Inclusive AI Education Supporting Neurodivergent Learners

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.

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.

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

NSF ExLENT: Inclusive AI Education Supporting Neurodivergent Learners

Abstract

This paper synthesizes two years of implementation and evaluation findings from PAS4AI: Preparing Autistic Students for the AI Workforce, an NSF ExLENT design-based research program preparing autistic community college students for AI and software careers. The program evolved from technical skills instruction to an integrated approach that addresses the gap between technical competence and workforce readiness for autistic learners. This shift responded to Year 1 and Year 2 evaluation findings that identified persistent interview anxiety, uneven instructor preparation, and the need for explicit social-emotional learning (SEL) support as barriers to student success. Using semi-structured interviews with project team members and thematic analysis, supplemented by project documentation, we report: (a) why generic communication instruction proved insufficient under stress, (b) how we operationalized SEL into coachable microskills, (c) why interviews emerged as a critical barrier requiring deliberate practice and transparent rehearsal supports, and (d) how scenario-based instructor preparation may support more consistent enactment of inclusive teaching practices. Because this was a small, purpose-built program for autistic learners, we do not generalize from these data to all instructors, all autistic students, or all neurodivergent learners. Instead, we identify adaptable design principles for AI education programs: build SEL supports into technical learning environments, use flexible participation structures that do not require disability disclosure, and assess technical competence through artifact-based demonstrations and high-pressure interviews.

Keywords: *neurodiversity; autism; AI education; social-emotional learning; self-efficacy; instructor preparation; internships*

Introduction

Autistic adults face disproportionately high unemployment despite technical aptitudes that align well with computing and AI work. While many autistic adults demonstrate strengths in pattern recognition, systematic thinking, and attention to detail [1], [2], their unemployment rate exceeds 85% [3]. This disparity persists not primarily because of insufficient technical ability, but because of invisible social access barriers embedded in how technical workplaces and hiring processes are structured [4]. Traditional computing education often assumes uniform participation norms, i.e., fast conversational turn-taking, rapid inference under time constraints, and socially fluent self-presentation, and these norms function as gatekeeping mechanisms that suppress demonstrations of competence even when technical ability is strong [4], [5].

Computing education has begun to recognize neurodiversity as a source of strength rather than a collection of deficits [6], [7]. Yet, few programs have addressed what we term the enactment problem: students who can describe communication strategies in low-stakes contexts often cannot deploy them under the performance pressure of career gateways, e.g., mock interviews or real-time collaborative debugging. Universal Design for Learning (UDL) [8] provides a useful framework for broadening access. Still, UDL-informed communication instruction has not consistently bridged the gap between strategy acquisition and situated enactment in the PAS4AI context. This paper reports on how PAS4AI addressed that gap across two design-based research cycles, contributing an integrated framework of SEL microskills, scenario-based instructor preparation, and ethical AI literacy that reframes workplace readiness as a stress-context performance problem. This work extends prior findings on instructor experience and social-emotional demands in neurodiversity-focused AI education [14], [15], while systematically responding to external evaluator recommendations [13].

Project Description

PAS4AI: Preparing Autistic Students for the AI Workforce is an NSF ExLENT community college-based bridge program preparing autistic students for entry-level AI and software careers. The program integrates technical training in Python, TensorFlow, Keras, and Jupyter Notebooks with structured internship experiences across an 18- to 22-week curriculum: a 10-week virtual training phase of approximately 20 hours per week followed by an 8- to 12-week paid internship emphasizing responsible AI design. Students are recruited nationwide from community colleges, and eligibility requires completion of two introductory programming courses and one mathematics course. Year 1, implemented in Summer 2024, served 16 students. Year 2, implemented in Spring and Summer 2025, served nine students.

The program is structured as design-based research [9], [10], in which successive implementation cycles surface design tensions that drive targeted redesign. Year 1 established a foundational model combining CS content with generic communication training. Year 2 executed a strategic pivot to an integrated SEL and ethical AI framework, informed directly by Year 1 evaluation findings. Year 3, planned for Spring 2026 as an unfunded sustainability extension, is testing train-the-trainer modules, portfolio-based assessment, and alumni peer mentorship.

Methodology

This study draws on qualitative data collected through Year 2 implementation and external evaluation cycles. The primary dataset consists of semi-structured interviews, approximately 45 to 60 minutes each, conducted by the research team and an independent external evaluator. Interviewed participants were project team members rather than students. The interview sample included five instructional-team members and mentors, primarily PhD students and a postdoctoral researcher, and two job coaches specializing in neurodivergent workplace support. Because instructional and mentoring roles overlapped in this small program, we report this group as the instructional team rather than making strong claims about role-based differences among instructors, mentors, and coaches.

We did not interview students as part of Year 2 data collection. Therefore, the findings reported here reflect instructor, mentor, job coach, evaluator, and documentation, not direct student accounts. This distinction is important: the study does not measure student SEL development, nor does it claim to compare instructors' practices through systematic observation.

We analyzed interview transcripts and project documentation using thematic analysis [11], following an iterative process of open coding to generate initial themes and axial coding to explore relationships among themes. The independent external evaluator conducted a parallel qualitative analysis focused on project efficacy, project improvement, partnerships, sustainability and scalability, and fidelity. The Year 2 evaluation identified several convergent themes, including the SEL/self-efficacy pivot, interview anxiety, portfolio-based assessment, ethical use of LLMs, and instructor preparation needs. The evaluation also noted that instructional staff did not consistently feel knowledgeable or autonomous enough to apply new SEL practices without explicit direction, suggesting an implementation-fidelity issue rather than a directly measured construct of instructor variability [13].

Because we conducted this study in a purpose-built program designed for autistic community college students, the program context differed from most general engineering or computing classrooms, where instructors typically do not know which students are autistic unless students choose to disclose. For that reason, we frame the implications of this study as design principles for inclusive learning environments rather than diagnostic guidance for identifying or responding to individual students.

Findings

Three core findings emerged from Year 2 implementation and evaluation data.

Finding 1: From Generic Communication Instruction to SEL Microskills

Year 2 data showed that UDL-informed communication skills instruction did not reliably translate into performance in high-stress contexts. A consistent pattern emerged across instructor accounts: students understood the concept of communication strategies but could not enact them under performance pressure, uncertainty, and heightened social demands. One instructor summarized the failure mode directly: “They know what to do but cannot do it in the moment.”

Interview anxiety caused capable learners to freeze during mock interviews; collaborative debugging required instructor moderation that students could not self-initiate; and students performed required activities because they were told to, rather than from an internalized understanding. In response, the program operationalized abstract SEL competencies into three discrete, coachable microskills functioning as prerequisites for successful enactment. Table 1 summarizes the breakdown scenarios, observable indicators, and coaching responses documented during Year 2.

Domain	Breakdown Scenario	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

Domain	Breakdown Scenario	Observable Indicator	Coaching Response
Self-advocacy	Anonymous routing	Student privately messages instructor to ask a question anonymously	Graduated disclosure practice; instructor models vulnerability
Emotional regulation	Camera shutdown	Student clams up and turns off camera and audio	Non-verbal Miro boards; text contribution as legitimate participation
Emotional regulation	Processing mismatch	Student falls behind in rapid turn-taking and cannot formulate responses quickly	Extended wait time; asynchronous ideation before synchronous discussion
Emotional regulation	Team conflict	Student needs external moderation during collaborative debugging	Instructor-facilitated resolution; structured turn-taking
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

Table 1. SEL Microskill Breakdown Scenarios and Coaching Responses

Finding 2: Instructor Preparation and Implementation Fidelity

A second Year 2 finding concerned instructor preparation and implementation fidelity. In the external evaluation, project team members reported that instructional staff did not always feel knowledgeable or autonomous enough to apply the new SEL interventions independently. The concern was not that the study directly measured differences among instructors through observation or comparative performance ratings. Rather, interview data suggested that inclusive practices had not yet been sufficiently internalized across the instructional team. Some instructors described needing clearer routines, shared language, and scenario-based preparation for responding to student stress, communication breakdowns, and team conflict.

We therefore treat this finding as evidence of an implementation problem: inclusive intent does not automatically translate into consistent inclusive practice, especially when instructors are asked to coach SEL microskills, support technical learning, and manage team dynamics within a compressed AI training program. In the Year 2 external evaluation, we concluded that scenario-based train-the-trainer modules were needed to support consistent implementation of SEL coaching and portfolio-based assessment [13].

The design response was a scenario-based train-the-trainer model covering SEL microskills coaching, ethical LLM rehearsal boundaries, and non-verbal collaboration scaffolds. Rather than asking instructors to infer students' needs from diagnostic categories, the model focuses on observable instructional situations that can arise for many learners: a student going silent during a high-pressure exchange, difficulty entering a rapid-turn conversation, uncertainty about how to ask for clarification, or reluctance to request help. We designed these scenarios to support instructor decision-making without requiring instructors to know, infer, or label a student's disability status.

This reframing is especially important for transfer beyond PAS4AI. In this program, participants were recruited specifically from the autistic community college student population preparing for AI careers. In most engineering and computing classrooms, instructors cannot assume they know which students are autistic, neurodivergent, anxious, or otherwise experiencing social-performance stress. The actionable implication is therefore not "identify autistic students and give them SEL support." The implication is to design classroom participation, collaboration, and assessment structures that make support available without requiring disclosure: multiple modes of participation, planned wait time, written or asynchronous contributions before synchronous discussion, explicit clarification scripts, low-stakes rehearsal, and artifact-based demonstrations of learning.

Finding 3: LLMs as Rehearsal Partners and Ethical AI Literacy

Instructors identified large language models (LLMs) as potentially valuable nonjudgmental rehearsal partners for practicing interview responses and debugging explanations. The Year 2 evaluation described LLMs as useful for low-stakes practice while also identifying the need for clearer ethical boundaries around prompt engineering, critical evaluation, and uncritical copying [13]. The program positions LLMs as one scaffold within a graduated rehearsal sequence:

1. Low social threat: LLM partners for initial rehearsal and self-paced practice.
2. Moderate social threat: Peer-to-peer practice with program cohort members.
3. Authentic stakes: Instructor-facilitated mock scenarios and artifact talk-throughs.

This sequencing is intentional: LLMs are an entry point for rehearsal, not the endpoint. Instructors also observed students using LLMs uncritically, e.g., copying outputs without verification or concealing AI use rather than engaging it as a learning tool. The resulting design response integrates ethical AI pedagogy from the first week of instruction, positioning AI use as an ethical decision-making practice requiring the same self-awareness and self-advocacy as the SEL microskills above: rehearsal versus substitution, transparent versus concealed use, and reflection after AI interactions.

Discussion and Implications

Taken together, these findings point to a core reframing: workplace readiness for neurodivergent learners is not primarily a content acquisition problem but a stress-context performance problem. Learners who demonstrate competence in low-stakes settings may fail to enact that competence in high-pressure career gateways, not because of skill deficits, but because performance conditions are structurally misaligned with how their competencies emerge. Technical interviews, rapid-turn collaborative debugging, and unsupported disclosure decisions function as invisible access barriers [4], [5] when they implicitly privilege particular interaction styles.

Implications for General Engineering and Computing Classrooms

The PAS4AI context differs importantly from general engineering and computing classrooms: students enter a program explicitly designed for autistic learners, whereas instructors in most courses do not know which students are autistic unless students disclose. Instructors also should not infer neurodivergence from outward signs of stress, withdrawal, communication differences, or anxiety. Such behaviors can arise for many reasons, including test anxiety, unfamiliarity with disciplinary norms, language background, prior educational experience, disability, or temporary stress.

For that reason, the most transferable implications of this work are pedagogical rather than diagnostic. The SEL microskills framework may be treated as a universal design support for high-pressure technical learning rather than as a separate intervention reserved for students with known diagnoses. Practices such as explicit expectations for teamwork, multiple participation modes, non-verbal collaboration boards, structured turn-taking, low-stakes rehearsal before high-stakes performance, and portfolio-based artifact talk-throughs can support autistic learners and students who experience anxiety, communication barriers, or uncertainty in technical collaboration. This framing avoids requiring students to disclose disability status to receive support and avoids asking instructors to determine which students "need" SEL. Instead, it treats SEL scaffolding as part of the instructional infrastructure of inclusive AI education.

Adaptable Design Elements

Three design elements appear transferable to other computing programs. First, the SEL microskills framing provides a practical unit of design that can be integrated into existing collaboration pedagogy: operationalizing abstract competencies into discrete, observable, coachable behaviors that instructors can notice, prompt, and reinforce during authentic technical tasks. Second, scenario-based instructor preparation targets enactment by preparing instructors to respond to predictable breakdowns with consistent coaching moves, rather than relying on generic one-time workshops. Third, portfolio-based assessment, i.e., replacing or augmenting high-stakes interviews with artifact talk-throughs, shifts evaluation from on-demand social performance toward structured demonstration of technical work. This is both an accommodation and a structural challenge to the assumption that time-pressured self-presentation is a valid proxy for technical competency [13].

The train-the-trainer approach also serves as a sustainability mechanism, shifting inclusive practice from individual instructor disposition toward shared, scenario-based preparation that can be maintained across cohorts and staffing changes. We developed open educational resources developed intended to support adaptation by other educators; however, we do not claim that someone can reproduce the full PAS4AI model without attention to local resources, instructor capacity, and student support structures.

Limitations

Several limitations constrain the interpretation of these findings. First, the Year 2 student cohort was small, with nine students, and the program was a single, purpose-built intervention for autistic community college students preparing for AI and software careers. The reader should not generalize the findings to all autistic learners, all neurodivergent learners, or all AI education contexts.

Second, we did not interview students as part of the Year 2 evaluation data collection. Findings reflect project team member perspectives, external evaluation analysis, and program documentation rather than direct student accounts. Future work should incorporate student interviews, student-facing surveys, and longitudinal follow-up to examine whether learners themselves experience the proposed SEL microskills and portfolio-based assessments as supportive.

Third, the evidence base for claims about instructor preparation and implementation fidelity is limited. The study did not use systematic classroom observations, video analysis, implementation checklists, or comparative measures of instructor practice. Interviewees reported differences in comfort, preparation, and enactment of inclusive practices, and the external evaluation identified instructor training as an area for improvement. However, this evidence does not support strong claims about measured instructor variability. We therefore frame the finding as a need for clearer instructor preparation and implementation support, rather than a quantified comparison of instructor behavior.

Fourth, the PAS4AI context differs from ordinary engineering and computing classrooms because students choose to enter into a program explicitly designed for autistic learners. In typical classrooms, instructors generally do not know which students are autistic unless a student discloses, and instructors should not infer disability status from behaviors such as silence, anxiety, reduced eye contact, or withdrawal. Autism is also highly heterogeneous; even when a student discloses, the instructional implications are not uniform. Some students may have extensive prior experience with self-advocacy and social navigation, while others may require more explicit scaffolding. Similar stress responses can also occur among students who are not autistic.

These limitations shape the paper's practical takeaway. We do not argue that instructors should identify autistic students or apply diagnosis-specific interventions based on observed behavior. Instead, we argue for inclusive instructional designs that reduce reliance on disclosure: flexible participation modes, explicit collaboration routines, low-stakes rehearsal, ethical LLM-supported practice, and portfolio-based assessment structures that allow students to demonstrate competence without depending solely on high-pressure social performance.

Conclusion

This paper documents how two implementation cycles of PAS4AI drove a shift from generic communication instruction toward enactment-oriented, SEL-integrated design. The central contribution is not a claim that instructor variability was measured as a generalizable construct. Rather, the contribution is a design argument: inclusive AI education requires instructional conditions that help learners enact technical, collaborative, and self-advocacy skills under stress.

By reframing workplace readiness as a stress-context performance problem, the focus moves from individual deficits to the learning conditions that shape access to AI careers. For PAS4AI participants, these conditions included mock interviews, collaborative debugging, disclosure decisions, and internship preparation. For broader engineering and computing education, the actionable takeaway is to design support structures that do not depend on knowing which students are autistic or neurodivergent. Instructors can make participation expectations explicit, provide multiple ways to contribute, build low-stakes rehearsal into technical courses, use LLMs transparently as bounded rehearsal tools, and assess competence through portfolios and artifact talk-throughs in addition to high-pressure interviews.

Creating equitable pathways into AI work requires more than access to technical content. It requires building learning environments where students can demonstrate competence, seek clarification, advocate for their needs, and participate in collaboration without being filtered out by misaligned participation norms. PAS4AI offers an adaptable design framework for this work while also underscoring the need for stronger future evidence, including direct student perspectives, observation-based fidelity measures, and longitudinal employment outcomes.

Acknowledgment

This material is based upon work supported by the National Science Foundation ExLENT program 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 National Science Foundation.

References

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

[2] 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.

[3] 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.

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

[5] 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.

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

[7] J. den Houting, "Neurodiversity: An insider's perspective," *Autism*, vol. 23, no. 2, pp. 271–273, 2019.

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

[9] Design-Based Research Collective, "Design-based research: An emerging paradigm for educational inquiry," *Educational Researcher*, vol. 32, no. 1, pp. 5–8, 2003.

[10] 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.

[11] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative Research in Psychology*, vol. 3, no. 2, pp. 77–101, 2006.

[12] 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.

[13] 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.

[14] A. Begel, D. M. Boyer, R. Kubina, S. Asadi, T. Mishra, R. Butler, and J. Jang, "Investigating Instructors' Experiences in a Neurodiversity-Focused AI Training Program," in *Proc. ASEE Annual Conference & Exposition*, 2025.

[15] R. Butler, D. M. Boyer, A. Begel, R. Kubina, S. Asadi, T. Mishra, and J. Jang, "Navigating the Social-Emotional Landscape of Neurodiversity in AI Education," in *Proc. ASEE Annual Conference & Exposition*, 2025.