Infrastructure as Curriculum: What Students Learn When AI Works Offline

Sai Gattupalli, Ph.D.
Society & AI Research Group
sai@societyandai.org

***Abstract*—Most AI-in-education deployments assume stable internet access and cloud-based computation, treating infrastructure as a neutral backdrop rather than a pedagogical force. This short paper offers a conceptual contribution rather than a new empirical study. We reframe infrastructural constraint as an instructional signal: when AI runs offline on local devices, students encounter visible limits, negotiate uncertainty, and engage in iterative troubleshooting and collaborative sensemaking. Drawing on observations from offline, on-device AI tutoring across multiple secondary computing contexts [1], we articulate a compact framework—*Infrastructure as Curriculum*—and derive design implications for AI literacy [7] and equitable learning technology integration. Rather than viewing offline AI as a degraded substitute for cloud intelligence, we argue that bounded systems can surface learning practices—verification, authorship, and calibrated trust—that are often muted by always-available tools.**

***Index Terms*—on-device AI, offline learning, AI literacy, equity, infrastructure, computing education, open science**

I. Introduction

Generative AI is entering educational practice through intelligent tutoring, feedback, and interactive support. Many current deployments assume stable connectivity and cloud-based inference as baseline infrastructure. This assumption excludes learners in low-connectivity settings [3] and, more subtly, treats infrastructure as a neutral substrate. In practice, infrastructure is never neutral—it encodes assumptions about users, normalizes certain practices, and renders others invisible [6]. When connectivity is abundant, AI assistance can seem frictionless and authoritative; when connectivity is absent, the system becomes bounded, imperfect, and visibly situated in the classroom.

This paper makes a simple claim: infrastructure is not merely logistical—it can function as curriculum. In offline, on-device deployments, learners encounter constraints (latency, limited model capacity, local-only resources) that reorganize how they ask questions, verify outputs, and collaborate with peers and teachers. These constraints are not only technical; they prompt classroom conversations about evidence, correctness, and responsibility. Rather than asking only how to bring more AI into classrooms, we ask what classrooms learn when AI itself has limits.

Our contribution is conceptual. We (1) introduce *Infrastructure as Curriculum* as a framework for treating infrastructural constraints as pedagogically meaningful signals; (2) describe learning shifts that become salient when AI support is local and bounded; and (3) distill design implications for AI literacy and equitable integration of advanced learning technologies. We synthesize a pattern observed across multiple offline deployments and argue that bounded AI can be designed to make verification, authorship, and calibrated trust part of the learning experience.

II. Background and Ongoing Deployments

In prior work, we demonstrated how an AI system operating without internet access could serve as a peer-programming companion in a secondary computing classroom [1]. The system ran locally on student devices using preloaded, quantized open-weight language models [4], [5], deliberately constrained in model capacity and interaction style. Students requested hints, clarifications, and partial reasoning steps; teachers retained control over task framing and evaluation. Crucially, the absence of connectivity was treated as a design condition, not a deficit.

Since that initial deployment, we have conducted observations across three additional secondary computing classrooms and two informal computing programs (2023–2025). Observations were gathered through structured classroom notes, teacher reflections, and periodic review of interaction logs. Settings varied in device availability and curricular context, but all shared the condition of preloaded, internet-free AI. This is not a controlled study; it is a pattern-oriented account intended to motivate a conceptual reframing.

Several consistent patterns emerged. Students developed a practical sense of what the AI could and could not do: when responses were incomplete or incorrect, they questioned the output, tested it against their own reasoning, or consulted peers. Prompting practices evolved from broad solution-seeking toward specific, task-oriented requests—explaining a line of code, diagnosing an error, or requesting a hint without full completion. Teachers adapted by pausing more frequently to provide shared clarification and model how to evaluate AI-generated suggestions. The offline system functioned less as a

This paper was accepted for presentation as a short paper at the 26th IEEE International Conference on Advanced Learning Technologies (ICALT 2026), Track 1: Technologies for Open Learning and Education (i-OPENLearn). The author was unable to attend due to the absence of institutional affiliation and funding support, and accordingly withdrew the submission from the conference proceedings. In the spirit of open science, this work is released as open access through the Society & AI Research Group (https://societyandai.org), freely available to educators, researchers, developers, and policymakers. The author welcomes correspondence, collaboration, and critique from any quarter—institutional affiliation is not a prerequisite for meaningful scholarly exchange.

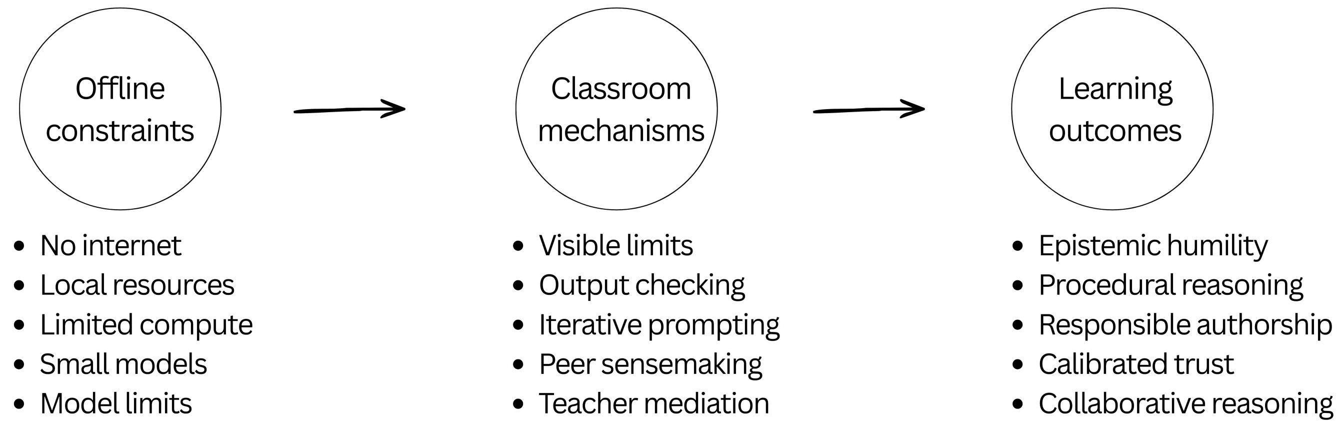

Fig. 1. Conceptual mechanism of *Infrastructure as Curriculum*: offline constraints shape classroom mechanisms (e.g., visible limits and output checking), which in turn surface distinct learning outcomes (e.g., epistemic humility and calibrated trust).

substitute for instruction and more as a catalyst for discussion about reasoning, correctness, and responsibility.

III. Infrastructure as Curriculum: A Conceptual Frame

A. Definition

We define *Infrastructure as Curriculum* as the idea that technological constraints—connectivity, compute, availability, latency—implicitly teach learners about knowledge production, responsibility, and agency. This reframing draws on infrastructural theory in science and technology studies [6], which argues that infrastructure encodes assumptions about who counts as a user, what practices are normalized, and what remains invisible. Applied to AI in education, the question shifts from what AI can do *for* learners to what learners come to understand *about* AI through encountering its material limits.

The claim is not that constraints are inherently beneficial, but that they can be designed and interpreted pedagogically. When limits are visible, learners must decide when to trust the system, how to verify outputs, and how to coordinate human judgment with machine suggestions.

B. Cloud AI vs. Offline On-Device AI

Cloud-based AI systems normalize always-available assistance, rapid responses, and perceived comprehensiveness—conditions that can encourage deference over deliberation [8]. Offline on-device systems foreground limits. These limits are not only technical but epistemic: they shape what counts as an acceptable answer, how verification occurs, and when human judgment becomes central. Offline systems make tradeoffs legible—coverage may be narrower and responses slower, but the system's boundedness encourages learners to treat AI output as a proposal to be tested rather than a conclusion to be accepted.

C. Learning Dimensions Activated by Constraint

Across offline deployments, we observe four recurring learning dimensions (see Fig. 1):

- **Epistemic humility:** Students recognize AI as fallible, prompting verification rather than deference.
- **Procedural reasoning:** Learners shift toward stepwise debugging and decomposition when outputs are imperfect.
- **Authorship and responsibility:** Higher friction and reduced copy-paste affordances increase reflective ownership of code and explanations.
- **Collaborative sensemaking:** Students cross-check responses and negotiate next prompts collectively, treating the AI as a shared artifact rather than a private assistant.

IV. Implications for AI Literacy and Learning Design

A. AI Literacy Beyond Prompting

AI literacy is often reduced to prompt skill [7]. Our framing foregrounds additional competencies: when to distrust outputs, how to verify claims, and how to responsibly integrate AI suggestions into authentic work. Offline AI surfaces these competencies precisely because errors and uncertainty are more visible. Learners must decide whether an output is plausible, how to test it, and when to seek human clarification. These are not peripheral skills; they are constitutive of what it means to use AI responsibly in academic and civic life.

B. Design Principles: Making Limits Pedagogically Meaningful

We propose three design principles for offline and low-resource AI integration:

- **Design for visible limits:** Make uncertainty, context boundaries, and confidence cues legible to learners. Avoid interfaces that imply omniscience.
- **Treat breakdowns as learning events:** Scaffold verification, testing, and reflection when AI fails, diverges, or produces ambiguous output.

- **Align AI behavior with curricular intent:** Prefer hinting, decomposition, and conceptual checks over direct solutions, especially in introductory computing contexts.

C. Equity as Epistemic Participation

Equity is not only access to advanced tools—it is participation in knowledge-building practices [3], [8]. Offline AI, designed intentionally, can support epistemic participation by foregrounding verification, authorship, and collaboration—capabilities that remain valuable even when infrastructure is scarce. This avoids a deficit narrative: the question becomes not what offline contexts *lack*, but what kinds of learning practices different infrastructures make more or less likely. Bounded systems, in this light, are not inferior versions of cloud intelligence; they are different pedagogical environments with distinct affordances.

V. Conclusion

Offline, on-device AI is often framed as a constrained substitute for cloud intelligence. We argue the opposite: constraint can be a pedagogical feature. By treating infrastructure as curriculum, scarcity becomes a signal rather than a deficit. Bounded AI can make verification, procedural reasoning, collaborative sensemaking, and responsible authorship more visible in classroom activity—forms of learning that always-available systems can inadvertently suppress.

This paper offers a framework for interpreting a recurring pattern: when AI has limits that learners can see, those limits shape how learners reason, collaborate, and assign responsibility. Future work can test this framework by examining how specific constraints interact with classroom norms and curricular goals, and by identifying design choices that reliably turn infrastructural limits into learnable moments—connecting to the broader grand challenges facing AI and education [2]. We share this work openly and invite educators, researchers, and developers—regardless of institutional affiliation—to adapt, critique, and extend it. The full paper and related resources are available at https://societyandai.org.

References

[1] S. Gattupalli, I. Arroyo, B. Woolf, and E. McEneaney, "SHIELD-ing Education: On-Device AI for Equitable, Offline Computing Education," in *Proc. 25th IEEE International Conference on Advanced Learning Technologies (ICALT)*, 2025, pp. 217–218.

[2] B. Woolf, D. Allessio, I. Arroyo, S. Gattupalli, and B. Zhang, "Grand challenges in AI and education beyond 2030," *Interaction Design & Architecture(s) Journal*, 2025. doi: 10.55612/s-5002-064-001sp.

[3] UNESCO, "Startling digital divides in distance learning emerge," 2020. [Online]. Available: https://www.unesco.org/en/articles/startling-digital-divides-distance-learning-emerge

[4] Meta, "Llama," 2024. [Online]. Available: https://www.llama.com/

[5] LM Studio, "LM Studio," 2024. [Online]. Available: https://github.com/lmstudio-ai

[6] G. C. Bowker and S. L. Star, *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: MIT Press, 1999.

[7] D. Long and B. Magerko, "What is AI literacy? Competencies and design considerations," in *Proc. CHI Conference on Human Factors in Computing Systems*, Honolulu, HI, USA, 2020, pp. 1–16. doi: 10.1145/3313831.3376727.

[8] N. Selwyn, *Should Robots Replace Teachers? AI and the Future of Education*. Cambridge: Polity Press, 2019.