



DOMAIN DEPTH

Depth Is Knowing Why

In a field where answers are free, the scarce skill is explaining them, and noticing when they are wrong.

Recall tells you what the answer is. Depth tells you when it stops being true.

KEY FINDINGS


- Answer-giving AI raised practice scores by 48% and cut later exam scores by 17%; a hint-based tutor avoided the harm.
- Depth is knowing why an answer is right and when it stops being true.
- Explaining and defending answers builds depth; reading them does not.

An answer to almost any cybersecurity question is now a prompt away. That makes recall cheap, and it makes stopping at recall tempting. This essay argues that domain depth has moved: it now lies in the ability to say why an answer is right, under what conditions it fails, and what an attacker would do next. That ability is built differently from recall, and the most convenient uses of AI can quietly erode it.

01 What answer-giving AI does to learning

The clearest evidence comes from a field experiment with nearly a thousand high-school students in Turkey. Bastani and colleagues gave some students a standard chat interface to GPT-4 during maths practice and others a tutor version built to give hints rather than answers. During practice, the standard interface lifted performance by 48 per cent and the tutor by 127 per cent. When the AI was taken away for the exam, students who had used the standard interface scored 17 per cent lower than students who never had AI at all. The tutor's safeguards largely removed that harm [1].

FIGURE 1 Answer-giving AI raised practice scores, then lowered exam scores



Field experiment, about 1,000 high-school maths students in Turkey. Source: Bastani et al., *PNAS* 2025 [1].

The authors' explanation is the important part. Students used the standard tool as a crutch: the answer arrived, the practice score rose, and the understanding that the practice was meant to build never formed. The study concerns mathematics, not security, and generalising from it is an inference. But the mechanism is not specific to maths. Anyone who lets the answer arrive before they have reasoned towards it is practising receiving answers.

02 What depth looks like

Stuart and Hubert Dreyfus described skill as a progression from novice to expert. Novices follow context-free rules; experts see situations whole and act from intuition built across thousands of cases [2]. The movement between those stages is not a matter of collecting more rules. It is a matter of learning when each rule applies and when it does not.

Micheline Chi's work on engagement points the same way. Across many studies, learning rises as activity moves from passive (receiving), to active (manipulating material), to constructive (generating explanations), to interactive (defending those explanations to someone who questions them) [3]. Explaining why an answer is right is constructive work. Being asked to defend it is interactive. Reading the answer is neither.

FIGURE 2 Four ways of engaging, from least to most learning



After Chi and Wylie's ICAP framework [3].

Recall tells you what the answer is. Depth tells you when it stops being true.

Security has its own version of this. David Bianco's Pyramid of Pain ranks the indicators a defender can detect by how much it costs an attacker to change them: a file hash is trivial to alter, an attacker's tactics and procedures are very hard to [4]. A detection built on hashes is recall; a detection built on understanding how an adversary operates is depth, and it is the only kind that keeps working after the attacker adapts.

03 A summary at half past one

ILLUSTRATIVE SCENE

A junior analyst is on shift at a Nairobi bank. An AI triage summary flags a login for a finance clerk from a new country and marks it "likely VPN, benign". Recall agrees: VPNs often cause location anomalies. Depth asks three further questions. Does the bank's own VPN exit in that country? Is tonight the month-end payroll run, when a finance account is most valuable? Did the session go on to request the payments module? The summary is fluent and plausible. Only the analyst who knows why the rule exists can tell whether this is the case where it fails.

04 How OCH builds it

On WhatsApp drills, OCH gives hints rather than answers, deliberately following the design that protected learning in the Turkish study. Learners answer with how confident they are. After a correct answer, they are asked to explain in one line why it is right, and that explanation is reviewed. Drills are levelled from Foundation to Expert so that the conditions get harder as a learner climbs.

It would be dishonest to claim results yet. OCH began recording explanations and confidence on WhatsApp drills on 16 October 2026. The data needed to show whether this builds depth faster than answer-giving does not exist yet.

05 What would change our mind

If learners who are simply given full answers perform as well on unaided, unfamiliar problems as learners who are made to explain, the argument fails. That is testable, and OCH intends to test it by comparing learners who complete explain-backs with those who skip them, on later drills they have never seen. If the AI tools themselves change so that they reliably teach rather than tell, the warning in this essay becomes less urgent, though the case for explanation as practice would stand.

06 What to do on Monday

For learners. After every answer you get right, write one line on why, and one condition under which it would be wrong. When you use AI, ask it for a hint first and the answer only after you have committed to your own.

For mentors and lecturers. Grade the explanation, not only the answer. A correct answer with a weak reason is a result that will not survive the next variation.

For employers. In interviews, follow every correct answer with one question: when would that be the wrong answer? The candidates who can answer it are the ones with depth.

REFERENCES

- [1] Bastani, H. et al. "Generative AI without guardrails can harm learning: Evidence from high school mathematics." *PNAS* 122(26), 2025.
- [2] Dreyfus, S. and Dreyfus, H. "A Five-Stage Model of the Mental Activities Involved in Directed Skill Acquisition." University of California, Berkeley, 1980.
- [3] Chi, M. and Wylie, R. "The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes." *Educational Psychologist* 49(4), 2014.
- [4] Bianco, D. J. "The Pyramid of Pain." 2013.

Drafted with AI assistance and reviewed before publication. External figures were checked against the sources cited; scenes marked illustrative are composites, not reports of real events; OCH figures are dated.