

# Self-Regulated Personal Contracts as a Harm Reduction Approach to Generative AI in Undergraduate Programming Education

Aadarsh Padiyath  
University of Michigan  
Ann Arbor, USA  
aadarsh@umich.edu

Jessica Shen  
University of Michigan  
Ann Arbor, USA  
jesishen@umich.edu

Barbara Ericson  
University of Michigan  
Ann Arbor, USA  
barbarer@umich.edu

## Abstract

Students learning programming exercise agency in deciding when and how to use GenAI tools like ChatGPT. However, this agency is often implicit and shaped by deadline pressure and peer behavior rather than explicit and conscious learning goals. We designed a GenAI Contract grounded in harm reduction and self-regulated learning theory to scaffold intentional decision-making: students articulated personal learning goals, created usage guidelines, and reflected on alignment at strategic points across an eleven-week semester. The contract was non-binding and graded only for completion, emphasizing self-awareness over enforcement. We implemented this with N=217 students in an intermediate Python course. For students still forming their relationship with GenAI, it worked, as 58% of students reported the intervention changing their thinking and created helpful accountability structures. However, awareness did not always translate to sustained behavior change. Some students who valued their guidelines still abandoned them under various pressures. Maintaining guidelines required constant self-control across hundreds of decisions, while using GenAI freely requires none. Many students could not sustain this burden despite this self-awareness. We discuss supporting student agency when GenAI tools and learning goals create tension.

## CCS Concepts

• **Social and professional topics** → **Computer science education**.

## Keywords

Programming Education, CS1, LLMs, Self-Regulated Learning

### ACM Reference Format:

Aadarsh Padiyath, Jessica Shen, and Barbara Ericson. 2026. Self-Regulated Personal Contracts as a Harm Reduction Approach to Generative AI in Undergraduate Programming Education. In *Proceedings of the 31st ACM Conference on Innovation and Technology in Computer Science Education V.1 (ITiCSE 2026)*, July 13–15, 2026, Madrid, Spain. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/XXXXXXX.XXXXXXX>

## 1 Introduction

Students learning programming today face a stream of decisions [7]: Should I ask ChatGPT to explain this error message? Should I use Copilot's suggestion, or write the code myself? Should I generate a

solution to check my work, or would that short-circuit my learning? These decisions happen constantly, often shaped by factors beyond the tools themselves [3]: peer behavior and career anxieties [24], late-night exhaustion and deadline pressure [6], as well as genuine curiosity [4]. Most of these choices are implicit; students rarely pause to ask whether their Generative AI (GenAI) use aligns with their learning goals [24, 27].

Instructors face their own choices. Some prohibit GenAI entirely [16], others embrace it enthusiastically, believing 'resistance is futile' [16, 28]. Some chart a middle ground, placing restrictions on certain uses [1, 26]. These are not neutral responses. Each policy choice influences student behavior and classroom culture [1]. Yet any policy should acknowledge the reality that programming students are not a monolith. Some are training for software development careers, others are curious about programming but do not expect themselves to code professionally [5]. Some arrive confident in their own programming skills; others are struggling with basic computational thinking. One-size-fits-all rules cannot serve this diverse population equally well [23]. Different students may need different GenAI approaches based on their unique goals, motivations, and futures. Yet GenAI use is shadowed by judgment: students worry GenAI reliance will be seen as incompetence [24, 30], while instructors debate which uses constitute academic dishonesty [16, 26]. This climate of judgment makes it difficult to have honest conversations about what students are actually doing and why.

*Harm Reduction* [21] offers a philosophical framework designed for exactly this tension: when something is widely accessible, when people may use it regardless of rules, and when judgment may drive behavior underground rather than preventing it. Originating in public health, the philosophy acknowledges that people make choices regardless of official policy [21]. Rather than mandating abstinence or morally condemning use, harm reduction focuses on supporting better decision-making. For example, comprehensive sex education provides information and access to contraception rather than relying solely on abstinence-only education [15, 22]. Applied to our context, this means recognizing that students encounter GenAI tools as readily available options [10, 14], that 'appropriate use' varies by student and goal [24], and our role as educators is helping them navigate their choices thoughtfully rather than policing behavior.

However, thoughtful navigation requires substantial self-regulatory skill [7]. Students need to monitor whether a strategy serves their goals, recognize when it is leading them astray, and adjust their approach accordingly [7, 29]. This kind of intentional practice takes time and experience to develop, and cannot be reduced to a universal rule ("always use GenAI for X, never for Y"). What we can do is create opportunities for students to build this capacity towards



This work is licensed under a Creative Commons Attribution 4.0 International License. *ITiCSE 2026, Madrid, Spain*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM  
<https://doi.org/10.1145/XXXXXXX.XXXXXXX>

aligning their GenAI use (and non-use) with their personal goals rather than following external rules they may resent or ignore.

This paper describes our implementation of a "*GenAI Contract*": a structured reflection tool grounded in self-regulated learning theory and harm reduction principles. Students set their own learning goals, create personal usage guidelines, and reflect multiple times throughout the semester on whether their choices align with their aspirations. Importantly, the contract is non-binding and not enforced. We designed four reflection prompts across eleven weeks, grading for completion rather than content.

We implemented this contract reflection process with  $N=217$  students in an intermediate Python course. Our goals were (1) to understand whether students would meaningfully engage with structured reflection about their GenAI choices; (2) to identify patterns in how they set goals and adjusted their approach over time; and (3) to determine what aspects of this intervention actually supported or hindered the development of intentional practice. This experience report shares what we learned: what worked, what didn't, and how other instructors might adapt this approach.

## 2 Prior and Related Work

### 2.1 Designing for Agency: Self-Regulation and Harm Reduction

Problem-solving and critical thinking skills are essential to learn in programming, but many studies show a tension between building these skills and using AI for learning assistance [3, 12]. The recent adoption of GenAI tools has made it increasingly easy for students to bypass active problem-solving steps. Although these tools can be helpful, they can also lead to over-reliance or over-trust in AI's help [3]. AI being easily accessible makes it difficult for students to judge when to rely on the tool versus their own reasoning, often leading to a lack of intention when using the tools [24, 27]. This has led to a widening gap between students who can and can't manage these tools effectively [27]. We draw from robust theoretical frameworks to guide students toward intentional AI use and increase self-awareness in learning practices.

We draw on Self-Regulated Learning (SRL): the active process of executing cognitive control during a task [19]. SRL depends on two foundational elements: metacognition and self-efficacy. Metacognition is a student's knowledge about their own cognitive control and effective strategies. This is important in the context of GenAI as students may over-trust the tool, leading to relying on the tool rather than building their own understanding. Self-efficacy is the belief in one's capacity to successfully execute a specific task [2, 19, 20]. This belief is crucial because self-efficacy influences a student's persistence and internal motivation when faced with challenges [19]. Without the students' awareness of their use of GenAI and the confidence in their own abilities, students lack the necessary tools to confidently navigate the learning process. To operationalize SRL, we use Zimmerman's continuous three-phase model, which breaks the learning process down into forethought (planning), performance (monitoring), and self-reflection (judgement) [32]. By implementing this loop, SRL can put students in charge of how they think and use GenAI, preventing them from falling into a passive role when using the tool. This approach directly scaffolds

the self-awareness and metacognition required to intentionally use GenAI tools in a responsible manner in their learning process.

Our approach is also rooted in harm reduction philosophy (HR) [21]. We chose this approach because GenAI tools are readily accessible, and research shows that students may turn to them because of time pressure and low motivation [10, 13], regardless of policies that try and control student behavior. HR acknowledges that GenAI use can be both helpful and harmful for learning and therefore shifts the focus. Rather than promoting abstinence or morally condemning use, it focuses on encouraging more responsible and intentional use, which respects student agency and supports informed decision-making. This intervention aims to close the gap between students who can and can't manage the tool effectively.

Finally, implementation intention plans combine the SRL framework and HR into an actionable practice [8]. These plans are action-planning strategies that link situational cues to a desired goal [8]. For example, a student's plan may be "If I am stuck on an error for more than 30 minutes, then I will ask GenAI to explain the concept and an example without asking for the exact solution". This commitment statement allows the student to maintain intentionality and discourages turning to GenAI too early. By having students formalize their planned GenAI usage into "contracts" ahead of time, we operationalize what Padiyath [23] theorized as necessary: making these implicit goals explicit. This allows students to consciously align their GenAI use with their learning goals and directly builds off the SRL forethought phase. Furthermore, revisiting and revising the contract scaffolds the self-monitoring phase and allows students to make adjustments to their strategies as they learn and gain experience. Lastly, asking students to evaluate their usage aligns with the self-reflection phase. By integrating our theoretical model into a structured tool in the form of a GenAI contract, we investigate how to empower programming students to become more self-aware of their learning process.

### 2.2 Programming Students and GenAI Tools

How programming students actually use generative AI is more nuanced than the utopian or dystopian narratives suggest [24]. Many students do report genuine benefits from GenAI tools. Many use them to quickly debug errors, understand error messages, or explore alternative approaches to a problem [24, 26]. For some, the tools lower barriers to larger projects as well as expressing themselves through programming [28, 31]. However, recent research also reveals patterns of use that may be troubling [3, 24, 27]. Students who use GenAI early and frequently report decreased confidence in their own abilities [20, 24]. Some describe a passivity that develops as they copy GenAI-created solutions without engaging with the underlying logic or code. Notably, many students either do not recognize this [27], or only recognize it after it affects their academic performance [24].

This pattern suggests the tools' affordances for speed and productivity may work against learning and understanding. After all, GenAI tools were not specifically designed with educational contexts in mind. Their affordances (generating complete, expert-like code instantly) combined with formal institutional structures and curricula, create a powerful pull toward productivity-focused use rather than learning-focused use [24]. While feeling efficient,

productivity-focused use may undermine students' learning goals. Students with substantial self-regulatory capacity can use GenAI in ways aligned with learning, but doing so requires conscious resistance to the tool's default affordances [7, 27]. Yet general self-regulation skills may not transfer automatically, as Margulieux et al. found no relationship between students' self-regulation strategies and their AI use [20]. In essence, students need to develop intentional use (and deliberate non-use) of GenAI: the ability to consciously choose how GenAI fits into their learning goals, rather than defaulting to passive use patterns [23].

## 2.3 Approaches to Supporting Intentional Use

Educators and researchers have proposed various integrations of GenAI in programming education to support intentional use and learning, yet most focus on modifying the tool itself – restricting access, scaffolding its outputs, or teaching students how to use it more effectively [11, 17, 18]. These tool-focused approaches often rest on the assumption that if we design better tools or teach better tool use, students will learn better. However, technology adoption and use are shaped by both technical *and* human factors including student goals, peer influence, institutional contexts, and crucially, student agency and intentionality [23, 24].

SRL research offers a more learner-centered lens, but applying it in programming education can be cognitively demanding. Parham et al. found that programming students tend to skip important metacognitive steps like planning and monitoring in favor of immediately writing code [25]. Wang et al.'s controlled study integrating ChatGPT with Zimmerman's SRL framework found improvements in students' motivation and self-efficacy, but no significant gains in learning outcomes, which the authors attribute to cognitive overload from managing SRL strategies alongside programming itself [29]. Critically, when scaffolding was minimal, students bypassed self-regulatory prompts to complete tasks more quickly [29]. These findings suggest that even when SRL is embedded within tool design, GenAI's frictionless affordances make reflective prompts easy to bypass. This matters because modifying tools alone cannot account for how students actually decide to use GenAI [11], and whether that use actually aligns with their learning goals [24, 27]. Few interventions directly address and support the metacognitive and self-regulatory dimensions of this decision-making [3]. Drawing on self-regulated learning theory and harm reduction frameworks, our approach places self-regulation outside the tool interaction entirely, making reflection the explicit focus rather than an add-on to an already demanding task.

## 3 Intervention Design and Context

### 3.1 Course Context

Our intervention took place in Fall 2025 in an intermediate-level Python programming course at a large public research university in the Midwest United States. The course enrolled N=217 students taught by one instructor alongside twelve undergraduate TAs. This course focused on the following Python skills: functions, conditionals, object-oriented programming basics, debugging, Git version control, Regular Expressions, I/O, APIs, and databases. Eight homework assignments and two projects gave students opportunities to

practice these skills. Two proctored midterm exams assessed students' abilities without external resources including GenAI. To take this course, students completed at least one of several prerequisite introductory programming courses.

Our student population contained approximately 60% Information Science majors (for whom the course was required), 10% Computer Science majors, and 30% from various other programs. The Information Science students typically envisioned futures as UX designers, data scientists, or "conversational programmers" [5]. Computer Science students, whose required sequence used C++, often take this course specifically to learn Python. Due to these varying degrees of expertise, a one-size-fits-all policy on GenAI use would fit no one particularly well. Of note, our syllabus stated: "You can also use generative AI like ChatGPT, but must list that as a thing you worked with and what it helped you with."

### 3.2 GenAI Contract Design

Given this context (diverse student goals, permissive GenAI policy, and free access to available AI tools), we needed an intervention that could support students in making thoughtful choices when using GenAI tools.

Four design principles, each grounded in our theoretical frameworks, shaped our contract. First, it would be non-binding, explicitly presented as a reflection tool rather than an enforced contract, reflecting HR's 'come as you are' philosophy of supported decision-making over judgment [21]. Second, it would be goal-driven: students articulate personal learning objectives before considering GenAI strategies, operationalizing Zimmerman's forethought phase [32]; the if-then guideline format then draws on implementation intention theory [8] to translate those goals into concrete, situationally-anchored action plans. Third, it would support iterative reflection, with students revisiting and revising guidelines at strategic points across the semester, scaffolding Zimmerman's performance monitoring and self-reflection phases [32] rather than producing a one-time artifact. Fourth, it would afford a diversity of approaches, consistent with both HR's recognition that appropriate use varies by individual [21] and with the reality that our students had genuinely different learning goals.

The contract structure translated these into four reflection prompts across eleven weeks (see Figures 1 and 2). In Week 2, students completed Steps 1 and 2 together as a single Google Doc assignment. Steps 3 and 4 were distributed as text for students to copy into their existing contract document. Step 3 was distributed in Week 5, after students had taken their first midterm. This timing was chosen due to its finding as a important time for reflection on GenAI usage [24]. The final reflection, Step 4, came in Week 11 after the second midterm. By this point, students had completed much coursework and taken both midterms. Each reflection was worth a small number of completion points to incentivize participation (1.5% of their final grade). We graded only for completion, never for content.

### 3.3 Data Collection and Analysis

This methodology was approved by our institution's review board. Our primary data source for analysis was student contract submissions. To understand what worked and what didn't, the research team read through all contract submissions, taking detailed memo

### GenAI Contract

**Background**  
Students don't just "use" GenAI tools like ChatGPT – they decide when, how, and why to use them. These choices are often implicit, shaped by goals, habits, and pressures. This contract is a simple reflection tool to help you make those decisions more deliberate.

**Important:** **This contract will not be enforced.** It is only for your personal reflection and growth. The goal isn't to get it "right" the first time; it is to increase your awareness of how GenAI fits with your learning.

**Step 1: Your Goals in This Course**  
Take a minute and actually read through the course description for this course, taken from the syllabus.

**COURSE DESCRIPTION**  
This course will focus on giving you a strong background in data-oriented programming, in the Python programming language. It is intended for students who have completed an introductory programming course and are moving on to the next step in a data-oriented fashion. Learning objectives for this course include:

- Develop intermediate programming skills in Python
- Practice using basic data structures (lists, tuples, dictionaries)
- Learn how to use a terminal window to run a Python program locally on your computer
- Familiarity with the basics of object-oriented programming: objects, classes, and inheritance
- Develop debugging and testing skills
- Use a distributed code repository (git and GitHub).
- Develop experience with using pattern matching (RegEx)
- Work with data from a variety of sources (files, web, APIs, JSON, and databases)
- Work with relational databases with SQL

Before thinking about how you might use GenAI, pause to consider what you want to get out of this course. Your goals might relate to your major, career plans, GPA, personal interests, skills you want to build, etc.  
Be specific. Vague statements like "I want to do well in this class" aren't useful as something you can act on and reflect on later.  
Start with at least one clear personal goal below, and feel free to list more if you like. You can use this format to describe your goals if you want: In this course, I *<your goal here>* because *<why that goal makes sense for you>*.

**Step 2: Your Personal Approach to GenAI**  
Think about how you might use GenAI – or not – throughout this course. The goal is to **make your thinking explicit** so you can revisit it later.  
**You can describe situations where you plan to use GenAI, situations where you plan not to, or a mix.**  
**These guidelines are for you and you alone.** You do not have to follow these guidelines perfectly; they exist only to help you reflect and make intentional choices. Focus on your reasoning – what matters to you, and why – rather than trying to follow a "right" pattern.  
Start with **at least one clear guideline about your personal approach to GenAI.** You can add as many additional guidelines as you wish. Just make sure each one is relevant to you and your goals in this course. You can use this format to describe your guidelines if you wish: I will *<the way you will use/not use GenAI>* when *<condition/constraint in which you will/won't use GenAI>* because *<reasoning that makes sense for you>*.

**Figure 1: The GenAI Contract template distributed to students through a Google Doc during Week 2.**

notes on patterns across students, surprising responses, contradictions and tensions, confusions, and interesting quotes. Our analysis was guided by our three goals: Did students engage meaningfully with the reflection prompts? What patterns emerged in how students aligned (or did not align) their GenAI use with their learning goals? And what aspects of the contract design seemed to support or hinder the development of intentional practice? As this is an experience report rather than a formal qualitative research study, we did not develop a formal codebook nor calculate inter-rater reliability. Instead, we report our observations on what worked and what didn't as instructor-researchers who read all students contracts. Statistics were generated by the first author.

## 4 Results and Takeaways

Student engagement with the contract process was relatively high across all three steps. Of the 217 enrolled students, 211 completed Steps 1-2 (97%), 200 completed Step 3 (92%), and 185 completed Step 4 (85%). While the quality of reflections varied, most responses contained sufficient detail for us to understand students' experiences and GenAI practices. The following subsections describe patterns in

students' responses. Since a single student could describe multiple experiences (e.g., finding the contract valuable while also abandoning guidelines under pressure) the percentages below are not mutually exclusive. It is also worth considering our analysis makes use of self-reported data, which has its own limitations.

### 4.1 What Worked: Helping Novices Form an Approach

More than half the students (58%, N=108/185) reported that the contract changed how they thought about GenAI Usage. Many of these students described arriving to the course without existing approaches to using AI tools. They had not given much thought to when, why, or how they engaged with GenAI in their learning process – they were simply using it as available. One student captured this shift: *"Having this contract process definitely changed how I think about using GenAI. It made me more aware of when and why I was using it, and helped me realize that setting personal boundaries with AI tools can actually make me a stronger learner."*

These students created concrete if-then rules tied to specific situations: time-based triggers (*"I will use GenAI after trying for 5*

**Step 3: Post-Midterm 1 Reflection**

- (1) Describe how you have actually used GenAI in this class so far (during homework, studying, debugging, learning concepts, etc.):
  - What types of situations prompted you to use it?
  - What generally did you ask it to do or help you with?
- (2) Looking back at the goal(s) you created in Step 1, do you think your GenAI usage helped or hurt you in working toward those goals? Has your goal(s) changed since the beginning of this course? Explain your reasoning.
- (3) What changes, if any, do you want to make to your GenAI usage guideline(s) you created in Step 2 going forward? You can use this format to describe your updated guidelines if you wish: I will *<the way you will use/not use GenAI>* when *<condition/constraint in which you will/won't use GenAI>* because *<reasoning that makes sense for you>*. OR I'm keeping my guidelines the same because *<explain how your current approach is working for your goals>*.

**Step 4: Post-Midterm 2 Reflection**

- (1) Before you reflect on the semester thus far, think about your recent GenAI usage:
  - In between Midterm 1 and Midterm 2, how often did you use GenAI for this class? [Not at all - Very frequently each week (10+ sessions)]
  - How does this compare to your usage earlier in the semester prior to Midterm 1? [Much less than before Midterm 1 - Much more than before Midterm 1]
- (2) You've now taken both midterms:
  - How confident do you feel in your programming skills compared to the start of the semester? [much less confident - much more confident]
  - Describe times when you chose to use GenAI, and separate times when you chose other resources (peers, lectures, office hours, etc.). Explain the reasoning behind your choices. What role do you think your GenAI usage played compared to other resources?
  - How do you feel about how you managed your GenAI use this semester? (Proud / Satisfied / Conflicted / Regretful / Other) Explain why.
- (3) Look back at the guidelines you created in Step 2 and/or 3.
  - Did you stick to them, modify them informally, or abandon them?
  - Were your guidelines effective strategies for your own learning goals you created in Step 1? Why or why not?
- (4) You might not have this contract reflection structure in future classes, but you'll keep making choices about GenAI.
  - Based on what you learned about yourself this semester, what's ONE insight about yourself as a learner that will shape how you approach GenAI tools going forward?
  - Complete this planning statement: "When learning programming in the future, I will use GenAI when \_\_\_\_\_ because \_\_\_\_\_. I will NOT use GenAI when \_\_\_\_\_ because \_\_\_\_\_."
- (5) This contract served as a self-regulation tool, as it asked you to plan (Steps 1-2), monitor (Step 3), and reflect (this step).
  - Did having this contract process change your GenAI usage or thinking about it? If so, how? If not, why not?
  - If you were re-designing this tool/process to help future students self-regulate their GenAI use when learning programming, what would you change?

**Figure 2: The GenAI Contract template distributed to students across two time points: Step 3 during Week 5, post-Midterm 1; and Step 4 during Week 11, post-Midterm 2. Students received Steps 3 and 4 as a text to copy/paste into their existing GenAI Contract Google Doc containing Steps 1 and 2.**

minutes", "I will ask GenAI ... if I am unable to solve the problem by myself within 30-40 minutes"), context-based conditions ("when office hours aren't available", "when I am working alone"), and action-based constraints ("I will use AI when I need a concept to be explained in very simple terms", "I will use GenAI, but carefully word the prompt so that GenAI only helps me without giving me the answer"). The written guidelines functioned as accountability structures. One student wrote, "Having a contract changed my brain when considering using AI, I'd always remember the rules I set for myself."

In contrast, some students (25%, N=46/185) arrived with established approaches to GenAI use. One student wrote, "I have always used it as a learning tool rather than a 'cheating' tool." Another noted: "I already knew what boundaries I wanted to set for myself for AI use." Other students had made deliberate choices to use GenAI more extensively. One explained: "I will use GenAI to find a solution ... because in the workforce, I will just be using AI anyway and correcting the work of the computer not the other way around so why not learn to use it." Another simply stated: "[I use GenAI] since I am only given a short amount of time. I rather get the points and not understand than not get the points and understand. ... I just want to pass the class and get my degree to get a good paying job." For these students – whether

they planned restricted or extensive use, whether they framed their approach through learning goals or pragmatic GPA/career goals – the contract documented existing practices rather than creating new ones. Thus, when asked if the contract changed their thinking, they reported "No."

This pattern aligns with implementation intention research. Gollwitzer and Sheeran's work demonstrates that if-then plans are most effective for initiating new behaviors, but less effective when strong habits (whether productive or unproductive) already exist [9]. Implementation intentions are traditionally tools for behavior initiation, not habit modification. Students without established practices used the contract to operationalize Zimmerman's SRL forethought phase [32] – translating their personal learning goals to concrete action plans. The iterative checkpoints created opportunities for performance monitoring and self-reflection. But students who had already systematized their GenAI use found the contract confirmatory rather than useful. As one put it: "I didn't really reflect or use this contract that much. I am generally pretty good about my AI usage and I don't think I needed this contract."

From a harm reduction perspective, this outcome supports the approach. Harm reduction emphasizes "meeting people where they

are" and supporting decision-making without imposing external control [21]. Students with established practices had already developed the decision-making capacity the intervention aimed to support, and students without those practices gained scaffolding they needed.

**Takeaway:** Self-regulation contracts work as helpful scaffolding for students initiating new behaviors, not for modifying existing habits. Instructors might see differential impact based on students' prior experiences with intentional GenAI use.

## 4.2 What Didn't Work: Forgetting about the Contract and Time Pressure

A smaller group of students (13%, N=23/185) reported forgetting about the contract between checkpoints. Students created guidelines in Week 2, returned to them in Week 5 after Midterm 1, and again in Week 11 after Midterm 2. In the weeks between, the students were not required to read or recall their guidelines. Students wrote: "I honestly haven't thought about this contract much.", and "It's very out of sight out of mind, once I finish it it's done and I never think about it again.". Some treated it as perfunctory: "honestly, this contract just seemed like another assignment to me. There wasn't ever really any impact." The contract lived in a Google Doc visited three times across the course, not where students made actual decisions about GenAI use. Students themselves identified this design flaw, suggesting "weekly check-ins" and ways to be "constantly reminded about" their guidelines.

More significantly, some students who valued their guidelines reported abandoning them under various pressures (11%, N=21/185). One wrote: "I use Gen AI when I am starting working on an assignment way too last minute, and feel pressured to just submit it on time." Another captured the tension: "There were absolutely situations where GenAI saved me a stupid amount of time and effort, like with debugging. At the same time, I do feel there were some instances where the "time-saving" argument was used in place of just critically thinking about what needs to be done."

**Takeaway:** Awareness and willpower alone may not be sufficient for sustained self-regulation under pressure. Future implementations could: (1) integrate guidelines into workflow where students make decisions, (2) increase reflection frequency to catch misalignment earlier, (3) more explicitly scaffold the revision process, and (4) consider assignment pacing/structures that reduce crunch time.

## 4.3 Student Suggestions for Redesigning the Contract

When asked for feedback on the contract process, students' suggestions clustered around three main ideas: increasing frequency and visibility, providing better scaffolding, and building in technical supports.

**More frequent touchpoints.** As discussed previously, students proposed bi-weekly or even weekly check-ins instead of midterm-based reflections: "I would include short weekly check-ins or reflection prompts instead of one big reflection" and "I would hold students even more accountable by maybe adding a reflection every week to talk about their use of AI weekly." Students recognized that infrequent reflections allowed problematic patterns to continue unchecked.

Although this may lead to intervention fatigue, more frequent touchpoints could address the forgetting problem by keeping guidelines active rather than allowing them to disappear for weeks at a time.

**Concrete examples and peer learning.** Students wanted more examples about what responsible use looked like. "I'd include more examples or short reflections from past students about how they used GenAI successfully," one wrote. Another suggested, "I might add a social sharing function (like in Piazza or other apps) that allows students to share their GenAI experience and skills with each other." Students also proposed showing "examples of good vs. bad usage" or "course-specific examples of when to use AI."

**Technical Scaffolding.** Students envisioned AI tools that enforced learning rather than affording shortcuts. One proposed "If I were to re-design a Gen-AI tool for learning programming I would make it ask you questions before giving you the correct code, almost like a rite of passage to make sure the person using it understands the concepts." Another proposed screen time limits: "Maybe introduce a feature that will allow students to put like a screen limit on their GenAI." These students recognized that the current process relied entirely on willpower without environmental design supporting intended behaviors, especially under time pressure.

## 5 Implications and Future Work

Aligning with implementation intention research [9], the contract supported the majority of our students' agency in better GenAI decision-making, but for a subset of our students, metacognitive awareness did not sustain intentional practice against the tool's tendency towards productivity-based use. Even students who valued their guidelines abandoned them because GenAI's frictionless alternatives created a low-cost pull away from effortful learning. Sustaining intentionality across hundreds of daily programming decisions requires support beyond self-regulation. This suggests that supporting genuine student agency may sometimes mean helping students recognize that intentional use is not possible for them right now – that the most agentic choice is deliberate non-use.

For instructors considering this approach, the contract successfully scaffolds students without established GenAI practices, making implicit choices explicit. Sustaining that intentionality requires much willpower and reflection. Students' suggestions in §4.3 may be helpful in reducing this burden. The contract helped students recognize the mismatch between their learning goals and GenAI's affordances, but resisting the tool's defaults still required self-regulation capacities not all students had developed. As such, instructors should stress that non-use is still a valid outcome of the intervention, not a failure to integrate the tool. We adopted a harm reduction philosophy to support student agency rather than impose rules, but our findings suggest this may require work from both sides: educators providing scaffolding and tools designed to work *with* students' goals rather than undermining them. The contract demonstrated that students can articulate these goals and reflect meaningfully on alignment. Our next step will test whether making these goals and reflections visible to the tools themselves can reduce the burden of self-regulation to sustainable levels. If tools cannot adapt to the context of students' learning needs, their usefulness in programming education may remain limited to students who least need the support [27].

## References

- [1] Areej Ali, Aayushi Hingle Collier, Umama Dewan, Nora McDonald, and Aditya Johri. 2025. Analysis of generative AI policies in computing course syllabi. In *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 1*. 18–24.
- [2] Albert Bandura. 1982. Self-efficacy mechanism in human agency. *American psychologist* 37, 2 (1982), 122.
- [3] Seth Bernstein, Ashfin Rahman, Nadia Sharifi, Ariunjargal Terbish, and Stephen MacNeil. 2025. Beyond the Benefits: A Systematic Review of the Harms and Consequences of Generative AI in Computing Education. In *Proceedings of the 25th Koli Calling International Conference on Computing Education Research*. 1–18.
- [4] Samuel Boguslawski, Rowan Deer, and Mark G Dawson. 2025. Programming education and learner motivation in the age of generative AI: student and educator perspectives. *Information and Learning Sciences* 126, 1/2 (2025), 91–109.
- [5] Kathryn Cunningham, Yike Qiao, Alex Feng, and Eleanor O'Rourke. 2022. Bringing "High-Level" Down to Earth: Gaining Clarity in Conversational Programmer Learning Goals. In *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education-Volume 1*. 551–557.
- [6] Paul Denny, Stephen MacNeil, Jaromir Savelka, Leo Porter, and Andrew Luxton-Reilly. 2024. Desirable characteristics for ai teaching assistants in programming education. In *Proceedings of the 2024 on Innovation and Technology in Computer Science Education V. 1*. 408–414.
- [7] Abbey Gandhi and Kasia Muldner. 2025. "That's Not the Way I Would Explain It": A Teacher-Researcher's Autoethnography of Learning to Program With ChatGPT. In *Proceedings of the 2025 ACM Conference on International Computing Education Research V. 1*. 227–239.
- [8] Peter M Gollwitzer. 1999. Implementation intentions: strong effects of simple plans. *American psychologist* 54, 7 (1999), 493.
- [9] Peter M Gollwitzer and Paschal Sheeran. 2006. Implementation intentions and goal achievement: A meta-analysis of effects and processes. *Advances in experimental social psychology* 38 (2006), 69–119.
- [10] Irene Hou, Sophia Mettill, Owen Man, Zhuo Li, Cynthia Zastudil, and Stephen MacNeil. 2024. The effects of generative ai on computing students' help-seeking preferences. In *Proceedings of the 26th australasian computing education conference*. 39–48.
- [11] Majeed Kazemitabaar, Runlong Ye, Xiaoning Wang, Austin Zachary Henley, Paul Denny, Michelle Craig, and Tovi Grossman. 2024. Codeaid: Evaluating a classroom deployment of an llm-based programming assistant that balances student and educator needs. In *Proceedings of the 2024 chi conference on human factors in computing systems*. 1–20.
- [12] Michael S Kirkpatrick. 2025. Challenging AI as Critical Thinking. In *International Conference on the Ethical and Social Impacts of ICT*. Springer, 239–251.
- [13] Shao-Heng Ko and Kristin Stephens-Martinez. 2025. Rethinking Computing Students' Help Resource Utilization through Sequentiality. *ACM Transactions on Computing Education* 25, 1 (2025), 1–34.
- [14] Shao-Heng Ko, Kristin Stephens-Martinez, Matthew Zahn, Yesenia Velasco, Lina Battestilli, and Sarah Heckman. 2025. Student perceptions of the help resource landscape. In *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 1*. 596–602.
- [15] Pamela K Kohler, Lisa E Manhart, and William E Lafferty. 2008. Abstinence-only and comprehensive sex education and the initiation of sexual activity and teen pregnancy. *Journal of adolescent Health* 42, 4 (2008), 344–351.
- [16] Sam Lau and Philip Guo. 2023. From "Ban it till we understand it" to "Resistance is futile": How university programming instructors plan to adapt as more students use AI code generation and explanation tools such as ChatGPT and GitHub Copilot. In *Proceedings of the 2023 ACM Conference on International Computing Education Research-Volume 1*. 106–121.
- [17] Mark Liffiton, Brad E Sheese, Jaromir Savelka, and Paul Denny. 2023. Codehelp: Using large language models with guardrails for scalable support in programming classes. In *Proceedings of the 23rd Koli Calling International Conference on Computing Education Research*. 1–11.
- [18] Rongxin Liu, Carter Zenke, Charlie Liu, Andrew Holmes, Patrick Thornton, and David J Malan. 2024. Teaching CS50 with AI: leveraging generative artificial intelligence in computer science education. In *Proceedings of the 55th ACM technical symposium on computer science education V. 1*. 750–756.
- [19] Dastyni Loksa, Lauren Margulieux, Brett A Becker, Michelle Craig, Paul Denny, Raymond Pettit, and James Prather. 2022. Metacognition and self-regulation in programming education: Theories and exemplars of use. *ACM Transactions on Computing Education (TOCE)* 22, 4 (2022), 1–31.
- [20] Lauren E Margulieux, James Prather, Brent N Reeves, Brett A Becker, Gozde Cetin Uzun, Dastyni Loksa, Juho Leinonen, and Paul Denny. 2024. Self-regulation, self-efficacy, and fear of failure interactions with how novices use llms to solve programming problems. In *Proceedings of the 2024 on Innovation and Technology in Computer Science Education V. 1*. 276–282.
- [21] G Alan Marlatt. 1996. Harm reduction: Come as you are. *Addictive behaviors* 21, 6 (1996), 779–788.
- [22] Michal A Naisteter and Justin A Sitron. 2010. Minimizing harm and maximizing pleasure: Considering the harm reduction paradigm for sexuality education. *American Journal of Sexuality Education* 5, 2 (2010), 101–115.
- [23] Aadarsh Padiyath. 2024. Do I Have a Say in This, or Has ChatGPT Already Decided for Me? *XRDS: Crossroads, The ACM Magazine for Students* 31, 1 (2024), 52–55.
- [24] Aadarsh Padiyath, Xinying Hou, Amy Pang, Diego Viramontes Vargas, Xingjian Gu, Tamara Nelson-Fromm, Zihan Wu, Mark Guzdial, and Barbara Ericson. 2024. Insights from Social Shaping Theory: The Appropriation of Large Language Models in an Undergraduate Programming Course. In *Proceedings of the 2024 ACM Conference on International Computing Education Research - Volume 1* (Melbourne, VIC, Australia) (ICER '24). Association for Computing Machinery, New York, NY, USA, 114–130. doi:10.1145/3632620.3671098
- [25] Jennifer Parham, Leo Gugerty, and DE Stevenson. 2010. Empirical evidence for the existence and uses of metacognition in computer science problem solving. In *Proceedings of the 41st ACM technical symposium on Computer science education*. 416–420.
- [26] James Prather, Paul Denny, Juho Leinonen, Brett A Becker, Ibrahim Albluwi, Michelle Craig, Hieke Keuning, Natalie Kiesler, Tobias Kohn, Andrew Luxton-Reilly, et al. 2023. The robots are here: Navigating the generative ai revolution in computing education. In *Proceedings of the 2023 working group reports on innovation and technology in computer science education*. 108–159.
- [27] James Prather, Brent N Reeves, Juho Leinonen, Stephen MacNeil, Arisoa S Randrianasolo, Brett A Becker, Bailey Kimmel, Jared Wright, and Ben Briggs. 2024. The widening gap: The benefits and harms of generative ai for novice programmers. In *Proceedings of the 2024 ACM Conference on International Computing Education Research-Volume 1*. 469–486.
- [28] Annapurna Vadaparty, Daniel Zingaro, David H Smith IV, Mounika Padala, Christine Alvarado, Jamie Gorson Benario, and Leo Porter. 2024. Csl-llm: Integrating llms into cs1 instruction. In *Proceedings of the 2024 on Innovation and Technology in Computer Science Education v. 1*. 297–303.
- [29] Zilin Wang, Di Zou, Ruofei Zhang, Lap-Kei Lee, Haoran Xie, and Fu Lee Wang. 2025. ChatGPT-enhanced self-regulated learning in programming education: impacts on motivation, self-efficacy, and learning outcomes. *Interactive Learning Environments* (2025), 1–26.
- [30] Ramazan Yilmaz and Fatma Gizem Karaoglan Yilmaz. 2023. Augmented intelligence in programming learning: Examining student views on the use of ChatGPT for programming learning. *Computers in Human Behavior: Artificial Humans* 1, 2 (2023), 100005.
- [31] Cynthia Zastudil, Magdalena Rogalska, Christine Kapp, Jennifer Vaughn, and Stephen MacNeil. 2023. Generative ai in computing education: Perspectives of students and instructors. In *2023 IEEE Frontiers in Education Conference (FIE)*. IEEE, 1–9.
- [32] Barry J Zimmerman. 2002. Becoming a self-regulated learner: An overview. *Theory into practice* 41, 2 (2002), 64–70.