 77. doi: 10.1145/3613904.3642894

[20] A. Pinto et al., "'A little bit of a life raft' – Exploring the Use and Experiences of ChatGPT as a Support Tool among Adults with ADHD," in Proc. OzCHI '25, New York, NY: ACM, 2025, pp. 50–67. doi: 10.1145/3764687.3764713

[21] H. Jamshed, M. Naseem, V. Potluri, and R. N. Brewer, "Rethinking Productivity with GenAI: A Neurodivergent Students' Perspective," in Proc. ASSETS '25, New York, NY: ACM, 2025, Art. 37. doi: 10.1145/3663547.3746329

[22] K. Glazko et al., "Autoethnographic Insights from Neurodivergent GAI 'Power Users'," in Proc. CHI '25, New York, NY: ACM, 2025, Art. 274. doi: 10.1145/3706598.3713670

[23] T. Armstrong, The Power of Neurodiversity: Unleashing the Advantages of Your Differently Wired Brain. Cambridge, MA: Da Capo Lifelong Books, 2011.

[24] J. den Houting, “Neurodiversity: An insider’s perspective,” Autism, vol. 23, no. 2, pp. 271–273, 2019. doi: 10.1177/1362361318820762

[25] A. Meyer, D. H. Rose, and D. T. Gordon, Universal Design for Learning: Theory and Practice. Wakefield, MA: CAST Professional Publishing, 2014.

[26] I. Scott, “Addressing the ‘hidden curriculum’ in higher education: The challenges for students with autism spectrum disorder,” Journal of Further and Higher Education, vol. 43, no. 3, pp. 303–315, 2019. doi: 10.1080/0309877X.2017.1359502

[27] E. Salas, N. J. Cooke, and M. A. Rosen, “On teams, teamwork, and team performance: Discoveries and developments,” Human Factors, vol. 50, no. 3, pp. 540–547, 2008. doi: 10.1518/001872008X288457

[28] J. A. Durlak, R. P. Weissberg, A. B. Dymnicki, R. D. Taylor, and K. B. Schellinger, “The impact of enhancing students’ social and emotional learning: A meta-analysis of school-based universal interventions,” Child Development, vol. 82, no. 1, pp. 405–432, 2011. doi: 10.1111/j.1467-8624.2010.01564.x

[29] Y. Lindsjørn, D. I. K. Sjøberg, T. Dingsøyr, G. R. Bergersen, and T. Dybå, “Teamwork quality and project success in software development: A survey of agile development teams,” Journal of Systems and Software, vol. 122, pp. 274–286, 2016. doi: 10.1016/j.jss.2016.09.028

[30] D. Long and B. Magerko, “What is AI literacy? Competencies and design considerations,” in Proc. CHI ‘20, New York, NY: ACM, 2020, pp. 1–16. doi: 10.1145/3313831.3376727

[31] D. T. K. Ng, J. K. L. Leung, S. K. W. Chu, and M. S. Qiao, “Conceptualizing AI literacy: An exploratory review,” Computers and Education: Artificial Intelligence, vol. 2, p. 100041, 2021. doi: 10.1016/j.caeai.2021.100041

[32] A. Cox, “Algorithmic literacy, AI literacy, and responsible generative AI literacy,” Communication Education, vol. 73, no. 4, pp. 347–362, 2024. doi: 10.1080/03634523.2024.2346049

[33] Design-Based Research Collective, “Design-based research: An emerging paradigm for educational inquiry,” Educational Researcher, vol. 32, no. 1, pp. 5–8, 2003. doi: 10.3102/0013189X032001005

[34] S. Barab and K. Squire, “Design-based research: Putting a stake in the ground,” Journal of the Learning Sciences, vol. 13, no. 1, pp. 1–14, 2004. doi: 10.1207/s15327809jls1301_1

[35] D. M. Boyer, "External Evaluation: Beginnings—Preparing Autistic Students for the AI Workforce. Year 2 Evaluation Report: Spring-Summer 2025 Implementation," Clemson Univ., Clemson, SC, Rep. NSF #2322554, Oct. 2025.

[36] V. Braun and V. Clarke, “Using thematic analysis in psychology,” Qualitative Research in Psychology, vol. 3, no. 2, pp. 77–101, 2006. doi: 10.1191/1478088706qp063oa

[37] A. Al-Azawei, F. Serenelli, and K. Lundqvist, “Universal design for learning (UDL): A content analysis of peer-reviewed journal papers from 2012 to 2015,” Journal of the

Scholarship of Teaching and Learning, vol. 16, no. 3, pp. 39–56, 2016. doi: 10.14434/josotl.v16i3.19295

[38] Stanford Human-Centered Artificial Intelligence (HAI), “Exploring the dangers of AI in mental health care,” Jun. 11, 2025. Online. Available: https://hai.stanford.edu/news/exploring-the-dangers-of-ai-in-mental-health-care

[39] American Psychological Association, “Health advisory: Use of generative AI chatbots and wellness applications for mental health,” 2025. Online. Available: https://www.apa.org/topics/artificial-intelligence-machine-learning/health-advisory-chatbots-well ness-apps

[40] The Guardian, “Mind launches inquiry into AI and mental health after Guardian investigation,” Feb. 20, 2026. Online Available: https://www.theguardian.com/technology/2026/feb/20/mind-inquiry-google-ai-overviews-mental -health-guardian-investigation

[41] M. Sharma et al., “Towards understanding sycophancy in language models,” arXiv:2310.13548, 2023. Online. Available: https://arxiv.org/abs/2310.13548

[42] J. Hong, G. Byun, S. Kim, and K. Shu, “Measuring sycophancy of language models in multi-turn dialogues,” in Findings of the Association for Computational Linguistics: EMNLP 2025, 2025, pp. 2239–2259, doi: 10.18653/v1/2025.findings-emnlp.121.

[43] N. Kosmyna et al., “Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task,” arXiv:2506.08872, 2025, doi: 10.48550/arXiv.2506.08872. Online. Available: https://arxiv.org/abs/2506.08872

[44] R. Schechner, Performance Theory. New York, NY: Routledge, 2003.