

# Who Will Become the Next Senior? How Generative AI Erodes the Development Pathway in Software Engineering

Anonymous submission

## Abstract

Generative AI (GenAI) is reshaping software engineering, raising concerns about how the development pathway through which juniors become seniors is being eroded. While macro statistics show a decline in junior hiring and controlled studies demonstrate the effects of AI on individual task performance, the *mechanisms* through which GenAI reshapes early-career *development* in real organizational and educational contexts have not been thoroughly examined. Through 14 semi-structured interviews with juniors at the threshold of entering software engineering and senior software engineers in South Korea, analyzed using Reflexive Thematic Analysis, we reveal a foundational pattern of Absorption—GenAI redirects entry-level work into senior-AI workflows—and three consequences: (1) juniors losing the productive struggle through which expertise once developed; (2) the structural reproduction of this loss through collective normalization of GenAI use in university classrooms; and (3) the perceptual asymmetry between seniors and juniors that prevents either side from correcting these dynamics on their own. By extending learning theory and situated cognition to organizational and institutional scales, we argue that GenAI absorbs not just specific categories of tasks but also the pathway through which the next generation of seniors is formed. Preserving this pathway will require deliberate institutional design across classrooms, workplaces, and the evaluation criteria for juniors.

## Introduction

The advent of generative AI (GenAI) has transformed software engineering at unprecedented speed. Tools like GitHub Copilot, ChatGPT, and Claude now perform substantial portions of code generation, debugging, and documentation that were frequently assigned to entry-level engineers. Concurrently, junior developer hiring has contracted sharply. In the United States, entry-level postings at the fifteen largest U.S. technology firms fell 25% between 2023 and 2024 (Brynjolfsson, Chandar, and Chen 2025), programmer employment in the United States declined 27.5% between 2023 and 2025 (Rak 2025). In South Korea, major technology companies have largely suspended open recruitment of junior software engineers, shifting to experience-based ad hoc hiring (Lee 2025).

At the individual level, GenAI clearly accelerates developer productivity (Brynjolfsson, Li, and Raymond 2025;

Peng et al. 2023; Cui et al. 2026). At the collective level, however, this raises one underexamined question: if entry-level work is no longer reaching junior engineers, who will become the next generation of seniors? Nevertheless, these hiring trends are commonly read as job displacement—efficient substitution of human labor with automated tools (i.e., GenAI), accompanied by corresponding adjustments in headcount. We do not dispute that this displacement is occurring. We would argue, however, that the tasks currently being absorbed by GenAI, such as minor bug fixes, routine implementation, and debugging, are not simply *entry-level tasks* that organizations can eliminate without further issues and concerns. What GenAI is absorbing is the *developmental scaffolding* through which juniors grow into seniors: the early stages of an apprenticeship-based development pathway, where juniors learn through real-world experience alongside more experienced developers, and expertise is passed down from one generation to the next (Lave and Wenger 1991). When this scaffolding is absorbed, what disappears is not only entry-level employment but also the very *mechanism* through which the next generation of seniors, i.e., skilled professionals, is formed.

We examine these dynamics through the case of South Korea, where this mechanism may operate strongly. Korean workers have adopted GenAI faster than nearly any other national context—work-related AI usage reaches 51.8%, almost double the U.S. rate (Bank of Korea 2025b). At the same time, open recruitment has collapsed, and Korea's labor market structures displace adjustment costs onto labor market entrants rather than incumbent workers (Bank of Korea 2025a). Therefore, rather than serving as a representative national case study, we note that South Korea functions as an important case in which these dynamics emerge earlier and more clearly.

Building on the macro-level hiring decline identified in numerous previous quantitative analyses, the contribution of our study lies not in confirming this phenomenon but in identifying its underlying mechanism and the factors contributing to its persistence. We address three research questions through semi-structured interviews with 14 software engineers—eight juniors at the threshold of workforce entry, and six seniors with at least six years of industry experience:

- **(RQ1)** What developmental capacities are juniors losing as GenAI restructures early-career work across work-

places and classrooms?

- (RQ2) Through what structural mechanism is junior productive struggle obstructed?
- (RQ3) Why do these dynamics persist without self-correction by individuals or organizations?

Our contribution is threefold. Empirically, we identify a foundational pattern of Absorption; GenAI redirects entry-level work into senior-AI workflows. From this pattern, three consequences follow. First, juniors are losing the productive struggle through which expertise was once developed (Kapur 2016; Bjork 1994). Second, this loss is reproduced as a structural condition through the collective normalization of GenAI use in university classrooms, foreclosing individual choice. Third, the perceptual asymmetry between seniors and juniors (Lave 1988; Brown, Collins, and Duguid 1989) blocks coordinated self-correction. Theoretically, we extend learning theory and situated cognition from the individual learner scale to the organizational and institutional scales where these dynamics actually occur. Normatively, we argue that the conditions for the growth of juniors cannot be restored through individual effort alone; they must be re-established through deliberate institutional design.

## Background and Related Work

### GenAI and the Changing Labor Market

The rapid proliferation of GenAI tools, such as ChatGPT, Claude, and GitHub Copilot, has prompted growing concern about their impact on labor markets, particularly for early-career workers. Global reports consistently show a structural shift in workforce demand, including sharp declines in U.S. entry-level technology hiring (Brynjolfsson, Chandar, and Chen 2025), as well as in overall programmer employment (Rak 2025). These trends do not appear to be merely the result of economic cycles, and employers are increasingly redefining what they expect from new hires. A large-scale analysis of approximately 11 million UK job postings from 2018 to 2024 found that demand for AI roles grew 21% while university degree requirements for those same roles declined by 15%, suggesting a structural shift toward skill-based hiring (Bone, Ehlinger, and Stephany 2025). Moreover, recent evidence further suggests that GenAI adoption is shifting work away from entry-level tasks across knowledge-intensive occupations, including software engineering, effectively narrowing the bottom stages of internal career ladders (Hosseini Maasoum and Lichtinger 2025).

South Korea offers a particularly clear lens on these dynamics. Alongside its unusually high workplace GenAI adoption (Bank of Korea 2025b), South Korea's open recruitment—in which major companies had long hired new software engineers in cohorts—has largely collapsed: leading technology companies have shifted to experience-based rolling hiring and halted open junior recruitment (Lee 2025), and IT-sector job postings decreased by 43% between 2023 and 2024, with entry-level positions accounting for only 4.4% of postings (Lee 2025). The Bank of Korea has framed this as a form of seniority-biased technological change, in which wage stickiness combined with patterns of seniority-biased adjustment concentrates costs on new entrants in the

labor market rather than existing workers (Bank of Korea 2025a).

However, it is important to note that establishing a causal link between GenAI adoption and the decline in junior hiring remains challenging. Empirical analyses suggest that the employment conditions for AI-exposed entry-level jobs had already begun to weaken in early 2022—months before ChatGPT's public release—suggesting that broader economic shifts, post-pandemic corrections, and monetary tightening were key contributing factors (Brynjolfsson, Chandar, and Chen 2025). Nevertheless, the convergence of labor market shifts and the rapid adoption of GenAI is reshaping junior roles and the competencies they require, regardless of how much of the broader hiring decline can be attributed to GenAI.

While existing research has documented these macro-level shifts through labor market statistics and employer surveys, such analyses leave critical questions unanswered; how do senior and junior engineers experience the changing role of early-career work, and where do their perceptions diverge? Through what mechanisms is GenAI reshaping early-career development in software engineering, and why do these mechanisms persist, preventing juniors who have completed their education from entering the workforce? This study addresses these gaps through qualitative interviews with *both* groups.

### GenAI and the Changing Nature of the Workplace

In addition to a broader impact on labor markets, GenAI is transforming the nature of work across many fields (Eloundou et al. 2024). In the realm of software engineering in particular, as software development practices change, there is growing interest in how developers can integrate GenAI into their day-to-day workflows (Liang, Yang, and Myers 2024; Barke, James, and Polikarpova 2023; Vaithilingam, Zhang, and Glassman 2022).

In terms of productivity, studies consistently find that AI-powered coding assistants accelerate task completion for individual developers (Brynjolfsson, Li, and Raymond 2025; Peng et al. 2023; Cui et al. 2026). Longitudinal research on a software development organization (Xiao et al. 2025) further suggests that while AI tools have not resolved persistent team-level problems such as accountability for results or fragile communication, they have helped reshape team norms: by accelerating individual tasks such as coding and documentation, they establish efficiency as a baseline expectation and the responsible use of AI as a marker of professionalism. Moreover, on the industry side, expectations of what new hires should bring are concurrently shifting (World Economic Forum 2025). Where previous cohorts of junior engineers were expected to bring foundational coding skills and develop further on the job, employers increasingly expect new entrants to arrive with AI literacy as a baseline competency (Maslej et al. 2025). In South Korea, a Delphi survey of software industry experts (Shin et al. 2025) found that creative problem-solving ability, AI utilization capability, and domain understanding have emerged as the key competencies employers now seek, displacing the technical coding skills that were previously prioritized. As a result, the

nature of entry-level positions is being redefined, while the barriers to entry are becoming even higher.

On the other hand, a recurring concern is deskilling (Braverman 1998), i.e., the erosion of human expertise, which occurs when automated systems absorb tasks that previously required deliberate practice and skill development. Recent research on GenAI revisits this concern at a new scale. People working in various knowledge industries anticipate that GenAI, even when framed as a tool for performing ‘menial work under human review,’ may exacerbate the existing forces of deskilling, dehumanization, and disconnection (Woodruff et al. 2024). This concern is particularly acute for novice and early-career practitioners: a diary study of young professionals found that ChatGPT use produced a diminished sense of ownership and challenge, captured in the question that titles the paper —“If the machine is as good as me, then what use am I?” (Kobiella et al. 2024).

Despite this growing body of work, most existing studies focus on individual productivity and team-level dynamics, and little empirical attention has been paid to how the introduction of GenAI is reshaping the *development processes* through which junior engineers develop expertise in real organizational contexts—a gap this study seeks to address.

Additionally, there is research analyzing the adoption of GenAI from a student perspective. A study of novice programmers, where the participants were students, found that while GenAI tools provided short-term productivity gains for some, they simultaneously widened the gap between those who struggled and those who progressed quickly by exacerbating pre-existing metacognitive difficulties and creating an illusion of competence—those who struggled often believed they had performed better than they had (Prather et al. 2024). Qualitative research on software engineering students documents recurring patterns of GenAI dependency and uncertainty about skill formation (Choudhuri et al. 2025). However, such studies have focused on student experiences in isolation, without incorporating the perspectives of industry seniors who can evaluate and hire these students.

Overall, existing research has largely examined this transition from either the educational side or the workplace side, not examining how those directly experiencing this structural change—such as junior engineers entering the workforce and senior engineers already within organizations—actually perceive it. This study contributes such an account. Specifically, we begin with the framework of the “skills gap between education and industry.”

The most recent research that is most similar to our research, which covers both juniors and seniors, is done by Feng, Yun, and Wang (2026): a recent mixed-methods study of agentic AI-mediated software engineering found that junior developers employed in industry oscillate between AI over-reliance and cautious avoidance and that agency in such work is shaped more by organizational policy than by individual preference. However, a key distinction is that, in our study, the term ‘junior’ refers to individuals preparing to enter the workforce, whereas in their definition, it refers to those who have already entered it. By defining ‘junior’ as ours, we can analyze the adoption of GenAI from *both* the workplace and educational perspectives.

## Theoretical Foundations of the Development Pathway

This developmental pathway has been theorized through the *situated cognition* tradition. Lave (1988) argues that cognition cannot be isolated from the social settings and practices in which it takes place. Brown, Collins, and Duguid (1989) extend this view: practitioners’ understanding of what counts as a problem or solution emerges from their direct engagement with practice, not from abstract reasoning that precedes action. Closely related work by Lave and Wenger (1991) develops the concept of *Legitimate Peripheral Participation* (LPP): an apprenticeship-like trajectory in which newcomers join a community as recognized members and learn by gradually moving from peripheral, lower-stakes tasks to more central roles—competence develops through participation itself, rather than through instruction. From this perspective, the developmental value of early-career work lies not only in the tasks themselves, but in the social and organizational context in which they are performed (Begel and Simon 2008; LaToza, Venolia, and DeLine 2006; Dagenais et al. 2010; Wenger 1999)—learning occurs through doing, observation, and incremental responsibility. The removal or reduction of lower-stakes entry-level tasks therefore represents not merely a change in workflow; it fundamentally disrupts the very structure through which expertise is reproduced. This mechanism has been documented in other professional domains as well (Beane 2019, 2024).

Research on the cognitive mechanisms of skill acquisition further explains *why* removing peripheral tasks is consequential. Kapur’s theory of *Productive Failure* posits that learning and performance are not commensurable: conditions that maximize short-term performance do not necessarily maximize long-term learning (Kapur 2016). Learners who first struggle with and fail to solve unfamiliar problems before receiving canonical instruction were found to have significantly greater conceptual understanding and application skills than those who receive instruction initially. In other words, generating imperfect or failed solutions enables learners to recognize critical features of knowledge that structured instruction alone does not allow. Bjork’s related concept of *desirable difficulties* similarly identifies conditions that impede short-term performance but enhance long-term retention and transfer (Bjork 1994). Crucially, both frameworks converge on a counterintuitive insight: smooth and successful performance during the early stages of learning is often a signal of superficial encoding rather than deep understanding. Taken together, these frameworks suggest that the junior-to-senior developmental pathway depends on a structured, incremental exposure to authentic organizational tasks, beginning with lower-stakes, error-prone work and gradually increasing in complexity and responsibility. This pathway has served as the primary mechanism through which organizations reproduce expertise across generations.

The question raised by the present study is whether GenAI adoption—by absorbing the lower stages of this ladder—is disrupting this process in ways that existing organizational theory has not yet fully accounted for. Our study examines




Figure 1: Temporal positioning of study participants. Juniors straddle the periods before and after the adoption of GenAI during their undergraduate years; seniors entered the industry before its adoption and continue to work in the field.

this question through the lived experiences of both juniors preparing to enter the workforce and seniors currently working in the field, in South Korea, offering an empirical analysis of how these theoretical disruptions manifest in practice.

## Method

### Study Design

We conducted a qualitative study using semi-structured interviews to investigate how the emergence of GenAI tools has reshaped skill expectations and career development pathways for early-career software engineers in South Korea. This study is detailed by three research questions (RQs), which are described in the Introduction section. Semi-structured interviews were chosen to enable participants to share their experiences while allowing the researchers to explore emerging themes in depth.

### Participants

We recruited 14 participants through purposive sampling drawn from the authors’ personal and professional networks. The participants were divided into two groups: eight juniors and six seniors. Participants were selected based on their temporal positioning relative to GenAI adoption, illustrated in Figure 1 and detailed below for each group. In accordance with the principle of information power (Malterud, Siersma, and Guassora 2016), we judged sample adequacy based on the specificity of our study population, the focus of our research objectives, and the analytic depth pursued.

*Junior participants* (J1–J8) were defined as individuals within two years of completing an undergraduate degree, representing those at the threshold of entering the professional workforce. All junior participants were undergraduate students or recent graduates majoring in Electrical and Computer Engineering (ECE) or Computer Science (CS), with career trajectories oriented toward software engineering. Notably, all junior participants experienced both the pre- and post-GenAI period during their undergraduate education, having taken courses both before and after the widespread availability of tools such as ChatGPT.

*Senior participants* (S1–S6) were defined as professionals with at least six years of industry experience, drawn from diverse organizational contexts: two from large enterprises, two from mid-sized companies, and two from startups, spanning domains including game development, e-commerce,

ID	Years	Organization Size	Domain	Major
J1–J8	(student)	N/A	N/A	ECE/CS
S1	9–11	Large Enterprise	Electronics	ECE
S2	9–11	Large Enterprise	Game	ECE
S3	6–7	Startup	E-commerce	ECE
S4	11–12	Startup	Robotics	ECE
S5	9–11	Mid-sized Company	Game	ECE
S6	6–7	Mid-sized Company	Re-commerce	CS

Table 1: Participant Overview. Junior participants (J1–J8) are current students or recent graduates; they have not yet entered the workforce, so Organization Size and Domain do not apply. Years indicates years of industry experience.

robotics, and others. We purposefully designed this diversity to further observe how the adoption of GenAI varies depending on organizational size and technical domain. Although senior participants were recruited from diverse organizational contexts, all held undergraduate degrees in ECE or CS, sharing the same educational foundation as the junior group, and are engaged in software engineering-related work in their respective fields. Table 1 provides an overview of participant demographics.

### Data Collection

Interviews were conducted between December 2025 and February 2026, in Korean, the native language of both the participants and the authors. Each interview lasted approximately 50 minutes and audio recordings were made for all interviews with participants’ consent. Ethical approval was granted by the Institutional Review Board (IRB) of the authors’ institution<sup>1</sup>. All recordings were transcribed using an automated speech recognition tool (CLOVA Note<sup>2</sup>), followed by manual correction by the authors. Since the authors are native Korean speakers, data analysis was conducted using the original Korean transcripts; quotes appearing in this paper are English translations of selected Korean excerpts, reviewed by the authors for accuracy and idiomatic fidelity.

### Interview Structure

Separate semi-structured interview guides were developed for junior and senior participants, each tailored to their respective roles and experiences. The junior guide focused on perceived skill gaps between university education and industry expectations, use of GenAI tools in academic and work contexts, and expectations about career entry. The senior guide focused on changes in task delegation and mentoring practices since the adoption of GenAI, perceptions of junior competencies, and hiring criteria changes. Both guides were designed to elicit participants’ perspectives on the evolving relationship between GenAI capabilities and human skill development. Detailed interview question scripts for junior and senior are described in the supplementary materials.

Although the two guides differed in group-specific framing, they were designed around shared anchor topics—

<sup>1</sup>Details are omitted to ensure anonymity.

<sup>2</sup><https://clovanote.naver.com/>

the changing nature of entry-level work, the conditions of skill formation, and expectations about junior–senior development—to which each group could respond from its respective vantage point. This shared topical structure enabled cross-group analysis while maintaining the specificity of the questions asked in each group. During analysis (described below), themes were constructed only when supported by accounts from both groups, except when asymmetry between groups itself constituted a meaningful analytic finding, as in our analysis of misaligned perspectives.

## Data Analysis

All transcripts were anonymised prior to analysis by replacing participant names and organisational identifiers with coded labels (e.g., J1, S1, and Enterprise A). We analyzed the data using Reflexive Thematic Analysis (RTA) (Braun and Clarke 2006, 2019), an inductive approach that is well-suited to exploratory qualitative research seeking to understand participants’ perspectives and experiences. We note here that in line with the theoretical orientation of RTA, we did not consider data saturation as a criterion for determining sample size (Braun and Clarke 2021).

Analysis proceeded through six phases: (1) familiarization with the data through repeated reading of all transcripts; (2) generating initial codes by annotating meaningful segments across the dataset; (3) constructing candidate themes by grouping related codes; (4) reviewing themes against the coded data and full dataset; (5) defining and naming themes; and (6) producing the final analysis.

Throughout our analysis, an audit trail was maintained to document how initial codes developed into final themes, supporting the transparency of the analytical process (Lincoln and Guba 1985). Disconfirming cases were actively sought to ensure that the themes reflected a wide range of participants’ perspectives, rather than only confirming instances.

Positionality, ethical considerations, and adverse impact statement are discussed in detail in the respective sections at the end of this paper.

## Findings

Having introduced the methods of our study, we present our analysis under four themes. The first—Absorption of Junior Opportunities into Senior Workflows—establishes the foundational pattern: GenAI expands senior capacity such that entry-level work no longer reaches juniors. Three subsequent themes address our research questions—*Erosion of Competence through Lost Failure Experiences* (RQ1) documents what juniors are losing across workplaces and classrooms; *Collective Pressure as a Barrier to Productive Struggle* (RQ2) examines why this loss is structurally reproduced rather than individually avoidable; and *Misaligned Perspectives Across Experience Levels* (RQ3) explains why these dynamics resist self-correction. In each section, we present quotes alongside the corresponding analysis results to illustrate our claims; unless otherwise noted, the patterns these quotes represent were observed across multiple participants.

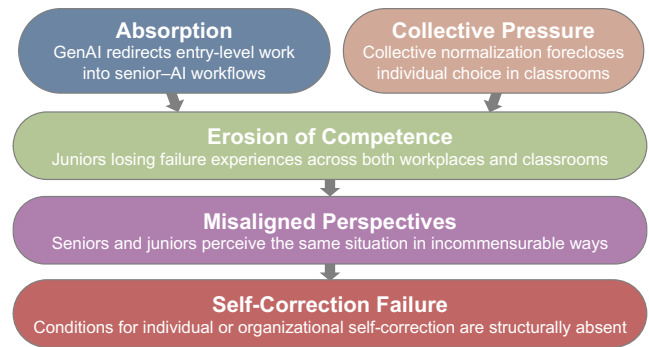


Figure 2: Overview of the paper’s argument. Absorption and collective pressure (RQ2) produce the erosion of competence (RQ1); misaligned perspectives prevent self-correction (RQ3).

### Absorption of Junior Opportunities into Senior Workflows

First, at the beginning of the interview, we examined how the emergence and adoption of GenAI has affected senior participants in their workplaces. A foundational pattern observed across senior participants was that GenAI tools had dramatically expanded their individual capacity, enabling them to handle tasks that would previously have been delegated to others, particularly entry-level engineers. This had a direct structural consequence: the entry-level work that had historically provided *learning opportunities* for entry-level engineers was no longer reaching them.

S4, a start-up founder with over twelve years of experience, described this shift in operational perspective. Having significantly reduced his junior headcount over the last year, he observed that their departure had had no measurable impact on productivity. As he put it,

“We let a lot of junior engineers go in the second half of last year. But there was no impact whatsoever.”

Following his remark that it was unavoidable in order to improve the company’s productivity, when asked what junior engineers could contribute that AI tools could not, S4 was blunt:

“What can a junior engineer do better than a 100,000-won Claude subscription? I don’t think there’s anything.”

This perspective—namely, junior employees were evaluated against the economic calculus regarding the possibility of being replaced by AI—was not unique to S4. Among senior participants (S1–S6), it was consistently observed that AI tools had effectively absorbed more routine, lower-complexity tasks, e.g., code generation, initial debugging, and documentation. As S5 observed,

“These days, things are being replaced so quickly that juniors are increasingly not being given the opportunity to do those things”.

On the other hand, one account made this pattern more complicated. S1, who worked at a large enterprise, explained

how, as a senior, he had become able to handle most of the tasks previously done by juniors, but at the same time, he noted the following:

“from the junior’s perspective, the work they do and what they learn hasn’t really changed much. The only difference is that AI now exists.”

This was the only such observation across our senior interviews. In contrast to the absorption pattern described by other seniors, S1’s account suggested that juniors in his organization continued to do similar work, with AI added as a supplement rather than as a replacement. Rather than reading it as a simple outlier, we consider it as indicating boundary conditions for the absorption mechanism. The seniors who described absorption most directly—notably S4 in a startup, where “we let a lot of juniors go in the second half of last year”—operated in settings where task allocation is filtered primarily through efficiency calculus. S1’s large enterprise context represents the opposite pole, where institutional structures such as formalized onboarding pipelines and structured training programs determine the allocation of work, regardless of immediate productivity considerations. While our sample cannot confirm this organizational difference as the underlying mechanism, the contrast between S4’s startup and S1’s large enterprise suggests a tentative hypothesis: absorption operates most freely where institutional structures for junior development are weak or implicit. We return to this point through S6’s onboarding design example in our Discussion section.

Based on the observations of the *absorption*, we continued to interview seniors and, by analysing their responses alongside those of juniors, we identified the main themes that will be discussed in the next section. The following two themes examine what juniors lose when absorption proceeds without discernment, and how such losses manifest across the workplace and educational sites that engage in their skill development.

### Erosion of Competence through Lost Failure Experiences

A central finding for both junior and senior participants was that, as they began using GenAI whether in school or at work, *failure* no longer served as a source of learning. Specifically, as GenAI absorbed the lower-complexity tasks that once characterized early-career workflows, the *productive struggle* embedded in those tasks also disappeared.

All junior participants (J1–J8) noted that, compared to the pre-GenAI period, they can now achieve good results easily without the struggle they once experienced when completing university assignments or preparing for exams, and described what remained in its place as follows: J6 said “I don’t know what not to do. I’ve only seen the good examples, so I don’t know what the bad ones look like,” and J1 described a more fundamental metacognitive issue, saying “I don’t know what I don’t know.” J2 elaborated on the same dynamics, saying:

“If AI gives you the answer directly, the time you spend struggling gets shorter—so when you’re faced

with a problem, you might not even recall that a relevant concept exists.”

As such, there was a consistent pattern among the participants regarding what emerged in place of *failure: output without understanding*. Anthony (2021) documented this phenomenon in adjacent professions: when novices use analytical tools without interpreting their outputs—what he terms ‘black boxing’—they cannot trace how outputs are generated, and therefore cannot learn from them. J8 provided the most illustrative example, recalling a project presentation at her internship where she used unverified, AI-generated content:

“There were things I was presenting that I didn’t understand myself. I was thinking, ‘Is this even right?’— and I hadn’t even checked, because there was no time.”

Taken together, these accounts reflect what Kapur’s theory of *Productive Failure* would suggest: when learners are shielded from the imperfect, effortful attempts that they must endure to achieve mastery, they not only lose knowledge, but also the ability to recognize the boundaries of their own competence (Kapur 2016). Similarly, Bjork’s concept of *desirable difficulties* suggests that the conditions that appear to slow progress, such as struggling, failing, and revising, produce durable and transferable learning (Bjork 1994). GenAI, by eliminating these conditions, yields what J4 described as a distinct difference in his experience of two university courses covering the same subject.

“My grades came out the same. But the growth I felt—I definitely studied much more in the machine learning course. The deep learning course just felt like I got better at using GPT.”

The students, the academic field, and the evaluation criteria are all the same; the only difference is whether the assignment can be completed using GenAI. The grade-equivalent outcomes are what make the case analytically significant: they show that the institutional signal of learning (the grade) can remain stable while what is actually learned diverges sharply. In one course, the work produced durable domain knowledge; in the other, it produced fluency in just operating a tool, i.e., GenAI. This case demonstrates the mechanism Kapur and Bjork would predict under nearly controlled conditions and reveals that the current assessment structure in education fails to identify it.

Even as growth through failure is fading as GenAI appears in the educational context, those failures remain an important element based on seniors’ experiences and perspectives. S6 explained what matters for professional development:

“The reason someone is called a senior is that they know how work should *not* be done, but now juniors can’t learn from the bad examples.”

Notably, junior and senior participants, standing at opposite ends of the experience spectrum, used nearly identical phrasing, “knowing what should not be done”, to describe the missing competency. Moreover, this perspective redefines *seniority* not merely as the accumulation of correct solutions but also as a process of acknowledging failure

through effort and challenging experience. S3 further confirmed this from personal experience, describing his own decision-making skills nowadays as something built through years of incremental, hard work from the time he was in his early career:

“Going through things one by one, making mistakes—the distress of having been wrong, accumulated over a very long period of time.”

Moreover, what people can do in workplaces fundamentally is related to the skills that can be formulated by this failure process. S5, a senior engineer at a mid-sized company, described this point as follows:

“The cost of code generation is so cheap by using GenAI, but the cost of code verification is so expensive.”

Such verification—knowing what is wrong, what does not work—is the human skill that failure-based learning produces. Yet juniors, by being shielded from failures, are no longer developing it. GenAI accelerates output production while eroding the very experiences through which juniors would learn to evaluate that output.

### **Collective Pressure as a Barrier to Productive Struggle**

Whereas the previous section describes the loss of failure experiences eroding junior developers’ competence, this section examines how those experiences are being eliminated in the first place. What we observed is that this is not due to individual choice, but rather stems from a structural dynamic that operates within the educational environment itself.

A consistent pattern observed across junior participants was that the use of GenAI in university coursework is no longer a matter of personal choice. Rather, it had become the norm for most students, meaning that individual students found it difficult to resist using GenAI, regardless of their own intentions. While this collective normalization reflects South Korea’s broader context of rapid AI adoption (Bank of Korea 2025b), the specific dynamics observed in the university classrooms of the junior participants are best traced through their own accounts.

The account provided by J4 was the most direct illustration of this dynamic: having initially chosen to work through assignments independently, he described the moment when this became impossible.

“Everyone except me was using GPT, so everyone except me was getting nearly perfect scores. I no longer had the freedom to work through assignments on my own, making mistakes as I went—so in the end, I had no choice but to use it too.”

The pressure J4 describes is structural. In other words, the grading curve itself, as well as the fact that peers are using GenAI, has become the mechanism of coercion. J5 expressed the same norm: “Not using it is just stupid.” J7 framed it as anxiety: “If I don’t use it, I feel like I’m falling behind on my own.” Accounts from all junior participants (J1–J8), including these three accounts, represent different

but convergent responses to the same underlying collective pressure, leaving little room for individual resistance.

Crucially, this pressure operated not only among students but also at the institutional level in their universities. J3 noted that faculty had effectively overlooked this practice:

“The professors know students are using ChatGPT anyway... I don’t think there was a single professor who told us not to use it.”

J6 explained the institutional logic behind this silence as follows:

“The courses that AI can already do well are the ones where the original learning objectives aren’t being met anymore. They should ban it—but you can’t stop everyone from using it.”

Consequently, a structural gap has emerged between the decisions about using GenAI in learning and what students actually are doing, leaving neither students nor institutions able to maintain the conditions necessary for *productive struggle*, which was introduced in our previous finding, thereby hindering the growth of juniors.

What makes this observation consequential for the skill development of juniors is precisely identified by the accounts of J5:

“The process has to be somewhat inefficient—I think that’s just what learning requires. But in a society that values achievement, falling behind by not using it seems certain.”

Albeit without using the theoretical term, this describes the *desirable difficulty* proposed by Bjork (1994). J5 recognized this—and used AI anyway, because the educational structure left no viable alternative.

Not all junior participants felt this pressure to the same extent. J2 described his ongoing personal resistance, stating, “I keep trying not to take those shortcuts,” while J6 identified open-ended, self-directed projects as spaces where independent thinking was still possible. These accounts suggest that structural pressure varied across learning situations. However, J2 also acknowledged the limits of such resistance:

“I’m trying not to, but I’m so pressed for time that I end up doing it anyway.”

Individual intention, in other words, was undermined by structural constraint, suggesting that collective pressure exerts a stronger influence than personal preference.

### **Misaligned Perspectives Across Experience Levels**

The preceding themes documented how this developmental pathway is being structurally obstructed. However, a critical question remains: why do these mechanisms persist without being resolved? Senior participants were not indifferent to the difficulties of juniors, and many expressed concern. This section addresses this question by examining how junior and senior participants perceive the same phenomenon in different terms and how that difference becomes a barrier to resolution.

A consistent pattern across senior participants was that the current situation appeared manageable. However, underlying this perception was an implicit assumption. Namely, they

assumed that juniors who had accumulated direct, hands-on experience would develop into competent professionals, combining that experience with AI fluency, much as seniors themselves had.

S5 articulated this most explicitly:

“The fact that a junior lacks domain knowledge—I actually think he can easily fill that gap using AI.”

Subsequently, in the hiring process, he indicated that the most important thing would be direct experiential learning:

“Even if there are no results to show, I think I’d give a lot of credit to someone who has directly worked with real data.”

What this framing does not account for is whether the conditions it assumes are actually in place where juniors are educated. As documented in the previous two sections, the educational contexts in which juniors are currently developing are ones in which GenAI has become the default mode of completing assignments and coursework—precisely the contexts in which the direct, effortful experience that seniors value would *otherwise* have been accumulated.

As noted earlier in our discussion of Absorption, S5 himself acknowledged that juniors no longer encounter the very experiences he treats as a prerequisite, observing that “the problem is that we can’t wait for that time anymore.” Yet this recognition coexists with the earlier assumption that juniors with direct, hands-on experience will emerge. S5 does not appear to recognize that these two statements are in tension. This means that, following the emergence of GenAI, the skills that seniors expect from juniors are extremely difficult to acquire given the pace and trends of society—including the educational environment. Seniors themselves are aware of this fact yet continue to hold those expectations.

Juniors experienced this same reality from a position that made its consequences *immediate*. While seniors viewed the situation as difficult but manageable, juniors experienced it as a process in which the opportunity to make knowledge their own had been denied. J2 described this directly:

“The less time you invest, the less you get back—it feels like the core is hollowed out.”

J4 articulated the same experience in terms of ownership:

“It becomes a skill you can only demonstrate when GPT is there. In the end, that’s not really mine—it’s not something I can draw on when I’m completely on my own.”

From the senior side, the same reality appeared in fundamentally different terms. S4 articulated this:

“We are in a position to evaluate AI, whether this output is right or wrong, because of twenty years of accumulated experience. The next generation won’t be in that position. Therefore, we are fine.”

This divergence is not simply a matter of differing opinions or experience levels. Rather, it reflects a related strand of the situated cognition tradition introduced earlier (Lave 1988; Brown, Collins, and Duguid 1989)—the idea that one’s position within a practice shapes what is perceptible and what remains invisible. Specifically, seniors who have

already crossed the developmental threshold perceive the current situation as one that they can handle with ease and that is straightforward to navigate. In contrast, juniors, who are on the verge of crossing the threshold, encounter the same situation as one in which crossing it has become more difficult. Each group sees what their position makes visible; what the other experiences remains, to a significant degree, outside their field of perception.

This asymmetry extends beyond the present moment. While several senior participants expressed confidence that the current disruption would eventually resolve, their optimism was not directed at the juniors currently navigating the transition, but at a future generation who will grow up with AI from the beginning. S2 drew this distinction explicitly:

“The current juniors are a bit in-between, honestly.”

Yet in the same breath, S2 expressed confidence:

“Kids who’ve been building things with AI since they were young will already be at a senior level.”

Junior participants recognized that they were not the intended beneficiaries of this optimism. J5 observed:

“We’ve experienced both worlds. But the ones who are doing well these days have been using it since middle and high school—I don’t think they’d even notice there’s a problem.”

The consequence of these asymmetries is that attempts to address this developmental disruption tend not to converge. Those in a position to act—seniors and organizations—do not fully perceive the structural conditions that juniors are navigating. Those who experience those conditions most directly—juniors—are not in a position to change them. S4, who recognized the problem with unusual clarity, captured this impasse:

“*Who is going to become the next senior?* This is a serious problem, and nobody has a solution.”

When the perception of the problem and the experience of the problem do not coincide, the conditions for self-correction are structurally absent.

## Discussion

### Absorption and Unprotected Pathway

According to our findings, the mechanisms through which GenAI affects the skill development pathways of juniors can be summarized as follows: not only has GenAI taken over the entry-level tasks that previously supported the growth of juniors, but the skills that seniors—who currently use GenAI in their work—believe juniors should possess are also being hindered due to juniors’ use of GenAI in university, thereby widening the gap between seniors and juniors. We go one step further to argue that this mechanism did not arise solely due to the emergence of GenAI, but rather that it has been added onto the *already unprotected* development pathway of juniors.

Senior participants in this study described their own developmental trajectories not as the product of formalized training, but as the incremental accumulation of hands-on

work that, in retrospect, was *not* systematized (Begel and Simon 2008). S1, reflecting on his own early-career trajectory, found that the source of his growth was the combination of willingly doing *menial* tasks and undertaking *self-directed* efforts to expand beyond those tasks:

“[They saw] that I really did the minor work diligently, exactly as told, without a word of complaint... And beyond that—since I didn’t know anything about the technical work at the time—I would look up tutorials on my own and try to expand what I was capable of. When people saw that, they began entrusting me with more challenging tasks, and [my senior] came to trust me and gave me harder things, so my scope kept expanding. Looking back, I think it wasn’t a matter of ‘this is just minor work’—it was the mindset that I could do even that kind of work, and that’s what allowed me to keep growing.”

S4 echoed this pattern from his own trajectory in software engineering, recalling that the same accumulation of *menial* entry-level tasks had constituted his own developmental path. Working as a part-time developer *without* specific assignments, he simply took on whatever came up:

“When I was an undergrad, I started as a part-time developer at [Large Enterprise A], and the first thing I did was what we call the *menial* work of development. Writing code as told, debugging every day, checking and testing whatever came back... I did a lot of so-called development grunt work. [...] that’s how I think I got to where I am now.”

Together, these accounts illustrate the unprotected character of this pathway in concrete terms: the conditions for development depended on individual circumstance, not institutional design.

The same S4, however, also reported that the contemporary equivalent of this work—namely, the entry-level tasks that juniors had previously been assigned in his startup, which had formed the very foundation of his cohort—had been absorbed by GenAI tools without any measurable productivity loss. Yet, S4 describes this very work—the kind that had constituted his own developmental trajectory—as substitutable by a “100,000-won Claude subscription.” The contradiction doesn’t seem to be a contradiction in his explanation: both statements appear obvious and make sense within their own contexts. The absorption proceeds through general organizational rationales—efficiency calculations, headcount planning, productivity assessments. As S4’s remarks suggest, no one needs to realize that this process is effectively blocking the pathways through which they themselves grew. This lack of awareness is made structurally possible because the costs of adjustments arise at a different point in time and for a different group than the decision-makers. In this way, the absorption continues, with no one in the right position to recognize or protect the pathway; what was already unprotected is further eroded.

## Extending Learning Theory and Situated Cognition

*Productive failure* (Kapur 2016) and *desirable difficulties* (Bjork 1994)—identifying struggle, error, and revision as the conditions of deep, transferable learning—were articulated at the scale of the individual learner; our findings suggest that the conditions these theories identify as essential for learning are determined by a much broader dimension—such as the organizational and institutional levels. Specifically, our findings extend these theories to two larger contexts—the workplace and university classroom—where the elimination of error-prone tasks and the collective normalization of GenAI use (Choudhuri et al. 2025) jointly forecloses productive struggle. As our findings showed, individual disposition is overridden by collective pressure under such conditions, as J2’s account of being unable to resist time pressure made explicit. Therefore, we suggest that productive failure be designed at organizational and institutional scales, through structural arrangements rather than through learner choice.

However, just pointing out that these conditions are getting weaker due to the use of GenAI doesn’t explain why this phenomenon continues to be left unaddressed, and this is where the idea of *situated cognition* is helpful. As applied in our analysis of misaligned perspectives, *situated cognition* (Brown, Collins, and Duguid 1989; Lave 1988; Lave and Wenger 1991) holds that what is perceptible to a practitioner is shaped by their situated engagement with practice. Building on this, our findings illuminate a more specific dynamic: when the practice itself is undergoing rapid transformation, the perception of that transformation diverges asymmetrically across experience levels. In other words, seniors who have already crossed the developmental threshold perceive the current moment based on their accumulated competence, whereas juniors mid-crossing the threshold perceive the same moment as one in which the threshold itself has receded. Each group identifies what they see from their own perspective, as the contrast between S4’s and J2’s accounts in our findings is made explicit. As a result, individual-level interventions, such as mentorship, advice, and encouragement, cannot, on their own, overcome this asymmetry. Combined with the organizational scale of the erosion itself described earlier, this implies again that sustaining the conditions for junior development requires structural intervention rather than individual remediation.

## Implications for Institutional Design

If self-correction does not occur on its own, the conditions for junior development must be re-established through institutional intervention rather than relying on individual initiative. S6’s account illustrates what such intervention can look like in practice. She described how her team integrates newcomers into active work, through what they call “first step, one step, big step.” The first step is small and concrete: “finding something like a missing period in one of our production projects, putting the period in, committing it, opening a pull request, getting reviewed, and pushing it all the way to pro-

duction”<sup>3</sup>—that is, taking even the smallest change through the team’s full product release process, so that the newcomer sees their own contribution become part of the final product. Subsequent step expands into larger tasks juniors must approach with increasing self-direction—defining the problem themselves, consulting relevant seniors, and sharing what they have done—so that they can naturally learn how work progresses within the team as well as the technical workflow. The team assigns these tasks without explaining how to solve them, and in a final stage, they simply ask newcomers to “find a problem and try to solve it.” S6 frames this design as “*a kind of learning by doing*”:

“rather than teaching them ‘here’s the process, learn it,’ we’ve designed it so that they feel it by doing it themselves.”

This approach concretizes one form of institutional intervention. The key lesson from this case—namely, that the environment for training juniors should not be left to develop naturally but must be systematically designed—has significant implications across three areas: university education, evaluation criteria for juniors, and workplace onboarding. At the educational stage, the question is how to preserve courses in which the learning objectives cannot be fulfilled by AI on the student’s behalf—an issue J6 raised. By revising educational standards to designate such courses as mandatory, “reducing AI dependency” can be established as a criterion for assessing educational quality rather than leaving a matter to faculty’s discretion. At workforce entry, hiring criteria and evaluation need to shift: as seniors are now recognized for their ability to solve higher-order problems using GenAI, evaluations for juniors should focus on their capacity to recognize their knowledge gaps and errors when using GenAI. Once juniors are inside the workplace, S6’s example suggests that organizations may consider preserving designated learning spaces. In other words, these spaces are not merely a way to improve work efficiency, but rather a place where employees can perform specific, small-scale tasks that are valued as a stepping stone for their growth.

What forms such interventions might look like can be motivated by examining two other fields that show similar patterns: in aviation, FAA SAFO 13002 urges airlines to preserve opportunities for manual flying so that pilot skills do not atrophy as cockpit automation expands (Federal Aviation Administration 2013), a concern confirmed by studies of cognitive degradation in manual flight (Casner et al. 2014; Ebbatson et al. 2010). In nuclear power, NRC requalification requirements mandate that licensed operators undergo periodic simulator training to maintain emergency-response capabilities that day-to-day automated operation rarely exercises (U.S. Nuclear Regulatory Commission 2025). In both cases, institutional design preserves developmental conditions that automation would otherwise erode.

Without a comparable design, the development pathway through which the next generation of seniors might be formed will continue to attenuate.

---

<sup>3</sup>All of these are simple and clear tasks for software engineers.

## Limitations

Our findings should be read alongside several limitations of the study’s design and scope. Our sample of 14 participants—eight juniors and six seniors—is appropriate for the kind of exploratory mechanism analysis we conducted (Braun and Clarke 2006, 2019), but does not support claims about the prevalence or magnitude of the patterns we identified. Nevertheless, we aimed to identify and interpret the mechanisms through which GenAI is reshaping junior development pathways, not to estimate their distribution.

Our cross-sectional design captures senior and junior perspectives at a single moment; the senior–junior comparison we present therefore reflects positional difference under conditions of situated cognition asymmetry rather than longitudinal change. Tracing how today’s juniors actually develop over time will require future longitudinal work.

Our findings are situated in South Korea—a context characterized by the recent collapse of structured graduate recruitment, rigid employment protection, and one of the world’s highest rates of GenAI adoption—and, as we note in the Discussion, whether the dynamics we documented emerge under different institutional configurations is a question for future comparative research.

Beyond sample size, our recruitment via purposive sampling through personal professional networks introduces selection effects. We deliberately constrained the sample to enable focused mechanism analysis: all participants were trained in electrical and computer engineering or computer science, and all junior participants had experienced both the pre- and post-GenAI periods during their undergraduate years. Details are demonstrated in our method section. While this homogeneity supports the comparative analysis we conducted, there may be limitations to directly generalizing our findings to other professional fields, educational backgrounds, or cohorts that emerged either before or after the widespread adoption of GenAI.

## Conclusion

Through 14 semi-structured interviews with juniors and seniors in software engineering, we have traced how GenAI is reshaping the early-career developmental pathway. The mechanisms we identified—visible through participants’ accounts while operating beyond their individual choice—connect workplace task allocation and classroom norm formation. The foundational pattern of Absorption together with the structural mechanism of collective normalization in the university classroom produces the erosion of competence through lost failure experiences across both workplaces and classrooms; and the perceptual asymmetry between seniors and juniors prevents self-correction through individual effort alone. Crucially, this pathway was already unprotected; GenAI has made it consequential. Deliberate institutional design in classrooms, in workplaces, and in the criteria recognizing junior competence will be required to form the next generation of senior practitioners. The question in our title—“Who will become the next senior?”—names a responsibility that ordinary organizational rationality is now structurally incapable of discharging on its own.

## Ethical Considerations

This study was approved by the Institutional Review Board (IRB) of the authors' institution. The IRB approval number is omitted to maintain anonymity during review and will be provided upon acceptance. All participants provided informed consent prior to participation and were assured of confidentiality. Participation was voluntary, and participants were free to withdraw at any time without consequence. Audio recordings and transcripts are stored securely and will be retained for 3 years following publication.

## Positionality

The authors are ML researchers with backgrounds in machine learning and AI systems, situated within the engineering tradition. This background provided sensitivity to the technical dimensions of GenAI capabilities and task boundaries, while requiring reflexive attention to potential biases toward technical explanations of sociotechnical phenomena (Berger 2015). Throughout the analysis, the authors actively sought interpretations grounded in participants' own framings rather than imposing engineering-centric categories onto the data.

## Adverse Impact Statement

This study documents structural dynamics through which GenAI is currently disrupting junior software engineer development. We anticipate three principal risks of misuse or misreading. First, the empirical findings could be misread as confirming that junior engineers are economically substitutable by AI, providing cover for hiring decisions that further reduce entry-level opportunities. This reading inverts our argument: we document the absorption of developmental work as a problem to be addressed, not as a fact to be optimized. The mechanism we identify suggests that reducing junior hiring exacerbates rather than resolves the structural condition, by accelerating the disappearance of the development pathway through which the next generation of seniors is formed. Second, our analysis of perceptual asymmetry could be cited to absolve senior engineers and decision-makers of responsibility—"seniors cannot perceive the problem because of situated cognition." The argument we make is the opposite: recognizing this asymmetry imposes a heightened, not diminished, duty of deliberate institutional design, because individual effort alone cannot restore the structural conditions that self-correction requires. Third, our findings are situated in the South Korean institutional context, treated as a critical case in which several conditions converge. Direct policy transfer to contexts with different labor protection regimes, education systems, or AI adoption rates is not warranted. The mechanism we identify likely operates in those contexts but its strength, distribution, and points of remediation may differ; comparative empirical work is required before specific recommendations can travel. We have framed our normative claims at the institutional rather than individual level precisely to discourage individualized interpretations—whether of blame (against juniors who use GenAI) or of optimism (that motivated individuals can resolve structural conditions on their own).

## References

- Anthony, C. 2021. When knowledge work and analytical technologies collide: The practices and consequences of black boxing algorithmic technologies. *Administrative science quarterly*, 66(4): 1173–1212.
- Bank of Korea. 2025a. AI Diffusion and Youth Employment. Technical Report BOK Issue Note 2025-30, Bank of Korea.
- Bank of Korea. 2025b. Rapid Adoption of Artificial Intelligence and Its Productivity Effects: The Case of Korea. Technical Report BOK Issue Note 2025-22, Bank of Korea.
- Barke, S.; James, M. B.; and Polikarpova, N. 2023. Grounded copilot: How programmers interact with code-generating models. *Proceedings of the ACM on Programming Languages*, 7(OOPSLA1): 85–111.
- Beane, M. 2019. Shadow learning: Building robotic surgical skill when approved means fail. *Administrative Science Quarterly*, 64(1): 87–123.
- Beane, M. 2024. *The Skill Code: How to Save Human Ability in an Age of Intelligent Machines*. New York: Harper Business.
- Begel, A.; and Simon, B. 2008. Novice software developers, all over again. In *Proceedings of the fourth international workshop on computing education research*, 3–14.
- Berger, R. 2015. Now I see it, now I don't: Researcher's position and reflexivity in qualitative research. *Qualitative research*, 15(2): 219–234.
- Bjork, R. A. 1994. Memory and Metamemory Considerations in the Training of Human Beings. In Metcalfe, J.; and Shimamura, A. P., eds., *Metacognition: Knowing about Knowing*, 185–205. Cambridge, MA: MIT Press.
- Bone, M.; Ehlinger, E. G.; and Stephany, F. 2025. Skills or degree? The rise of skill-based hiring for AI and green jobs. *Technological Forecasting and Social Change*, 214: 124042.
- Boyatzis, R. E. 1998. *Transforming qualitative information: Thematic analysis and code development*. Sage.
- Braun, V.; and Clarke, V. 2006. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2): 77–101.
- Braun, V.; and Clarke, V. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health*, 11(4): 589–597.
- Braun, V.; and Clarke, V. 2021. To saturate or not to saturate? Questioning data saturation as a useful concept for thematic analysis and sample-size rationales. *Qualitative research in sport, exercise and health*, 13(2): 201–216.
- Braverman, H. 1998. *Labor and monopoly capital: The degradation of work in the twentieth century*. nyu Press.
- Brown, J. S.; Collins, A.; and Duguid, P. 1989. Situated cognition and the culture of learning. *Educational researcher*, 18(1): 32–42.
- Brynjolfsson, E.; Chandar, B.; and Chen, R. 2025. Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Technical report, Stanford Digital Economy Lab.

- Brynjolfsson, E.; Li, D.; and Raymond, L. 2025. Generative AI at work. *The Quarterly Journal of Economics*, 140(2): 889–942.
- Casner, S. M.; Geven, R. W.; Recker, M. P.; and Schooler, J. W. 2014. The retention of manual flying skills in the automated cockpit. *Human factors*, 56(8): 1506–1516.
- Choudhuri, R.; Ramakrishnan, A.; Chatterjee, A.; Trinkenreich, B.; Steinmacher, I.; Gerosa, M.; and Sarma, A. 2025. Insights from the frontline: Genai utilization among software engineering students. In *2025 IEEE/ACM 37th International Conference on Software Engineering Education and Training (CSEE&T)*, 1–12. IEEE.
- Cui, K. Z.; Demirel, M.; Jaffe, S.; Musolff, L.; Peng, S.; and Salz, T. 2026. The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers. *Management Science*.
- Dagenais, B.; Osher, H.; Bellamy, R. K.; Robillard, M. P.; and De Vries, J. P. 2010. Moving into a new software project landscape. In *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering-Volume 1*, 275–284.
- Ebbatson, M.; Harris, D.; Huddleston, J.; and Sears, R. 2010. The relationship between manual handling performance and recent flying experience in air transport pilots. *Ergonomics*, 53(2): 268–277.
- Eloundou, T.; Manning, S.; Mishkin, P.; and Rock, D. 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702): 1306–1308.
- Federal Aviation Administration. 2013. Safety Alert for Operators SAFO 13002: Manual Flight Operations. Technical Report SAFO 13002, FAA Flight Standards Service.
- Feng, D.; Yun, B.; and Wang, A. Y. 2026. From Junior to Senior: Allocating Agency and Navigating Professional Growth in Agentic AI-Mediated Software Engineering. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, 1–24.
- Hosseini Maasoum, S. M.; and Lichtinger, G. 2025. Generative AI as seniority-biased technological change: Evidence from US résumé and job posting data. Available at SSRN 5425555.
- Kapur, M. 2016. Examining productive failure, productive success, unproductive failure, and unproductive success in learning. *Educational psychologist*, 51(2): 289–299.
- Kobiella, C.; Flores López, Y. S.; Waltenberger, F.; Draxler, F.; and Schmidt, A. 2024. "If the machine is as good as me, then what use am I?"—How the use of ChatGPT changes young professionals' perception of productivity and accomplishment. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–16.
- LaToza, T. D.; Venolia, G.; and DeLine, R. 2006. Maintaining mental models: a study of developer work habits. In *Proceedings of the 28th international conference on Software engineering*, 492–501.
- Lave, J. 1988. *Cognition in practice: Mind, mathematics and culture in everyday life*. Cambridge University Press.
- Lave, J.; and Wenger, E. 1991. *Situated learning: Legitimate peripheral participation*. Cambridge university press.
- Lee, K. 2025. Korea's tech hiring drops as AI replaces entry-level developer jobs. *ChosunBiz (Chosun Ilbo)*.
- Liang, J. T.; Yang, C.; and Myers, B. A. 2024. A large-scale survey on the usability of ai programming assistants: Successes and challenges. In *Proceedings of the 46th IEEE/ACM international conference on software engineering*, 1–13.
- Lincoln, Y. S.; and Guba, E. G. 1985. *Naturalistic inquiry*. sage.
- Malterud, K.; Siersma, V. D.; and Guassora, A. D. 2016. Sample size in qualitative interview studies: guided by information power. *Qualitative health research*, 26(13): 1753–1760.
- Maslej, N.; Fattorini, L.; Perrault, R.; Gil, Y.; Parli, V.; Kariuki, N.; Capstick, E.; Reuel, A.; Brynjolfsson, E.; Etchemendy, J.; Ligett, K.; Lyons, T.; Manyika, J.; Niebles, J. C.; Shoham, Y.; Wald, R.; Walsh, T.; Hamrah, A.; Santarlasci, L.; Lotufo, J. B.; Rome, A.; Shi, A.; and Oak, S. 2025. The AI Index 2025 Annual Report. Technical report, AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA.
- Peng, S.; Kalliamvakou, E.; Cihon, P.; and Demirel, M. 2023. The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590*.
- Prather, J.; Reeves, B. N.; Leinonen, J.; MacNeil, S.; Randrianasolo, A. S.; Becker, B. A.; Kimmel, B.; Wright, J.; and Briggs, B. 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.
- Rak, G. 2025. How to Stay Ahead of AI as an Early-Career Engineer. IEEE Spectrum.
- Shin, S.; Yun, B.; Chin, H.; and Cho, Y. 2025. The Impact of Generative AI on the Future of Software Developer Hiring: Insights from an Expert Delphi Survey. *JOURNAL OF KOREAN ASSOCIATION FOR REGIONAL INFORMATION SOCIETY*, 28(2): 69–110.
- U.S. Nuclear Regulatory Commission. 2025. 10 CFR § 55.59 — Requalification. Code of Federal Regulations. Last amended 89 FR 106253, Dec. 30, 2024; eCFR current as of Jan. 29, 2025.
- Vaithilingam, P.; Zhang, T.; and Glassman, E. L. 2022. Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models. In *Chi conference on human factors in computing systems extended abstracts*, 1–7.
- Wenger, E. 1999. *Communities of practice: Learning, meaning, and identity*. Cambridge university press.
- Woodruff, A.; Shelby, R.; Kelley, P. G.; Rousso-Schindler, S.; Smith-Loud, J.; and Wilcox, L. 2024. How knowledge workers think generative ai will (not) transform their industries. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–26.

World Economic Forum. 2025. The Future of Jobs Report 2025. Insight report, World Economic Forum, Geneva, Switzerland.

Xiao, Q.; Hu, X. E.; Whiting, M. E.; Karunakaran, A.; Shen, H.; and Cao, H. 2025. AI Hasn't Fixed Teamwork, But It Shifted Collaborative Culture: A Longitudinal Study in a Project-Based Software Development Organization (2023-2025). *arXiv preprint arXiv:2509.10956*.

## Semi-Structured Interview Protocols

The following guides reflect the semi-structured interview protocols used with junior and senior participants. They were developed around shared anchor topics—the changing nature of entry-level work, the conditions of skill formation, and expectations about junior–senior development—to which each group could respond from its respective vantage point. As semi-structured protocols, not all questions were asked verbatim or in fixed order; the order in which topics were taken up was adapted to each participant’s responses and conversational style, and when accounts opened onto themes relevant to our research questions, follow-up questions were developed in the moment to deepen those threads. Some questions used scenario-based framings—for example, asking senior participants to recall specific moments of considering AI versus junior delegation—to elicit concrete accounts rather than general opinions. Representative follow-up probes are included below under some questions, but these were asked flexibly (e.g., with additional probes sometimes added) during the interview and did not follow a fixed schedule. The original protocols were written and conducted in Korean; the translations below are provided for accessibility.

### Junior Interview Guide

#### Opening / Background

- When did you first begin using GenAI tools in your studies?
- Could you describe the courses you have taken before and after the widespread availability of GenAI tools?

#### Reflections on your educational trajectory

- Looking back on your undergraduate experience, can you compare what learning was like before widespread GenAI availability versus after?
  - *Probe:* Were there specific courses, projects, or moments that crystallized this difference for you?
  - *Probe:* Have you noticed differences in your own study habits, sense of effort, or sense of accomplishment?
- Have you experienced any course or project where the role of GenAI differed significantly, and what was that experience like?
- Are there capabilities you feel you developed primarily during the pre-GenAI period that you would not have developed in the same way today?

#### GenAI integration in academic work

- How often do you currently use GenAI in your learning process?
- Are there tasks where you feel that, because of GenAI, you have had fewer opportunities to actually experience the work yourself?
  - *Probe:* Could you describe a specific course or assignment where this was most pronounced?
- Have you ever felt that you should be working through a problem on your own to develop your skills, but ended up relying on AI? How did you feel about that?

- *Probe:* What pressures made it hard to resist using AI in that moment?

- Has your way of studying changed over the past year, with LLMs becoming widely available?
- Looking at the topics you have learned so far—through coursework, projects, internships—are there parts where you think, “GenAI could just do all of this”?

- *Probe:* Does this realization change how you think about what you should focus on going forward?

#### Anticipated transition to the workforce

- Among the topics you have studied in university, which do you think will directly apply to professional work?
- During your job search, have you noticed expressions or requirements in job descriptions that strike you as substantially different from what you have learned in school?

#### Hiring criteria and AI literacy

- Do you sense that what companies expect from new hires has changed compared to earlier periods? If so, how?
- In the era of GenAI, have you reconsidered what you should now be learning?
- How do you think you can demonstrate your AI competence to employers—through projects, portfolios, certifications, or other forms?

#### Closing

- Is there anything important about your experience with GenAI, your education, or your career development that we have not yet discussed?
- Is there anyone in your network—peers, recent graduates, mentors—whom you would suggest we speak with for this study?

### Senior Interview Guide

#### Opening / Background

- Could you describe your current role, the organization you work in, and the team you work with?
- Have you been directly involved in training or mentoring juniors? Was this before or after GenAI tools became widely available—or both?

#### Reflections on your own developmental trajectory

- Could you walk us through your own development from a junior to a senior engineer? What kinds of work did you do during your first one to three years?
- Looking back, were there things you worked on as a junior that did not seem meaningful at the time but turned out to be important for your development?
  - *Probe:* Was there a particular moment of recognition or change for you during that process?
- As a junior, were there tasks that a senior could have completed themselves but deliberately delegated to you for your development?
  - *Probe:* From the senior’s perspective, what cost or risk did they take on by delegating these tasks?

- In the team you currently work with as a senior, is junior development handled in a similar way to how it was when you were a junior—or has it changed?
  - *Probe*: If similar, why do you think these practices have persisted despite the broader changes in the field?
  - *Probe*: If different, when and why do you think this changed?

### GenAI integration and changes in team workflow

- How often do you currently use GenAI in your own work?
- Based on your own experience, what kinds of tasks does GenAI do well, and what kinds does it not do well?
- Have you recently had the experience of thinking, “I could have delegated this task to a junior, but I just handled it myself with AI”? If you do not work closely with juniors, have you observed this pattern in other parts of your organization?
  - *Probe*: What was the deciding factor for not delegating (e.g., time, explanation, feedback cost, quality, or responsibility)?
- Has the kind of work assigned to juniors changed since GenAI became available?
  - *Probe*: Have the criteria become higher? Has the nature of the work changed?

### Junior development pathway

- Earlier we discussed the kinds of work that have been most important for junior development. Looking at things now, who currently does that work? Has it disappeared from junior portfolios?
- Have you considered how these changes might affect the organization in the longer term?
  - *Probe*: Are you or your team thinking about how to address this?

### Hiring criteria and future competency needs

- In recent hiring of junior engineers, would you say the criteria are closer to “they need to be more capable than before,” or closer to “we are now looking for different capabilities”?
  - *Probe*: What specifically are these new or different capabilities?
- Do you think these capabilities can be adequately developed through university education?
  - *Probe*: If not, where do you think they should be developed?
- Looking three to five years ahead, what competencies do you think will become most important for juniors?

### Closing

- Is there anything important about GenAI’s impact on junior engineering, team formation, or career development that we have not yet discussed?
- Is there anyone in your network whom you would suggest we speak with for this study?

## Analytical Approach

We analyzed the dataset using Reflexive Thematic Analysis (RTA) following the framework articulated by Braun and Clarke (2006, 2019). RTA was chosen over alternative qualitative approaches—including codebook-style thematic analysis (Boyatzis 1998) and Grounded Theory—for two reasons. First, our research questions concerned the identification of mechanisms within an emergent phenomenon, not the construction of a new substantive theory from the ground up; this aligned with RTA’s interpretive engagement orientation rather than Grounded Theory’s theory-building program. Second, we are committed to the view that themes are constructed through analytical engagement rather than discovered through reliability procedures; this aligned with RTA’s reflexive, single-analyst orientation rather than codebook approaches relying on inter-rater reliability.

Consistent with this orientation:

- Analysis was conducted by the authors throughout, with iterative engagement across multiple readings of the full dataset.
- Inter-rater reliability scoring and saturation counts were not employed; Braun and Clarke (2019) discuss why such procedures, drawn from different qualitative paradigms, are not constitutive of analytical rigor in RTA.
- Reflexive engagement was supported by the authors’ positionality (described in the main paper), with explicit attention to potential engineering-centric biases in interpretation.
- Disconfirming cases (notably S1, examined as a boundary condition in Theme 1) were actively sought during phase 4 and integrated into the analysis rather than treated as outliers.

## Ethical Procedures

**IRB approval.** This study received approval from the Institutional Review Board (IRB) at the authors’ home institution. Specific IRB identifying details have been redacted to comply with double-blind review. The approval covered semi-structured interviews with software engineering professionals and students, audio recording, transcription, and analysis. No modifications to the approved protocol were made during the study.

**Recruitment.** Participants were recruited through purposive sampling from the authors’ professional network. Initial contact was made via direct messaging using a brief recruitment script that described the study’s purpose, expected time commitment (approximately 50 minutes), voluntary nature, and anonymization procedures.

**Informed consent.** All participants provided informed consent prior to participation. Consent was obtained both verbally (recorded) and via written form in Korean, the participants’ native language. The consent process explicitly covered: the study’s purpose; what data would be collected and how it would be used; anonymization and storage procedures; the right to withdraw at any time without consequence; and the right to decline to answer specific questions.

**Confidentiality and anonymization.** Participant names and direct identifiers were replaced with coded labels (J1–J8 for

juniors, S1–S6 for seniors) immediately upon transcription. Organizational identifiers were generalized to category labels (e.g., "Large Enterprise A," "Mid-sized Company"). Quotes selected for use in this paper were reviewed by the authors for any residual identifying details—specific project names, distinctive team sizes, technology stacks, or product references—and modified or omitted where re-identification risk was present.

**Data storage and retention.** Audio recordings and transcripts are stored on a password-protected device accessible only to the authors. Identifying information (names, contact details, organization names) is stored separately from de-identified transcripts. All study materials will be retained for three years following publication, after which they will be permanently deleted in accordance with the IRB-approved retention plan.

**Handling sensitive topics.** Several topics covered in interviews were professionally sensitive—for junior participants, reflections on faculty practices and anticipated job market entry under uncertainty; for senior participants, reflections on hiring decisions and team management changes. Interviewers explicitly reminded participants of the confidentiality protections at moments where sensitive material arose, and offered participants the option to mark specific statements as off-the-record. Statements marked off-the-record were not transcribed and are not represented in our analysis.

**Researcher positioning.** The authors conducted interviews from a position of relative seniority with respect to junior participants and relative juniority with respect to senior participants. This dual positioning was attended to during interviewing—particular care was taken not to lead junior participants toward critical accounts of their education, and to avoid deferential acquiescence to senior accounts.