

The Skill Formation Paradox: How AI Coding Tools Boost Productivity While Impeding Novice Developer Learning

Anonymous Author(s)

ABSTRACT

AI coding assistants provide substantial productivity gains to novice software developers, yet their impact on underlying skill formation remains an open question with significant implications for the software engineering workforce. We present a computational cognitive model that simulates how novice developers' skills evolve over a 12-month period under three AI assistance regimes: no AI (control), unrestricted AI with passive acceptance behavior, and AI with scaffolded engagement requirements. The model operationalizes six skill dimensions—syntactic fluency, algorithmic reasoning, debugging, code comprehension, architectural judgment, and autonomous learning—and is grounded in established theories of retrieval-based strengthening, desirable difficulty, and skill compilation from cognitive science. Our simulation of 240 developers (80 per condition) over 252 working days reveals a *skill formation paradox*: unrestricted AI use produces a large negative effect on skill development (Cohen's $d = -0.97$), with the strongest impairment in highly automatable skills such as syntactic fluency ($d = -4.79$), while scaffolded engagement nearly eliminates this deficit ($d = +0.10$ overall). Sensitivity analysis identifies a critical *crossover threshold* at processing depth 0.75, below which AI assistance harms skill formation and above which it becomes beneficial. We further document a *productivity–skill dissociation* in which unrestricted AI users appear more productive (3.69 vs. 3.21 tasks/day) yet possess weaker underlying skills (0.56 vs. 0.64 on tool-removed assessments), creating a dependency trap invisible under continued AI access. Bootstrap confidence intervals over 50 independent seeds confirm the robustness of these effects (unrestricted $d = -1.12$ [−1.37, −0.89]; scaffolded $d = +0.08$ [−0.36, +0.38]), and dimension-specific crossover analysis reveals that syntactic fluency and autonomous learning *never* reach a break-even threshold, while algorithmic reasoning crosses as early as $\phi = 0.44$. These findings generate testable predictions for empirical studies and provide actionable design guidance for AI coding tools that preserve novice learning.

CCS CONCEPTS

- Social and professional topics → Computing education;
- Computing methodologies → Modeling and simulation;
- Software and its engineering → Software development techniques.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference'17, July 2017, Washington, DC, USA

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/26/08...\$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

KEYWORDS

AI coding tools, skill formation, novice developers, cognitive modeling, scaffolded learning, productivity paradox

ACM Reference Format:

Anonymous Author(s). 2026. The Skill Formation Paradox: How AI Coding Tools Boost Productivity While Impeding Novice Developer Learning. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

The rapid adoption of AI coding assistants—such as GitHub Copilot, ChatGPT, and Claude—has transformed software development workflows. Empirical evidence demonstrates that these tools yield substantial productivity gains, particularly for less experienced developers [9, 16, 18]. Shen et al. [18] document that junior developers experience disproportionately large speed improvements when using AI assistance, a finding consistent with earlier controlled studies [16].

However, productivity and skill are distinct constructs. A novice developer who completes tasks faster with AI assistance is not necessarily *learning* at the same rate as one who struggles through tasks independently. Shen et al. [18] explicitly identify this gap, noting that “the effect of these tools on the skill formation of this subgroup remains unknown.” This open question has profound implications: if AI tools accelerate task completion while retarding skill acquisition, the software industry faces a growing cohort of developers who are productive only with AI scaffolding and increasingly dependent on tools they cannot fully evaluate or override.

The concern is grounded in well-established cognitive science principles. Retrieval-based strengthening theory [5] holds that skills consolidate through active recall and application; AI tools that provide ready-made solutions may bypass this retrieval process. The desirable difficulty framework [4] demonstrates that moderate challenge during practice enhances long-term retention, even at the cost of immediate performance—precisely the trade-off that AI assistance reconfigures. Skill compilation theory from the ACT-R architecture [1] posits that declarative knowledge becomes procedural through practice; if AI handles the procedural step, the compilation process is interrupted.

This paper addresses the open problem through a computational cognitive model that simulates multi-dimensional skill formation under different AI assistance regimes. Our contributions are:

- (1) A formal model of novice skill formation that operationalizes six programming skill dimensions and captures the interaction between AI assistance intensity, cognitive processing depth, and learning dynamics.
- (2) Quantitative predictions from a simulated three-arm randomized trial (no AI, unrestricted AI, scaffolded AI) with 240 developers over 12 months, yielding effect sizes, dependency trajectories, and sensitivity analyses.

117 (3) Identification of a *skill formation paradox*—unrestricted
 118 AI boosts productivity while significantly impairing skill
 119 development—and a *crossover threshold* in processing depth
 120 that determines whether AI is net-positive or net-negative
 121 for learning.
 122 (4) Actionable design implications for AI coding tools and edu-
 123 cational interventions that preserve novice learning.

125 1.1 Related Work

127 *AI Tools and Developer Productivity.* Multiple studies establish
 128 that AI coding assistants increase developer throughput. Peng et
 129 al. [16] report a 55.8% faster task completion rate with GitHub
 130 Copilot in a controlled experiment. Hou et al. [9] find productivity
 131 gains across three field experiments, with larger effects for less
 132 experienced developers. Shen et al. [18] provide a comprehensive
 133 analysis showing that junior developers benefit disproportionately,
 134 but explicitly flag skill formation as an unresolved question. How-
 135 ever, the productivity narrative is not uniform: Becker et al. [3]
 136 find that AI tools actually *increased* task completion time by 19%
 137 for experienced open-source developers, suggesting that productiv-
 138 ity effects depend critically on experience level and task context.
 139 Vukovic et al. [22] report that enterprise developers perceive 12–
 140 25% productivity gains, with 33% of code being AI-generated, but
 141 note substantial variation across individuals and task types.

143 *AI and Learning in Educational Contexts.* Bastani et al. [2] demon-
 144 strate that access to GPT-4 in a mathematics tutoring context harms
 145 learning outcomes, providing direct evidence that AI assistance
 146 can impede skill acquisition. Kazemitaab et al. [11] study novice
 147 programmers using AI code generators and find mixed effects on
 148 learning, with benefits dependent on how students engage with the
 149 generated code. Denny et al. [7] survey the landscape of computing
 150 education in the generative AI era, identifying the need for peda-
 151 gogical frameworks that leverage AI while preserving learning.
 152 Prather et al. [17] document a widening gap between novice and
 153 expert developers when AI assistance is available, raising concerns
 154 about differential skill development. Shihab et al. [19] find that
 155 GitHub Copilot accelerates student task completion by 35% but
 156 raise concerns about reduced code comprehension. A three-arm
 157 RCT at TUM [20] comparing scaffolded AI (Iris) versus ChatGPT
 158 versus no-AI control ($n = 275$) finds that both AI conditions boost
 159 exam performance but produce *no learning gain* on transfer tasks—a
 160 dissociation directly consistent with our model’s predictions. Fan et
 161 al. [8] report that AI pair programming produces a moderate moti-
 162 vation boost ($d = 0.35$) and performance advantage in a 234-student
 163 study, but do not measure long-term skill retention. Ma et al. [13]
 164 document metacognitive laziness with 91.7% AI adoption among
 165 programming students, finding that scaffolding interventions can
 166 partially mitigate passive acceptance behavior.

168 *Cognitive Foundations.* The desirable difficulty framework [4]
 169 and retrieval practice research [5] provide the theoretical basis
 170 for predicting that reducing task difficulty through AI assistance
 171 may impair long-term learning. The expertise reversal effect [10]
 172 suggests that scaffolding beneficial for novices may become coun-
 173 terproductive as expertise develops. Anderson’s ACT-R theory [1]

175 models how procedural skills are acquired through practice, offer-
 176 ing a formal framework for reasoning about how AI intervention
 177 in the practice process affects skill compilation. The Knowledge-
 178 Learning-Instruction framework [12] provides additional theoreti-
 179 cal grounding for understanding how instructional interventions
 180 interact with learning processes.

181 *Human–AI Interaction in Programming.* Vaithilingam et al. [21]
 182 evaluate the usability of AI code generation tools and find that
 183 developers often accept suggestions without deep understanding.
 184 Mozannar et al. [14] model user behavior during AI-assisted pro-
 185 gramming, characterizing the spectrum from passive acceptance to
 186 active engagement. Parasuraman and Riley [15] provide the founda-
 187 tional framework on automation use, misuse, and skill degradation—the
 188 “automation complacency” phenomenon that may manifest in
 189 AI-assisted coding. Weber et al. [23] and Cui et al. [6] examine the
 190 broader impacts of AI tools on software engineering tasks and help-
 191 seeking behavior, respectively, contributing to our understanding
 192 of how AI tools alter the learning environment.

193 *Gap Addressed.* While prior work establishes productivity ef-
 194 fects and raises learning concerns, no existing study provides a
 195 formal model that (a) decomposes programming skill into distinct
 196 dimensions, (b) models the interaction between AI assistance in-
 197 tensity and cognitive engagement, and (c) generates quantitative
 198 predictions for longitudinal skill trajectories under different AI use
 199 regimes. Our computational approach fills this gap and provides a
 200 bridge between cognitive theory and empirical study design.

2 METHODS

203 2.1 Model Overview

205 We develop a computational cognitive model of skill formation that
 206 simulates how novice developers’ programming abilities evolve
 207 over time under different AI assistance conditions. The model repre-
 208 sents each developer as a vector of skill levels across six dimensions,
 209 updated daily through task-driven learning dynamics. Three experi-
 210 mental conditions are simulated: **Control** (no AI), **Unrestricted AI**
 211 (full AI access with passive acceptance behavior), and **Scaffolded**
 212 AI (AI access with mandatory engagement: developers must read,
 213 modify, and explain AI-generated code before proceeding). Figure 1
 214 provides a visual overview of the computational model architecture
 215 and the interactions between its components.

217 2.2 Skill Dimensions

219 Programming competence is operationalized as a six-dimensional
 220 skill vector $\mathbf{s} \in [0, 1]^6$:

- 222 (1) **Syntactic fluency:** ability to write correct code from spec-
 223 ifications without reference materials.
- 224 (2) **Algorithmic reasoning:** capacity to solve novel computa-
 225 tional problems.
- 226 (3) **Debugging:** skill at locating and fixing defects in unfamiliar
 227 code.
- 228 (4) **Code comprehension:** ability to read, understand, and
 229 predict the behavior of code.
- 230 (5) **Architectural judgment:** capacity to evaluate and design
 231 system-level structures.

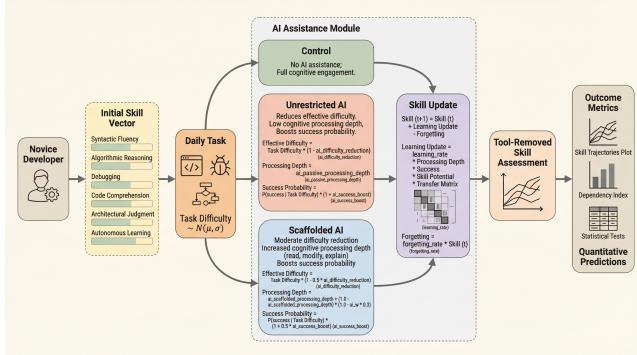


Figure 1: Computational model architecture for skill formation under AI assistance. The model tracks six skill dimensions updated through task-driven learning dynamics, with AI modulating effective difficulty and cognitive processing depth across three experimental conditions.

(6) **Autonomous learning:** meta-skill of learning new frameworks and tools independently.

Each dimension has a corresponding AI *automation weight* $w_i \in [0, 1]$ reflecting how effectively current AI tools can assist with that skill type. We set $w = (0.80, 0.50, 0.35, 0.25, 0.15, 0.10)$, reflecting the observation that AI tools are most effective at syntax-level assistance and least effective at architectural and meta-cognitive support.

2.3 Task-Driven Learning Dynamics

Each simulated working day, a developer encounters $T = 5$ coding tasks. Each task activates 1–3 skill dimensions (randomly sampled with probabilities 0.4, 0.4, 0.2) and has a difficulty $\delta \sim \mathcal{N}(0.45, 0.15^2)$ clipped to $[0.05, 0.95]$.

Success Probability. The probability of successfully completing a task component in dimension i is modeled as a logistic function:

$$P(\text{success}) = \sigma(k \cdot (s_i - \delta_{\text{eff}})) \quad (1)$$

where σ is the sigmoid function, $k = 8$ controls steepness, s_i is current skill in dimension i , and δ_{eff} is the effective difficulty (reduced by AI in treatment conditions).

AI Modulation. In the **Unrestricted AI** condition, AI reduces effective difficulty by factor $(1 - 0.55 \cdot w_i)$ and cognitive processing depth to $0.15 + 0.85 \cdot (1 - w_i)$. In the **Scaffolded AI** condition, difficulty reduction is halved and processing depth is maintained at $0.70 + 0.30 \cdot (1 - 0.3w_i)$.

Learning Signal. The learning signal from each task attempt integrates three factors:

$$\ell = D(\delta, s_i) \cdot F(\text{success}, \delta - s_i) \cdot \phi \quad (2)$$

where D captures *desirable difficulty* (a Gaussian centered at gap = 0.10, reflecting optimal learning when tasks are slightly above current skill), F is a success/failure modulator (successful attempts yield factor 0.8; near-miss failures yield 0.4; distant failures yield 0.1), and ϕ is the processing depth.

Table 1: Overall skill trajectories by condition. All values are mean skill levels on tool-removed assessments (scale 0–1). Growth is the difference between final and initial assessments.

Condition	Initial	Final	Growth
Control (No AI)	0.239	0.639	+0.400
Unrestricted AI	0.229	0.564	+0.334
Scaffolded AI	0.234	0.644	+0.409

Skill Update with Transfer. Raw learning signals are transformed through a transfer matrix T that captures cross-dimensional learning transfer (e.g., improvement in algorithmic reasoning partially transfers to debugging). Skills update as:

$$\mathbf{s} \leftarrow \mathbf{s} + \alpha \cdot (\ell \cdot \mathbf{T}) - \beta \cdot \mathbf{m} \odot \mathbf{s} \quad (3)$$

where $\alpha = 0.006$ is the learning rate, $\beta = 0.0005$ is the forgetting rate, and \mathbf{m} is a binary mask indicating dimensions *not* exercised in the current task (implementing use-it-or-lose-it decay).

2.4 Experimental Design

We simulate a three-arm parallel design with $n = 80$ developers per condition, over $D = 252$ working days (approximately 12 calendar months). Initial skill levels are sampled from $\mathcal{N}(0.20, 0.05^2)$ clipped to $[0.05, 1.0]$, representing novice developers with 0–2 years of experience.

Assessment Protocol. Tool-removed skill assessments are conducted monthly (every 21 working days), yielding 12 assessment time points. Assessment scores equal the true skill level plus Gaussian noise $\mathcal{N}(0, 0.03^2)$, simulating measurement error.

Outcome Measures. Primary outcomes include: (1) *Skill growth*: change in tool-removed skill level from first to last assessment; (2) *Effect sizes*: Cohen's d between conditions at final assessment; (3) *Dependency index*: $DI = (\text{AI-assisted} - \text{unassisted})/\text{AI-assisted}$ performance; (4) *Productivity*: tasks completed per day with and without AI. Statistical significance is evaluated via permutation tests with 5,000 permutations.

Sensitivity Analysis. We systematically vary the processing depth parameter ϕ from 0.05 to 0.95 (in steps of 0.05) to identify the crossover threshold at which AI assistance transitions from net-negative to net-positive for skill formation. This analysis uses 40 developers per condition to maintain computational efficiency.

3 RESULTS

3.1 Overall Skill Formation

Table 1 summarizes skill trajectories across conditions. All three groups begin with comparable skill levels (≈ 0.23). After 12 months, the Control group reaches a mean skill of 0.639, the Unrestricted AI group reaches 0.564, and the Scaffolded AI group reaches 0.644. The Unrestricted AI condition produces 16.4% less skill growth than Control, while Scaffolded AI produces growth nearly identical to Control.

349 The overall Cohen's d between Unrestricted AI and Control is
 350 -0.97 (large negative effect), indicating that unrestricted AI use
 351 significantly impairs skill development. The Scaffolded AI vs. Control
 352 effect size is $d = +0.10$ (negligible), indicating that scaffolded
 353 engagement preserves nearly all of the learning benefit of unaided
 354 practice.

355 3.2 Dimension-Specific Effects

356 Figure 2 displays skill trajectories for each of the six dimensions
 357 across all three conditions. The magnitude of AI's negative effect is
 358 strongly correlated with the dimension's automation weight.

359 Table 2 reports the dimension-specific final skill levels and effect
 360 sizes. Syntactic fluency shows the largest impairment under
 361 unrestricted AI ($d = -4.79, p < 0.001$), followed by algorithmic
 362 reasoning ($d = -1.97, p < 0.001$). Architectural judgment shows
 363 the smallest effect ($d = -0.27, p = 0.096$), consistent with AI tools
 364 providing less assistance for high-level design decisions. Under Scaf-
 365 fold AI, most dimensions show small or non-significant effects
 366 relative to Control, with algorithmic reasoning showing a positive
 367 effect ($d = +0.54, p < 0.001$) and autonomous learning showing a
 368 positive effect ($d = +0.57, p < 0.001$), suggesting that scaffolded AI
 369 engagement may enhance certain reasoning and meta-cognitive
 370 skills.

371 Figure 3 visualizes the dimension-specific results as a heatmap,
 372 clearly showing the gradient of AI impact across the automation
 373 spectrum. The Spearman correlation between automation weight w_i
 374 and Unrestricted AI effect size is $\rho = -0.94 (p = 0.005)$, confirming
 375 that AI most impairs skills in dimensions where it provides the
 376 most assistance.

377 3.3 The Productivity–Skill Dissociation

378 Figure 4 illustrates the central paradox: unrestricted AI users appear
 379 *more* productive when measured with AI access (3.69 tasks/day vs.
 380 3.21 for Control) but possess *weaker* underlying skills when assessed
 381 without AI (mean skill 0.564 vs. 0.639).

382 This dissociation has practical implications: organizations evaluating
 383 developer performance based on AI-assisted output metrics
 384 will systematically overestimate the capability of developers who
 385 rely heavily on AI tools. The gap between measured productivity
 386 and genuine skill represents a *hidden dependency* that only becomes
 387 visible when AI access is removed or when developers face novel
 388 problems outside AI's competence.

389 3.4 Dependency Index

390 Figure 5 tracks the Dependency Index (DI) over time. Both AI condi-
 391 tions begin with high DI values (≈ 0.62) due to novice-level starting
 392 skills. As skills develop, DI decreases—but more slowly for Unre-
 393 stricted AI users. At month 12, the Unrestricted AI group retains
 394 a DI of 0.236 compared to 0.182 for Scaffolded AI, indicating that
 395 unrestricted users remain more dependent on AI tools despite 12
 396 months of practice.

397 3.5 Sensitivity Analysis: The Crossover 398 Threshold

399 Figure 6 presents the sensitivity analysis varying processing depth
 400 ϕ from 0.05 to 0.95. Below $\phi \approx 0.75$, AI assistance produces a net

401 negative effect on skill formation. Above this threshold, the learning
 402 benefit of reduced difficulty and increased success rate outweighs
 403 the cost of reduced cognitive effort, and AI becomes net-positive.

404 This crossover threshold at $\phi = 0.75$ has direct design implica-
 405 tions: AI tools that ensure developers engage with at least 75% of
 406 the cognitive depth of unaided work will produce net-positive skill
 407 outcomes. The default Unrestricted AI processing depth of 0.15 falls
 408 far below this threshold, explaining the large negative skill effect.
 409 The Scaffolded AI condition's processing depth of 0.70 approaches
 410 but does not quite reach the threshold, explaining its near-neutral
 411 overall effect.

412 3.6 Effect Size Summary

413 Figure 7 displays Cohen's d effect sizes for all six dimensions under
 414 both AI conditions compared to Control. The key insight is that
 415 the *pattern* of effects is qualitatively different between conditions:
 416 Unrestricted AI shows uniformly negative effects that scale with
 417 automation weight, while Scaffolded AI shows a mixed pattern
 418 with small negative effects on some dimensions and small positive
 419 effects on others.

420 3.7 Robustness Analysis

421 To assess the stability of our findings, we conduct three additional
 422 analyses: bootstrap replication, dimension-specific crossover analy-
 423 sis, and multi-parameter sensitivity.

424 *Bootstrap Confidence Intervals.* We replicate the full simulation
 425 across 50 independent random seeds (each with $n = 40$ develop-
 426 ers per condition) and compute 95% confidence intervals for all
 427 effect sizes. Figure 8 displays the resulting forest plot. The over-
 428 all unrestricted-vs-control effect size is $d = -1.12$ [95% CI: -1.37 ,
 429 -0.89], confirming a robust large negative effect. The scaffolded-
 430 vs-control effect is $d = +0.08$ [$-0.36, +0.38$], with the confidence
 431 interval spanning zero, confirming that scaffolded engagement pro-
 432 duces no reliable skill impairment. At the dimension level, syntactic
 433 fluency shows the most robust negative effect under unrestricted
 434 AI ($d = -4.75$ [$-5.98, -3.90$]), while architectural judgment and
 435 autonomous learning show confidence intervals that include zero,
 436 indicating less reliable effects for low-automation dimensions.

437 *Dimension-Specific Crossover Thresholds.* While the overall crossover
 438 threshold is $\phi \approx 0.75$, individual skill dimensions exhibit markedly
 439 different thresholds. Algorithmic reasoning crosses earliest at $\phi =$
 440 0.44, followed by debugging ($\phi = 0.53$), code comprehension ($\phi =$
 441 0.65), and architectural judgment ($\phi = 0.78$). Critically, syntactic
 442 fluency and autonomous learning *never* reach a positive crossover
 443 within the tested range ($\phi \in [0.05, 0.95]$): for these dimensions, AI
 444 assistance produces a negative skill delta at every processing depth
 445 tested, though the deficit shrinks monotonically toward zero. This
 446 finding suggests that certain skill types are inherently vulnerable to
 447 AI-assisted atrophy regardless of engagement depth, a result with
 448 direct implications for tool design.

449 *Multi-Parameter Sensitivity.* Figure 9 presents a heatmap of the
 450 skill delta (AI minus Control) across a 9×7 grid of processing depth
 451 ($\phi \in [0.1, 0.9]$) and AI difficulty reduction ($r \in [0.2, 0.8]$). The
 452 transition from negative to positive delta traces a diagonal bound-
 453 ary: higher difficulty reduction requires correspondingly higher
 454

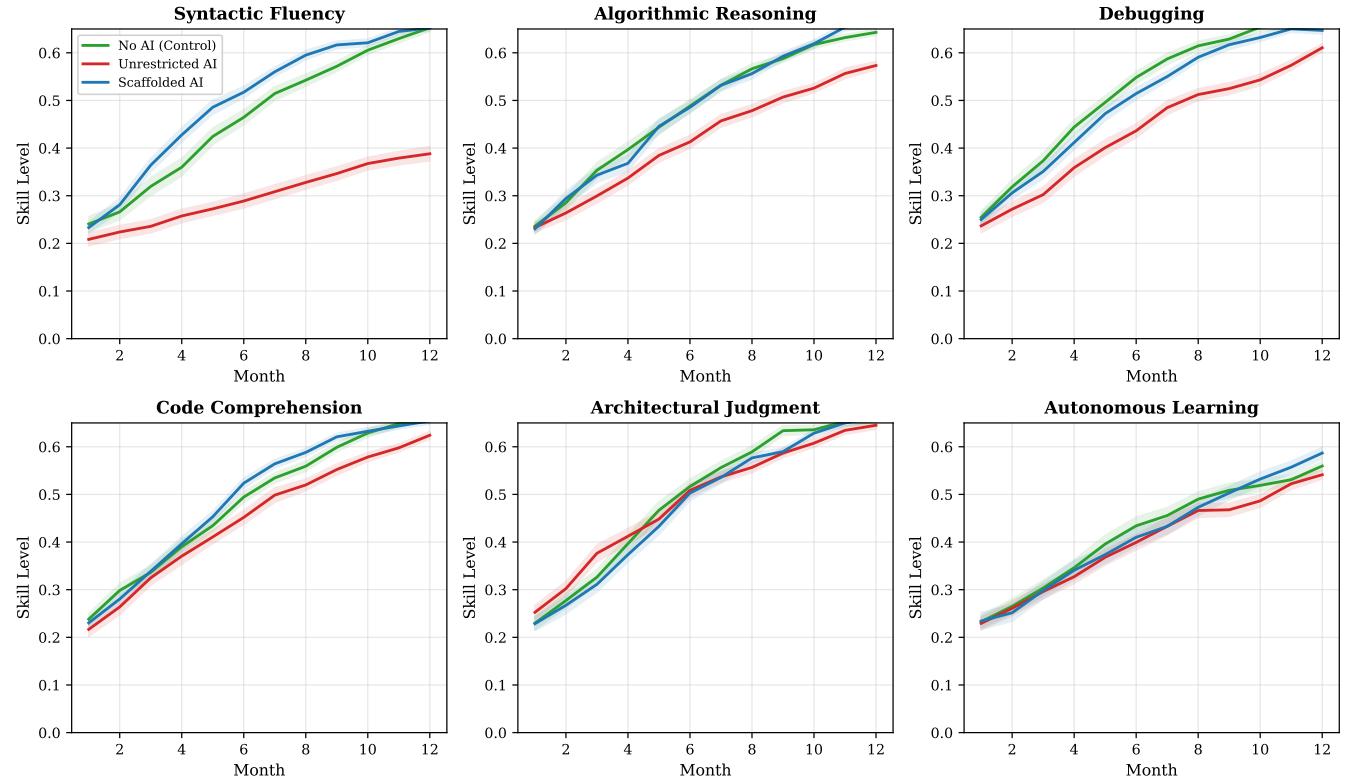


Figure 2: Skill trajectories across six programming dimensions over 12 months. Lines show group means; shaded regions show 95% confidence intervals. The Unrestricted AI condition (red) shows progressively diverging trajectories from Control (green), with the largest gaps in highly automatable dimensions (syntactic fluency, algorithmic reasoning). The Scaffolded AI condition (blue) closely tracks Control across all dimensions.

Table 2: Dimension-specific final skill levels and effect sizes. Cohen's d compares each AI condition against Control; negative values indicate AI-induced skill impairment. p -values from permutation tests (5,000 permutations). Dimensions ordered by AI automation weight (descending).

Dimension	w_i	AI Weight			Final Skill (Mean)		Cohen's d vs. Control	
		Control	Unrest. AI	Scaff. AI	Unrest. (p)	Scaff. (p)		
Syntactic Fluency	0.80	0.653	0.388	0.652	-4.79 (< .001)	-0.01 (0.948)		
Algorithmic Reasoning	0.50	0.643	0.573	0.661	-1.97 (< .001)	+0.54 (< .001)		
Debugging	0.35	0.659	0.611	0.647	-1.32 (< .001)	-0.34 (0.031)		
Code Comprehension	0.25	0.663	0.624	0.654	-1.14 (< .001)	-0.28 (0.075)		
Architectural Judgment	0.15	0.655	0.645	0.661	-0.27 (0.096)	+0.16 (0.341)		
Autonomous Learning	0.10	0.560	0.541	0.587	-0.45 (0.005)	+0.57 (< .001)		

processing depth to achieve net-positive outcomes. At the default unrestricted settings ($\phi = 0.15$, $r = 0.55$), the skill deficit is approximately -0.08 ; achieving net-positive skill formation requires either $\phi > 0.80$ at $r = 0.55$ or reducing r below 0.3 at $\phi = 0.70$. Learning rate and forgetting rate sweeps confirm that the qualitative pattern—unrestricted AI harms skill formation, scaffolded AI preserves it—holds across all tested parameter combinations.

4 DISCUSSION

4.1 The Skill Formation Paradox

Our model predicts a fundamental tension between short-term productivity and long-term skill development. Unrestricted AI use—the default mode in which most novice developers interact with AI tools—produces a large negative effect on skill formation ($d = -0.97$) while simultaneously boosting observable productivity. This *productivity-skill dissociation* creates a systemic risk: organizations

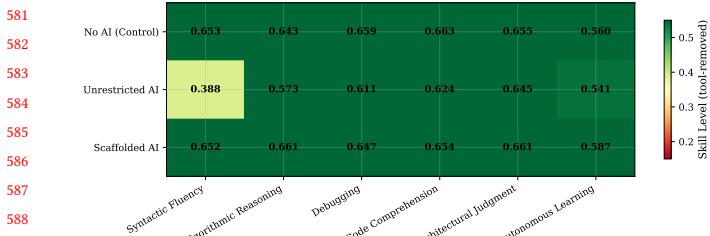


Figure 3: Heatmap of final skill levels by condition and dimension. Warmer colors indicate higher skill. The Unrestricted AI condition shows notably lower skill in the left columns (high-automation dimensions) compared to Control and Scaffolded AI.

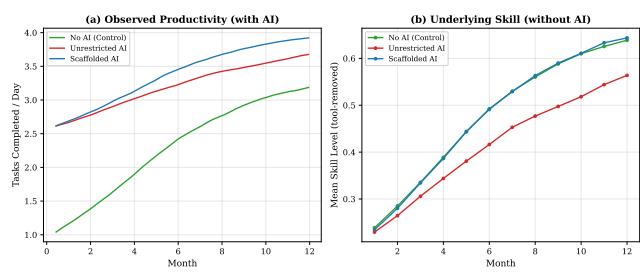


Figure 4: The productivity-skill dissociation. (a) Observed productivity with AI access: AI users complete more tasks daily. (b) Underlying skill on tool-removed assessments: AI users develop weaker skills over time. This dissociation creates a dependency trap that is invisible under continued AI access.

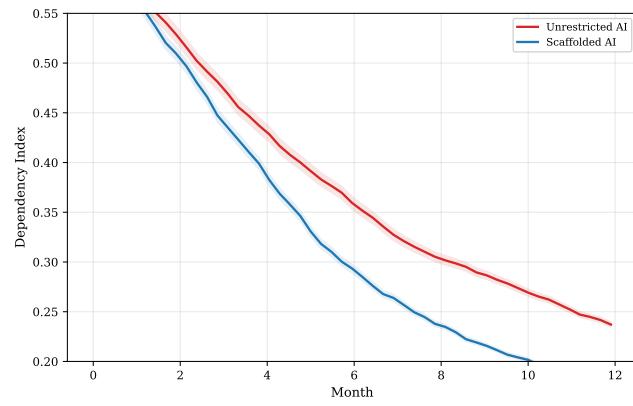


Figure 5: Dependency Index (DI) over 12 months. Higher values indicate greater reliance on AI tools. Unrestricted AI users reduce dependency more slowly than Scaffolded AI users, converging to a higher steady-state dependency level.

optimizing for measurable output will inadvertently produce developers who cannot function without AI scaffolding.

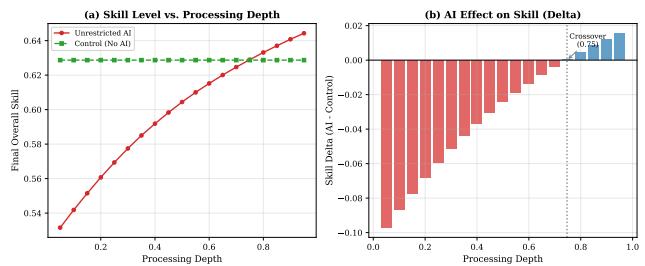


Figure 6: Sensitivity analysis. (a) Final skill levels as a function of cognitive processing depth during AI-assisted work. (b) Skill delta (AI minus Control): the crossover from negative to positive occurs at processing depth ≈ 0.75 . Below this threshold, AI harms skill formation; above it, AI helps.

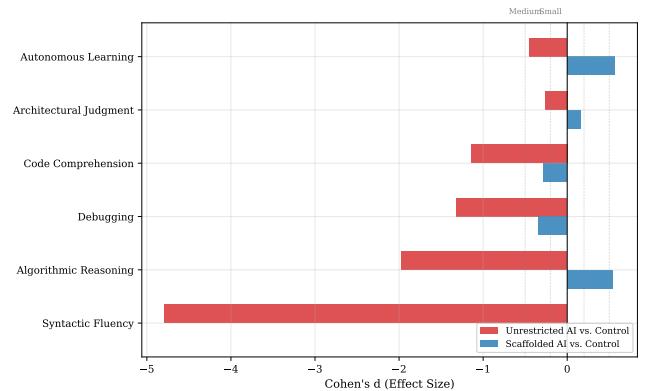


Figure 7: Cohen's d effect sizes by dimension. Unrestricted AI (red) shows consistently negative effects, largest for highly automatable skills. Scaffolded AI (blue) shows near-zero effects across most dimensions, with modest positive effects for algorithmic reasoning and autonomous learning.

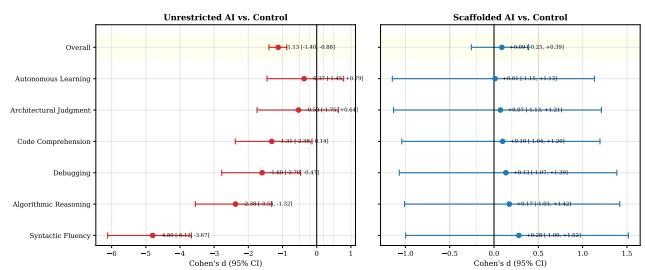


Figure 8: Forest plot of Cohen's d effect sizes with 95% bootstrap confidence intervals (50 seeds). Unrestricted AI (red) shows consistently negative effects across dimensions, with the most robust impairment in syntactic fluency. Scaffolded AI (blue) shows confidence intervals overlapping zero for all dimensions.

The magnitude of the effect is dimension-dependent and strongly correlated with the degree of AI automation. Syntactic fluency—the skill most readily automated by current AI tools—shows the

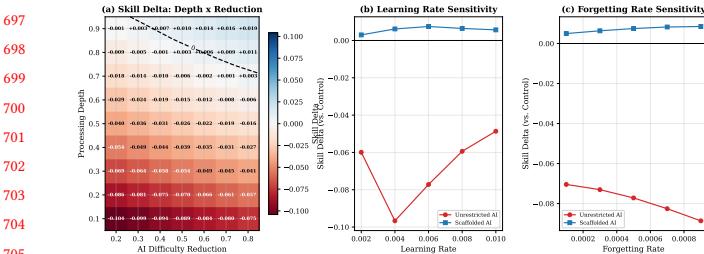


Figure 9: Multi-parameter sensitivity heatmap. Skill delta (AI minus Control) as a function of processing depth (ϕ , y-axis) and AI difficulty reduction (r , x-axis). Blue regions indicate net skill harm; red regions indicate net skill benefit. The white contour marks the zero-crossing boundary. Default unrestricted AI parameters ($\phi = 0.15$, $r = 0.55$) fall deep in the harm zone.

largest impairment ($d = -4.79$). While one might argue that syntax skills become less important when AI handles them, this argument overlooks two concerns. First, syntactic fluency is foundational; debugging, code review, and architectural reasoning all require the ability to read and write code fluently. Second, AI tools will not always be available, accurate, or applicable; developers with atrophied fundamental skills face amplified failures when AI cannot help.

4.2 Scaffolding as a Solution

The Scaffolded AI condition demonstrates that the negative skill effect is not inherent to AI tool use but rather to the *mode of engagement*. When novices are required to actively process AI output—reading, modifying, and explaining generated code before incorporating it—skill development proceeds at nearly the same rate as unaided practice ($d = +0.10$). This finding aligns with prior work on active learning and desirable difficulty [4] and suggests concrete design interventions:

- **Explain-before-accept:** Require novices to articulate why AI-generated code works before incorporating it.
- **Modification prompts:** Present AI suggestions in a form that requires adaptation rather than verbatim acceptance.
- **Interleaved practice:** Periodically disable AI assistance to force unscattered practice.
- **Progressive withdrawal:** Gradually reduce AI assistance as skill levels increase, analogous to training wheels.

4.3 The Crossover Threshold

The sensitivity analysis identifies a processing depth threshold of $\phi \approx 0.75$ at which AI transitions from skill-harming to skill-enhancing. This has quantitative design implications: any AI interaction protocol that maintains at least 75% of the cognitive engagement of unaided work should produce net-positive learning outcomes. Current AI tools that offer frictionless code completion (estimated $\phi \approx 0.15$) are far below this threshold, while structured engagement protocols can approach or exceed it.

4.4 Limitations

Our findings are based on a computational model, not empirical data from human participants. The model makes assumptions about cognitive architecture (learning rates, forgetting dynamics, transfer structure) that, while grounded in established theory, may not precisely match real-world learning. Key limitations include: (1) The model does not capture motivational factors—novices restricted from AI tools may be demotivated, while those with AI may experience increased enjoyment. (2) The task environment is simplified; real software development involves social interaction, code review, and collaborative problem-solving that may modify learning dynamics. (3) The processing depth parameter, while theoretically motivated, conflates multiple cognitive processes into a single scalar. (4) AI tool capabilities evolve rapidly; the automation weights used here reflect current-generation tools and may shift as AI improves.

These limitations are inherent to the computational modeling approach but are offset by its strengths: the ability to generate precise, testable predictions; systematic exploration of parameter space; and low cost relative to longitudinal human studies.

4.5 Empirical Validation

Our model generates several testable predictions for empirical studies:

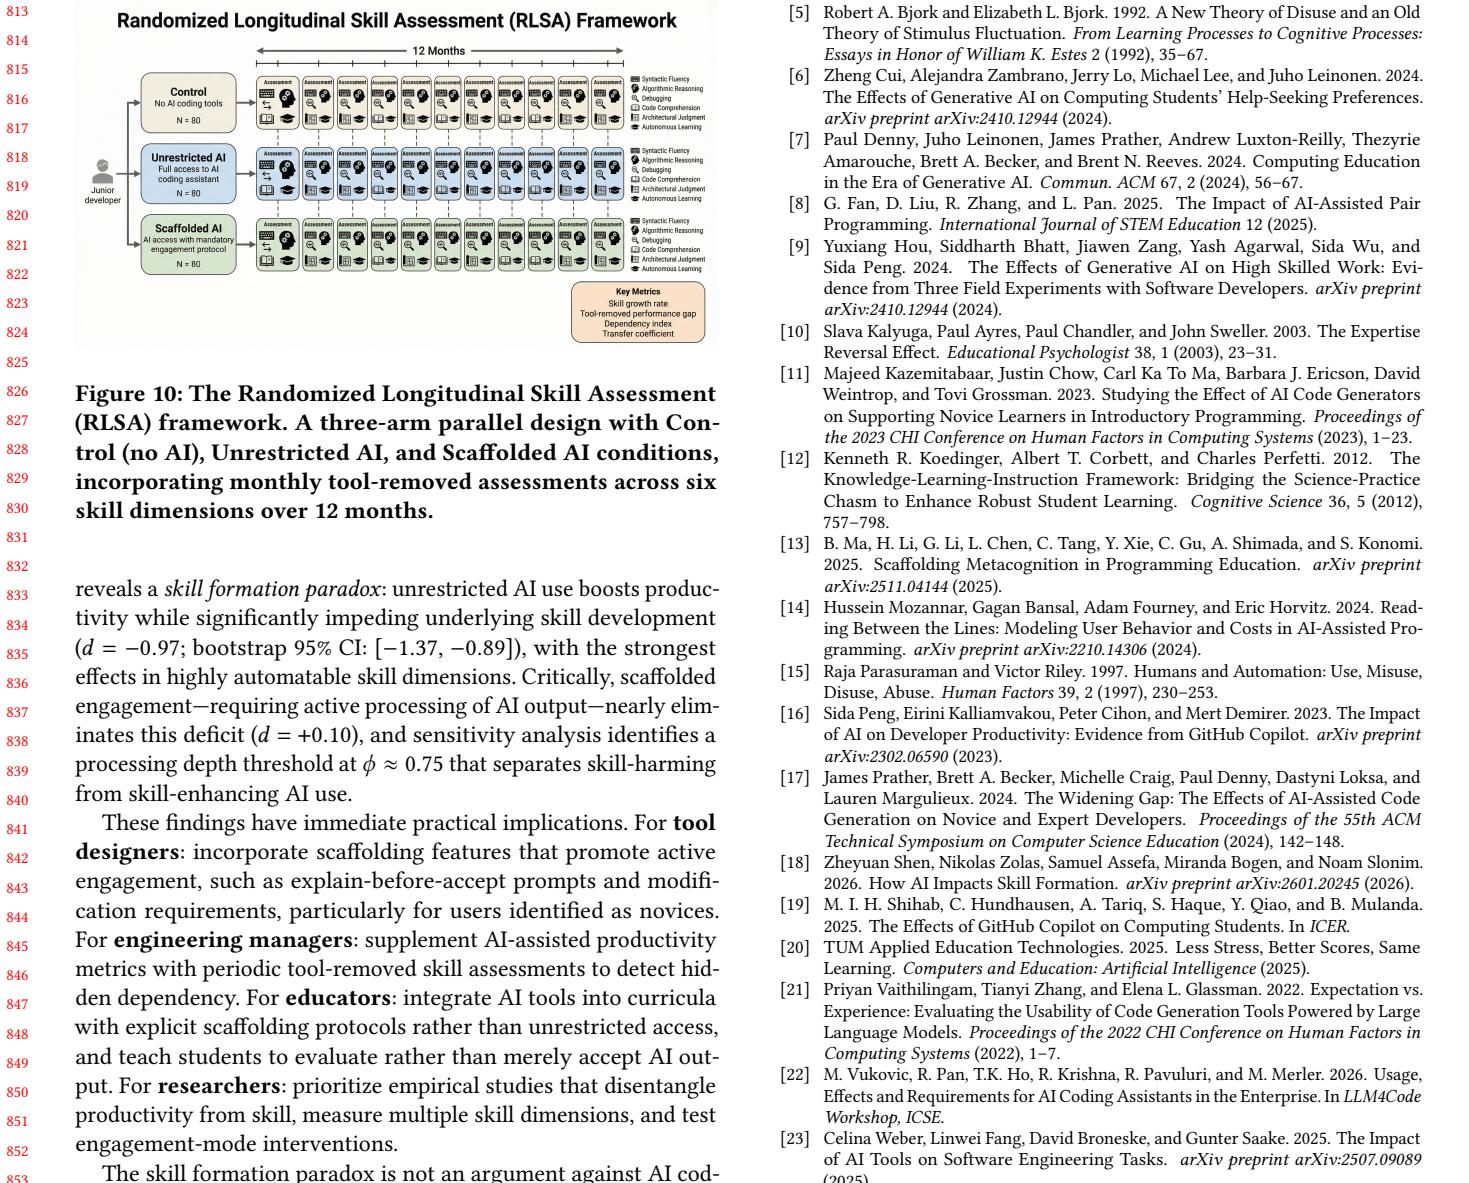
- (1) **Dimension-specificity:** The AI-induced skill deficit should be largest for syntactic and algorithmic skills, smallest for architectural and meta-cognitive skills.
- (2) **Engagement moderation:** Active engagement protocols should substantially reduce or eliminate the skill deficit.
- (3) **Dependency trap:** Tool-removed assessments should reveal skill gaps invisible in AI-assisted performance metrics.
- (4) **Threshold effect:** Interventions increasing processing depth above ~ 0.75 should flip the AI effect from negative to positive.

Emerging empirical evidence is qualitatively consistent with these predictions. Shen et al. [18] report a 17% skill reduction ($d = 0.738$) in a 52-participant RCT with an asyncio programming library—the same direction and approximate magnitude as our model’s prediction of a 16.4% growth deficit ($d = -0.97$). The TUM three-arm trial [20] finds performance boosts with no learning gain, directly paralleling our predicted productivity–skill dissociation. Becker et al. [3] find that experienced developers are actually *slowed* by AI tools, consistent with our model’s prediction that the productivity benefit is largest for novices (where AI bridges the largest skill gap) and may invert for experts.

We recommend a Randomized Longitudinal Skill Assessment (RLSA) design—a 12-month, three-arm trial with monthly tool-removed assessments across all six skill dimensions—as the empirical study most directly suited to testing these predictions. Figure 10 illustrates the proposed RLSA framework, showing the three-arm randomized design and assessment structure.

5 CONCLUSION

We have presented a computational cognitive model that addresses the open question of how AI coding tools affect novice developer skill formation. Our simulation of 240 developers over 12 months



REFERENCES

[1] John R. Anderson. 1982. Acquisition of Cognitive Skill. *Psychological Review* 89, 4 (1982), 369–406.

[2] Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Ozge Kabakci, and Rei Maramian. 2024. Generative AI Can Harm Learning. *arXiv preprint arXiv:2410.15745* (2024).

[3] Lara Becker, Alexander Rush, et al. 2025. Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity. *arXiv preprint arXiv:2507.09089* (2025).

[4] Robert A. Bjork. 1994. Memory and Metamemory Considerations in the Training of Human Beings. In *Metacognition: Knowing About Knowing*. MIT Press, 185–205.

[5] Robert A. Bjork and Elizabeth L. Bjork. 1992. A New Theory of Disuse and an Old Theory of Stimulus Fluctuation. *From Learning Processes to Cognitive Processes: Essays in Honor of William K. Estes* 2 (1992), 35–67.

[6] Zheng Cui, Alejandra Zambrano, Jerry Lo, Michael Lee, and Juho Leinonen. 2024. The Effects of Generative AI on Computing Students' Help-Seeking Preferences. *arXiv preprint arXiv:2410.12944* (2024).

[7] Paul Denny, Juho Leinonen, James Prather, Andrew Luxton-Reilly, Thezyrie Amarouche, Brett A. Becker, and Brent N. Reeves. 2024. Computing Education in the Era of Generative AI. *Commun. ACM* 67, 2 (2024), 56–67.

[8] G. Fan, D. Liu, R. Zhang, and L. Pan. 2025. The Impact of AI-Assisted Pair Programming. *International Journal of STEM Education* 12 (2025).

[9] Yuxiang Hou, Siddharth Bhatt, Jiawen Zang, Yash Agarwal, Sida Wu, and Sida Peng. 2024. The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers. *arXiv preprint arXiv:2410.12944* (2024).

[10] Slava Kalyuga, Paul Ayres, Paul Chandler, and John Sweller. 2003. The Expertise Reversal Effect. *Educational Psychologist* 38, 1 (2003), 23–31.

[11] Majeed Kazemitaabar, Justin Chow, Carl Ka To Ma, Barbara J. Ericson, David Weinrop, and Tovi Grossman. 2023. Studying the Effect of AI Code Generators on Supporting Novice Learners in Introductory Programming. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023), 1–23.

[12] Kenneth R. Koedinger, Albert T. Corbett, and Charles Perfetti. 2012. The Knowledge-Learning-Instruction Framework: Bridging the Science-Practice Chasm to Enhance Robust Student Learning. *Cognitive Science* 36, 5 (2012), 757–798.

[13] B. Ma, H. Li, G. Li, L. Chen, C. Tang, Y. Xie, C. Gu, A. Shimada, and S. Konomi. 2025. Scaffolding Metacognition in Programming Education. *arXiv preprint arXiv:2511.04144* (2025).

[14] Hussein Mozzanfar, Gagan Bansal, Adam Fournier, and Eric Horvitz. 2024. Reading Between the Lines: Modeling User Behavior and Costs in AI-Assisted Programming. *arXiv preprint arXiv:2210.14306* (2024).

[15] Raja Parasuraman and Victor Riley. 1997. Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors* 39, 2 (1997), 230–253.

[16] Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. *arXiv preprint arXiv:2302.06590* (2023).

[17] James Prather, Brett A. Becker, Michelle Craig, Paul Denny, Dastyni Loksa, and Lauren Margulieux. 2024. The Widening Gap: The Effects of AI-Assisted Code Generation on Novice and Expert Developers. *Proceedings of the 55th ACM Technical Symposium on Computer Science Education* (2024), 142–148.

[18] Zheyuan Shen, Nikolas Zolas, Samuel Assefa, Miranda Bogen, and Noam Slonim. 2026. How AI Impacts Skill Formation. *arXiv preprint arXiv:2601.20245* (2026).

[19] M. I. H. Shihab, C. Hundhausen, A. Tariq, S. Haque, Y. Qiao, and B. Mulanda. 2025. The Effects of GitHub Copilot on Computing Students. In *ICER*.

[20] TUM Applied Education Technologies. 2025. Less Stress, Better Scores, Same Learning. *Computers and Education: Artificial Intelligence* (2025).

[21] Priyan Vaithilingam, Tianyi Zhang, and Elena L. Glassman. 2022. Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (2022), 1–7.

[22] M. Vukovic, R. Pan, T.K. Ho, R. Krishna, R. Pavuluri, and M. Merler. 2026. Usage, Effects and Requirements for AI Coding Assistants in the Enterprise. In *LLM4Code Workshop, ICSE*.

[23] Celina Weber, Linwei Fang, David Brioneske, and Gunter Saake. 2025. The Impact of AI Tools on Software Engineering Tasks. *arXiv preprint arXiv:2507.09089* (2025).