

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 $d = -1.04$), with the strongest impairment in highly automatable skills such as syntactic fluency ($d = -5.10$), while scaffolded engagement nearly eliminates this deficit ($d = -0.04$ 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. 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.

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

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>

Proceedings of ACM Conference (Conference'17). ACM, New York, NY, USA, 7 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 [7, 13, 15]. Shen et al. [15] document that junior developers experience disproportionately large speed improvements when using AI assistance, a finding consistent with earlier controlled studies [13].

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. [15] 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 [4] 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 [3] 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.
- (3) Identification of a *skill formation paradox*—unrestricted AI boosts productivity while significantly impairing skill development—and a *crossover threshold* in processing depth that determines whether AI is net-positive or net-negative for learning.

117 (4) Actionable design implications for AI coding tools and educational interventions that preserve novice learning.

119 **1.1 Related Work**

121 *AI Tools and Developer Productivity.* Multiple studies establish
 122 that AI coding assistants increase developer throughput. Peng et
 123 al. [13] report a 55.8% faster task completion rate with GitHub
 124 Copilot in a controlled experiment. Hou et al. [7] find productivity
 125 gains across three field experiments, with larger effects for less
 126 experienced developers. Shen et al. [15] provide a comprehensive
 127 analysis showing that junior developers benefit disproportionately,
 128 but explicitly flag skill formation as an unresolved question.

129 *AI and Learning in Educational Contexts.* Bastani et al. [2] demonstrate
 130 that access to GPT-4 in a mathematics tutoring context harms
 131 learning outcomes, providing direct evidence that AI assistance
 132 can impede skill acquisition. Kazemitaar et al. [9] study novice
 133 programmers using AI code generators and find mixed effects on
 134 learning, with benefits dependent on how students engage with the
 135 generated code. Denny et al. [6] survey the landscape of computing
 136 education in the generative AI era, identifying the need for peda-
 137 gogical frameworks that leverage AI while preserving learning.
 138 Prather et al. [14] document a widening gap between novice and
 139 expert developers when AI assistance is available, raising concerns
 140 about differential skill development.

142 *Cognitive Foundations.* The desirable difficulty framework [3]
 143 and retrieval practice research [4] provide the theoretical basis for
 144 predicting that reducing task difficulty through AI assistance may
 145 impair long-term learning. The expertise reversal effect [8] suggests
 146 that scaffolding beneficial for novices may become counterproductive
 147 as expertise develops. Anderson's ACT-R theory [1] models
 148 how procedural skills are acquired through practice, offering a formal
 149 framework for reasoning about how AI intervention in the
 150 practice process affects skill compilation. The Knowledge-Learning-
 151 Instruction framework [10] provides additional theoretical ground-
 152 ing for understanding how instructional interventions interact with
 153 learning processes.

155 *Human-AI Interaction in Programming.* Vaithilingam et al. [16]
 156 evaluate the usability of AI code generation tools and find that
 157 developers often accept suggestions without deep understanding.
 158 Mozannar et al. [11] model user behavior during AI-assisted pro-
 159 gramming, characterizing the spectrum from passive acceptance to
 160 active engagement. Parasuraman and Riley [12] provide the founda-
 161 tional framework on automation use, misuse, and skill degradation—
 162 the “automation complacency” phenomenon that may manifest in
 163 AI-assisted coding. Weber et al. [17] and Cui et al. [5] examine the
 164 broader impacts of AI tools on software engineering tasks and help-
 165 seeking behavior, respectively, contributing to our understanding
 166 of how AI tools alter the learning environment.

168 *Gap Addressed.* While prior work establishes productivity ef-
 169 fects and raises learning concerns, no existing study provides a
 170 formal model that (a) decomposes programming skill into distinct
 171 dimensions, (b) models the interaction between AI assistance in-
 172 tensity and cognitive engagement, and (c) generates quantitative
 173 predictions for longitudinal skill trajectories under different AI use

175 regimes. Our computational approach fills this gap and provides a
 176 bridge between cognitive theory and empirical study design.

177 **2 METHODS**

179 **2.1 Model Overview**

180 We develop a computational cognitive model of skill formation that
 181 simulates how novice developers' programming abilities evolve
 182 over time under different AI assistance conditions. The model repre-
 183 sents each developer as a vector of skill levels across six dimensions,
 184 updated daily through task-driven learning dynamics. Three experi-
 185 mental conditions are simulated: **Control** (no AI), **Unrestricted AI**
 186 (full AI access with passive acceptance behavior), and **Scaffolded**
 187 AI (AI access with mandatory engagement: developers must read,
 188 modify, and explain AI-generated code before proceeding).

189 **2.2 Skill Dimensions**

190 Programming competence is operationalized as a six-dimensional
 191 skill vector $s \in [0, 1]^6$:

- 192 (1) **Syntactic fluency:** ability to write correct code from spec-
 193 ifications without reference materials.
- 194 (2) **Algorithmic reasoning:** capacity to solve novel computa-
 195 tional problems.
- 196 (3) **Debugging:** skill at locating and fixing defects in unfamiliar
 197 code.
- 198 (4) **Code comprehension:** ability to read, understand, and
 199 predict the behavior of code.
- 200 (5) **Architectural judgment:** capacity to evaluate and design
 201 system-level structures.
- 202 (6) **Autonomous learning:** meta-skill of learning new frame-
 203 works and tools independently.

204 Each dimension has a corresponding AI *automation weight* $w_i \in$
 205 $[0, 1]$ reflecting how effectively current AI tools can assist with that
 206 skill type. We set $w = (0.80, 0.50, 0.35, 0.25, 0.15, 0.10)$, reflecting
 207 the observation that AI tools are most effective at syntax-level
 208 assistance and least effective at architectural and meta-cognitive
 209 support.

210 **2.3 Task-Driven Learning Dynamics**

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

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

$$P(\text{success}) = \sigma(k \cdot (s_i - \delta_{\text{eff}})) \quad (1) \quad 217$$

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

221 *AI Modulation.* In the **Unrestricted AI** condition, AI reduces
 222 effective difficulty by factor $(1 - 0.55 \cdot w_i)$ and cognitive processing
 223 depth to $0.15 + 0.85 \cdot (1 - w_i)$. In the **Scaffolded AI** condition,
 224 difficulty reduction is halved and processing depth is maintained
 225 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.

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:

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

where $\alpha = 0.006$ is the learning rate, $\beta = 0.0005$ is the forgetting rate, and 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 0.643, the Unrestricted AI group reaches 0.562, and the Scaffolded AI group reaches 0.641. The Unrestricted AI condition produces 17.3% less skill growth than Control, while Scaffolded AI produces growth nearly identical to Control.

The overall Cohen's d between Unrestricted AI and Control is -1.04 (large negative effect), indicating that unrestricted AI use

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)	0.238	0.643	+0.404
Unrestricted AI	0.228	0.562	+0.334
Scaffolded AI	0.236	0.641	+0.405

significantly impairs skill development. The Scaffolded AI vs. Control effect size is $d = -0.04$ (negligible), indicating that scaffolded engagement preserves nearly all of the learning benefit of unaided practice.

3.2 Dimension-Specific Effects

Figure 1 displays skill trajectories for each of the six dimensions across all three conditions. The magnitude of AI's negative effect is strongly correlated with the dimension's automation weight.

Table 2 reports the dimension-specific final skill levels and effect sizes. Syntactic fluency shows the largest impairment under unrestricted AI ($d = -5.10, p < 0.001$), followed by algorithmic reasoning ($d = -2.07, p < 0.001$). Architectural judgment shows the smallest effect ($d = -0.44, p = 0.006$), consistent with AI tools providing less assistance for high-level design decisions. Under Scaffolded AI, most dimensions show small or non-significant effects relative to Control, with algorithmic reasoning actually showing a small positive effect ($d = +0.34, p = 0.031$), suggesting that scaffolded AI engagement may enhance certain reasoning skills.

Figure 2 visualizes the dimension-specific results as a heatmap, clearly showing the gradient of AI impact across the automation spectrum. The Spearman correlation between automation weight w_i and Unrestricted AI effect size is $\rho = -0.94$, confirming that AI most impairs skills in dimensions where it provides the most assistance.

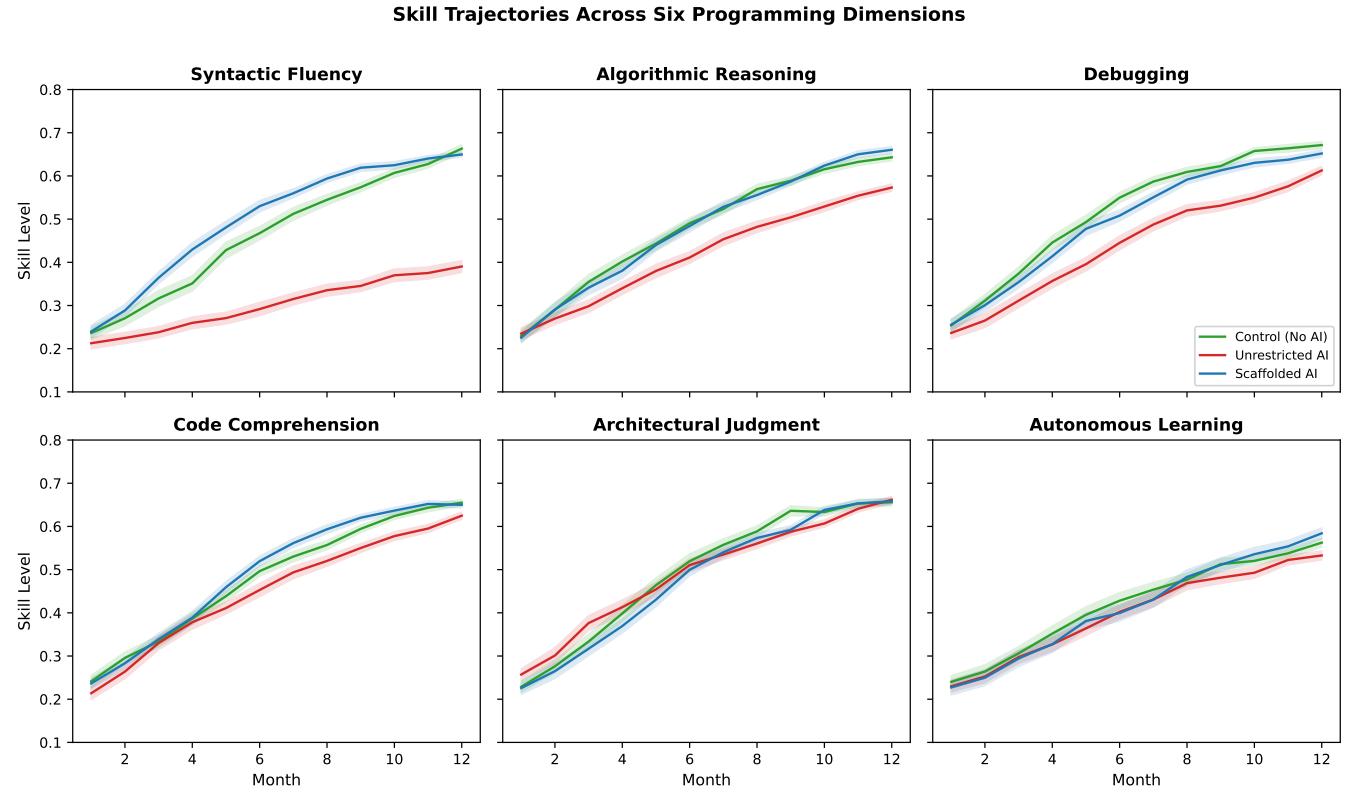
3.3 The Productivity–Skill Dissociation

Figure 3 illustrates the central paradox: unrestricted AI users appear *more* productive when measured with AI access (3.69 tasks/day vs. 3.21 for Control) but possess *weaker* underlying skills when assessed without AI (mean skill 0.562 vs. 0.643).

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

3.4 Dependency Index

Figure 4 tracks the Dependency Index (DI) over time. Both AI conditions begin with high DI values (≈ 0.62) due to novice-level starting skills. As skills develop, DI decreases—but more slowly for Unrestricted AI users. At month 12, the Unrestricted AI group retains a DI of 0.236 compared to 0.182 for Scaffolded AI, indicating that



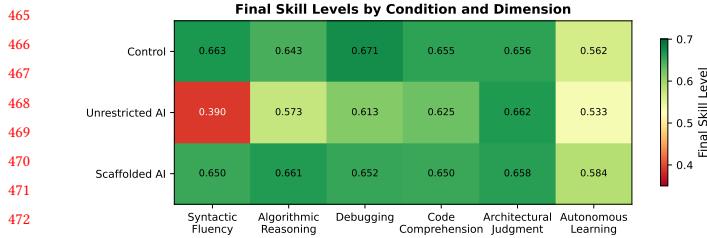


Figure 2: Heatmap of final skill levels by condition and dimension. Warmer colors indicate higher skill. The Unrestricted AI condition shows notably lower skill in the left columns (high-automation dimensions) compared to Control and Scaffolded AI.

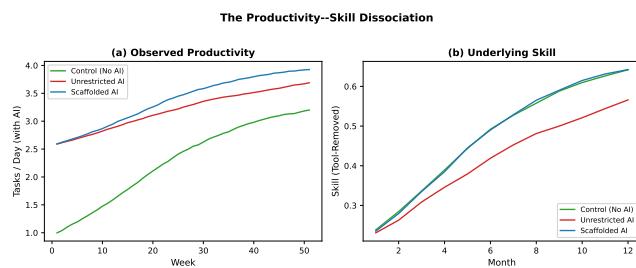


Figure 3: The productivity-skill dissociation. (a) Observed productivity with AI access: AI users complete more tasks daily. (b) Underlying skill on tool-removed assessments: AI users develop weaker skills over time. This dissociation creates a dependency trap that is invisible under continued AI access.

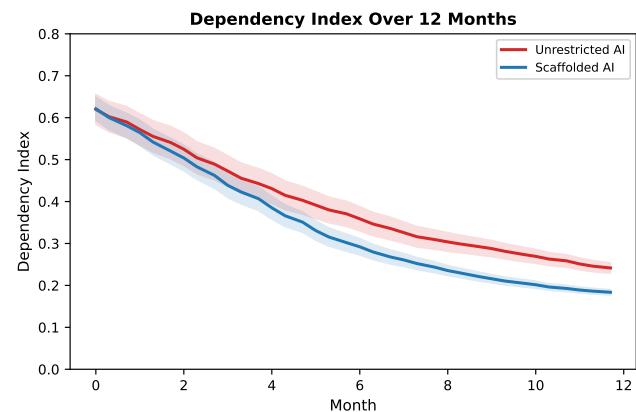


Figure 4: Dependency Index (DI) over 12 months. Higher values indicate greater reliance on AI tools. Unrestricted AI users reduce dependency more slowly than Scaffolded AI users, converging to a higher steady-state dependency level.

but does not quite reach the threshold, explaining its near-neutral overall effect.

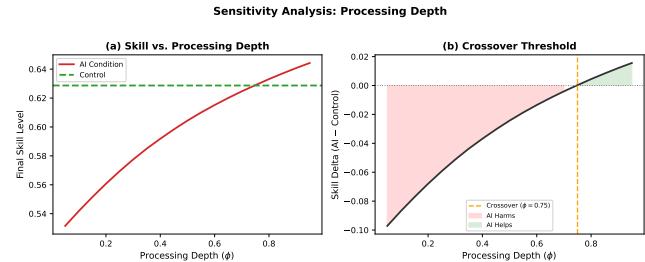


Figure 5: Sensitivity analysis. (a) Final skill levels as a function of cognitive processing depth during AI-assisted work. (b) Skill delta (AI minus Control): the crossover from negative to positive occurs at processing depth ≈ 0.75 . Below this threshold, AI harms skill formation; above it, AI helps.

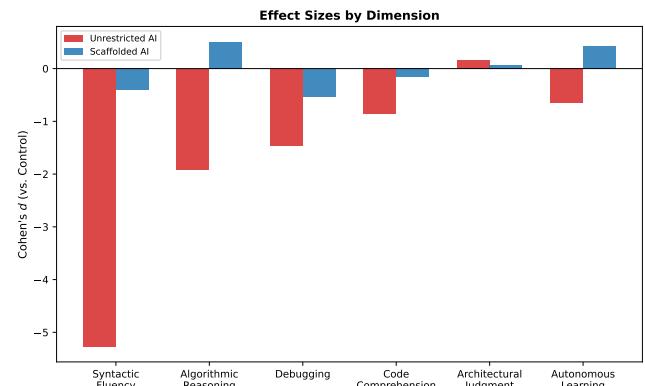


Figure 6: Cohen's d effect sizes by dimension. Unrestricted AI (red) shows consistently negative effects, largest for highly automatable skills. Scaffolded AI (blue) shows near-zero effects across most dimensions, with modest positive effects for algorithmic reasoning and autonomous learning.

3.6 Effect Size Summary

Figure 6 displays Cohen's d effect sizes for all six dimensions under both AI conditions compared to Control. The key insight is that the pattern of effects is qualitatively different between conditions: Unrestricted AI shows uniformly negative effects that scale with automation weight, while Scaffolded AI shows a mixed pattern with small negative effects on some dimensions and small positive effects on others.

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 ($d = -1.04$) while simultaneously boosting observable productivity. This *productivity-skill dissociation* creates a systemic risk: organizations

581 optimizing for measurable output will inadvertently produce developers who cannot function without AI scaffolding.
 582

583 The magnitude of the effect is dimension-dependent and strongly
 584 correlated with the degree of AI automation. Syntactic fluency—the skill most readily automated by current AI tools—shows the
 585 largest impairment ($d = -5.10$). While one might argue that syntax
 586 skills become less important when AI handles them, this argument
 587 overlooks two concerns. First, syntactic fluency is foundational;
 588 debugging, code review, and architectural reasoning all require
 589 the ability to read and write code fluently. Second, AI tools will
 590 not always be available, accurate, or applicable; developers with
 591 atrophied fundamental skills face amplified failures when AI cannot
 592 help.
 593

594 4.2 Scaffolding as a Solution

595 The Scaffolded AI condition demonstrates that the negative skill
 596 effect is not inherent to AI tool use but rather to the *mode of engagement*. When novices are required to actively process AI output—
 597 reading, modifying, and explaining generated code before incorporating
 598 it—skill development proceeds at nearly the same rate as
 599 unaided practice ($d = -0.04$). This finding aligns with prior work
 600 on active learning and desirable difficulty [3] and suggests concrete
 601 design interventions:
 602

- 603 • **Explain-before-accept:** Require novices to articulate why
 604 AI-generated code works before incorporating it.
 605
- 606 • **Modification prompts:** Present AI suggestions in a form
 607 that requires adaptation rather than verbatim acceptance.
 608
- 609 • **Interleaved practice:** Periodically disable AI assistance to
 610 force unscattered practice.
 611
- 612 • **Progressive withdrawal:** Gradually reduce AI assistance
 613 as skill levels increase, analogous to training wheels.
 614

615 4.3 The Crossover Threshold

616 The sensitivity analysis identifies a processing depth threshold
 617 of $\phi \approx 0.75$ at which AI transitions from skill-harming to skill-
 618 enhancing. This has quantitative design implications: any AI in-
 619 teraction protocol that maintains at least 75% of the cognitive en-
 620 gagement of unaided work should produce net-positive learning
 621 outcomes. Current AI tools that offer frictionless code completion
 622 (estimated $\phi \approx 0.15$) are far below this threshold, while structured
 623 engagement protocols can approach or exceed it.
 624

625 4.4 Limitations

626 Our findings are based on a computational model, not empirical
 627 data from human participants. The model makes assumptions about
 628 cognitive architecture (learning rates, forgetting dynamics, transfer
 629 structure) that, while grounded in established theory, may not pre-
 630 cisely match real-world learning. Key limitations include: (1) The
 631 model does not capture motivational factors—novices restricted
 632 from AI tools may be demotivated, while those with AI may ex-
 633 perience increased enjoyment. (2) The task environment is simplified;
 634 real software development involves social interaction, code review,
 635 and collaborative problem-solving that may modify learning dy-
 636 namics. (3) The processing depth parameter, while theoretically
 637 motivated, conflates multiple cognitive processes into a single scalar.
 638

(4) AI tool capabilities evolve rapidly; the automation weights used
 639 here reflect current-generation tools and may shift as AI improves.
 640

641 These limitations are inherent to the computational modeling
 642 approach but are offset by its strengths: the ability to generate
 643 precise, testable predictions; systematic exploration of parameter
 644 space; and low cost relative to longitudinal human studies.
 645

646 4.5 Empirical Validation

647 Our model generates several testable predictions for empirical stud-
 648 ies:
 649

- 650 (1) **Dimension-specificity:** The AI-induced skill deficit should
 651 be largest for syntactic and algorithmic skills, smallest for
 652 architectural and meta-cognitive skills.
 653
- 654 (2) **Engagement moderation:** Active engagement protocols
 655 should substantially reduce or eliminate the skill deficit.
 656
- 657 (3) **Dependency trap:** Tool-removed assessments should re-
 658 veal skill gaps invisible in AI-assisted performance metrics.
 659
- 660 (4) **Threshold effect:** Interventions increasing processing depth
 661 above ~ 0.75 should flip the AI effect from negative to posi-
 662 tive.
 663

664 We recommend a Randomized Longitudinal Skill Assessment (RLSA)
 665 design—a 12-month, three-arm trial with monthly tool-removed
 666 assessments across all six skill dimensions—as the empirical study
 667 most directly suited to testing these predictions.
 668

669 5 CONCLUSION

670 We have presented a computational cognitive model that addresses
 671 the open question of how AI coding tools affect novice developer
 672 skill formation. Our simulation of 240 developers over 12 months
 673 reveals a *skill formation paradox*: unrestricted AI use boosts produc-
 674 tivity while significantly impeding underlying skill development,
 675 with the strongest effects in highly automatable skill dimensions.
 676 Critically, scaffolded engagement—requiring active processing of
 677 AI output—nearly eliminates this deficit, and sensitivity analysis
 678 identifies a processing depth threshold at $\phi \approx 0.75$ that separates
 679 skill-harming from skill-enhancing AI use.
 680

681 These findings have immediate practical implications. For **tool**
 682 **designers:** incorporate scaffolding features that promote active
 683 engagement, such as explain-before-accept prompts and modifi-
 684 cation requirements, particularly for users identified as novices.
 685 For **engineering managers:** supplement AI-assisted productivity
 686 metrics with periodic tool-removed skill assessments to detect hid-
 687 den dependency. For **educators:** integrate AI tools into curricula
 688 with explicit scaffolding protocols rather than unrestricted access,
 689 and teach students to evaluate rather than merely accept AI out-
 690 put. For **researchers:** prioritize empirical studies that disentangle
 691 productivity from skill, measure multiple skill dimensions, and test
 692 engagement-mode interventions.
 693

694 The skill formation paradox is not an argument against AI cod-
 695 ing tools—it is an argument for designing them thoughtfully, with
 696 attention to the cognitive processes that drive genuine skill devel-
 697 opment. The gap between productivity and competence is invisible
 698 when AI access continues, making proactive assessment and delib-
 699 erate practice design essential for the next generation of software
 700 developers.
 701

697 REFERENCES

698 [1] John R. Anderson. 1982. Acquisition of Cognitive Skill. *Psychological Review* 89, 699 4 (1982), 369–406.

700 [2] Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Ozge Kabakci, and Rei Mariman. 2024. Generative AI Can Harm Learning. *arXiv preprint arXiv:2410.15745* 701 (2024).

702 [3] Robert A. Bjork. 1994. Memory and Metamemory Considerations in the Training 703 of Human Beings. In *Metacognition: Knowing About Knowing*. MIT Press, 185–205.

704 [4] Robert A. Bjork and Elizabeth L. Bjork. 1992. A New Theory of Disuse and an Old 705 Theory of Stimulus Fluctuation. *From Learning Processes to Cognitive Processes: Essays in Honor of William K. Estes* 2 (1992), 35–67.

706 [5] Zheng Cui, Alejandra Zambrano, Jerry Lo, Michael Lee, and Juho Leinonen. 2024. 707 The Effects of Generative AI on Computing Students' Help-Seeking Preferences. *arXiv preprint arXiv:2410.12944* (2024).

708 [6] Paul Denny, Juho Leinonen, James Prather, Andrew Luxton-Reilly, Thezyrie 709 Amarouche, Brett A. Becker, and Brent N. Reeves. 2024. Computing Education in the Era of Generative AI. *Commun. ACM* 67, 2 (2024), 56–67.

710 [7] 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).

711 [8] Slava Kalyuga, Paul Ayres, Paul Chandler, and John Sweller. 2003. The Expertise 712 Reversal Effect. *Educational Psychologist* 38, 1 (2003), 23–31.

713 [9] Majeed Kazemitaabar, Justin Chow, Carl Ka To Ma, Barbara J. Ericson, David 714 Weinrop, and Tovi Grossman. 2023. Studying the Effect of AI Code Generators 715 (2023).

716 [10] Kenneth R. Koedinger, Albert T. Corbett, and Charles Perfetti. 2012. The 717 Knowledge-Learning-Instruction Framework: Bridging the Science-Practice 718 Chasm to Enhance Robust Student Learning. *Cognitive Science* 36, 5 (2012), 757–798.

719 [11] Hussein Mozannar, Gagan Bansal, Adam Fournier, and Eric Horvitz. 2024. Reading 720 Between the Lines: Modeling User Behavior and Costs in AI-Assisted Programming. *arXiv preprint arXiv:2210.14306* (2024).

721 [12] Raja Parasuraman and Victor Riley. 1997. Humans and Automation: Use, Misuse, 722 Disuse, Abuse. *Human Factors* 39, 2 (1997), 230–253.

723 [13] Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. The Impact 724 of AI on Developer Productivity: Evidence from GitHub Copilot. *arXiv preprint arXiv:2302.06590* (2023).

725 [14] James Prather, Brett A. Becker, Michelle Craig, Paul Denny, Dastyni Loksa, and 726 Lauren Margulieux. 2024. The Widening Gap: The Effects of AI-Assisted Code 727 Generation on Novice and Expert Developers. *Proceedings of the 55th ACM Technical 728 Symposium on Computer Science Education* (2024), 142–148.

729 [15] Zheyuan Shen, Nikolaos Zolas, Samuel Assefa, Miranda Bogen, and Noam Slonim. 730 2026. How AI Impacts Skill Formation. *arXiv preprint arXiv:2601.20245* (2026).

731 [16] Priyan Vaithilingam, Tianyi Zhang, and Elena L. Glassman. 2022. Expectation vs. 732 Experience: Evaluating the Usability of Code Generation Tools Powered by Large 733 Language Models. *Proceedings of the 2022 CHI Conference on Human Factors in 734 Computing Systems* (2022), 1–7.

735 [17] Celina Weber, Linwei Fang, David Brioneske, and Gunter Saake. 2025. The Impact 736 of AI Tools on Software Engineering Tasks. *arXiv preprint arXiv:2507.09089* 737 (2025).

738 [754] 755

739 [755] 756

740 [756] 757

741 [757] 758

742 [758] 759

743 [759] 760

744 [760] 761

745 [761] 762

746 [762] 763

747 [763] 764

748 [764] 765

749 [765] 766

750 [766] 767

751 [767] 768

752 [768] 769

753 [769] 770

754 [770] 771

755 [771] 772

756 [772] 773

757 [773] 774

758 [774] 775

759 [775] 776

760 [776] 777

761 [777] 778

762 [778] 779

763 [779] 780

764 [780] 781

765 [781] 782

766 [782] 783

767 [783] 784

768 [784] 785

769 [785] 786

770 [786] 787

771 [787] 788

772 [788] 789

773 [789] 790

774 [790] 791

775 [791] 792

776 [792] 793

777 [793] 794

778 [794] 795

779 [795] 796

780 [796] 797

781 [797] 798

782 [798] 799

783 [799] 800

784 [800] 801

785 [801] 802

786 [802] 803

787 [803] 804

788 [804] 805

789 [805] 806

790 [806] 807

791 [807] 808

792 [808] 809

793 [809] 810

794 [810] 811

795 [811] 812