CONVERSATIONAL CURRICULUM: FOUR DESIGN CONSTRAINTS FOR LEARNING PATHS THAT EMERGE FROM DIALOGUE

***Hrozian S.**, Staff Product Designer, Speak, San Francisco, California, USA*

Abstract. *This thesis introduces the Conversational Curriculum: Four Constraints (CC4C) that keep agentic-tutor learning paths coherent across sessions. Drawing on Speak's voice-first lessons and tutoring dialogues, the framework names objective anchor, scope contract, return path, and mastery checkpoint as the constraints that distinguish a learning conversation from an aimless chat.*

Keywords: *conversational AI, agentic tutoring, curriculum design, learning UX, scaffolding, intelligent systems.*

A premade curriculum is set in advance, making it static. A learner moves through prepared units. Each unit is evaluated against fixed objectives. Agentic tutors break the static curriculum. The path is generated turn by turn, shaped by what the learner says, the mistake made, and how the model responds [1, 2]. The result, in current production tutors, is a tradeoff teams are losing on both ends. Strict curricula are coherent and disengaging. Pure dialogue is engaging and unevaluable.

This thesis introduces the Conversational Curriculum: Four Constraints (CC4C). Four design constraints that hold an emergent path together without forcing it back into a fixed sequence. The four are minimal by design. Anything fewer collapses to chat. Anything more collapses back to a static curriculum. The four constraints were derived from production observation of voice-first tutoring at Speak.

Table 1 – Conversational Curriculum: Four Constraints (CC4C)

Constraint	**Definition**	**Failure Mode Prevented**
Objective Anchor	Each session opens by surfacing the learning objective in the dialogue itself.	Aimless chat without measurable goal.
Scope Contract	Agent and learner agree what is in and out of scope.	Silent scope creep into unevaluable territory.
Return Path	Every digression carries a structured way back.	Lost-in-tangent sessions where the original objective is abandoned.
Mastery Checkpoint	Progress is anchored to evaluable mastery, not session length.	Time-spent illusion of learning.

Objective Anchor opens the session by surfacing the learning objective in the dialogue itself. Not in a header. Not in a hidden system prompt. In the first turn the learner can see and correct. "Today we work on past simple. Sound right?" That anchor is what the rest of the session is evaluated against. The cost is that the learner cannot tell at session end what they were supposed to have learned.

Scope Contract is the explicit agreement about what is in and out for the session. A Speak lesson on past simple may permit a brief detour into past continuous (in scope: contrastive grammar) but not into IELTS test strategy (out of scope: different objective). The contract is not a wall. It is a vocabulary for negotiation. When the learner asks an out-of-scope question, the agent surfaces the trade and lets the learner choose. Silently switching contexts is the failure mode.

Return Path is the design rule that every digression carries a way back. The most useful learning moment in a conversation is often the off-topic question. "Why does English use 'do' in questions?" That can open a deeper grammatical insight. Without a return path, the digression becomes the new direction. The objective is lost. With a return path, the digression becomes a recorded enrichment, and the agent surfaces the bridge: "Good question. Short answer first, then back to past simple." This extends the scaffolding concept [1] – coined for human tutors – to emergent dialogue.

Mastery Checkpoint anchors progress to evaluable mastery. Not session length, words spoken. This is the hardest constraint to implement and the most consequential. A 30-minute session that produced no demonstrable mastery is not learning, however engaging the conversation felt. The agentic tutor must apply it to generated content. The agent has to construct a mastery check on the fly that matches the path the learner walked through.

The four constraints together describe a conversation that knows where it is going. None of them require the agent to follow a script. All of them require the agent to hold structure most production prompts do not enforce.

The implementation gap is real. Most LLM based tutors today have none of these constraints. Or one or two without the others. A tutor with an objective anchor and no return path becomes rigid. A tutor with a scope contract and no mastery checkpoint becomes plausible sounding but unevaluable. The four reinforce each other. Partial adoption produces partial coherence.

CC4C sits inside a longer tradition of intelligent tutoring research. Cognitive tutors [2] showed that the structure of practice matters more than the medium's surface fluency. VanLehn [3] found that intelligent tutors approach human-tutor effectiveness when feedback is granular and recovery is well designed. Agentic tutors inherit those findings. CC4C is an attempt to operationalize them under conditions where the path itself is generated.

The framework has limits. New learners need harder objective anchors and advanced learners productively burst the scope contract. Mastery checkpoints constructed on the fly are the hardest engineering problem and may degrade with model variance.

The next generation of agentic tutors will be judged less on how fluent the dialogue is and more on whether the path it traced makes sense looking back. That is the real job of a mastery check. The learner has to be able to look at the session afterward and recognize their own progress, even when the path was emergent. If they cannot, the lesson did not really happen, no matter how good the conversation sounded.

References

1. Bruner, J. S. (1966). Toward a theory of instruction. Harvard University Press.

2. Anthropic. (2024). Model Context Protocol specification. https://modelcontextprotocol.io.

3. Wood, D., Bruner, J.S., & Ross, G. (1976). The role of tutoring in problem solving. Journal of Child Psychology and Psychiatry, 17(2), 89-100.

4. Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. The Journal of the Learning Sciences, 4(2), 167-207.

5. VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. Educational Psychologist, 46(4), 197-221.