

Holding the Room: Designing Interaction without Technology in Large Chinese University Lectures

Boran Liu¹, Lingzhi Tan², Haiying Dong^{3*}

1. School of Marxism, Taizhou University, Taizhou, Zhejiang, 318000, China

2. Jiaojiang District Education Bureau, Taizhou, Zhejiang, 318000, China

3. The Affiliated School of Taizhou University, Taizhou, Zhejiang, 318000, China

*Corresponding author: Haiying Dong, 87978348@qq.com

Copyright: 2026 Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY-NC 4.0), permitting distribution and reproduction in any medium, provided the original author and source are credited, and explicitly prohibiting its use for commercial purposes.

Abstract: In Chinese universities, undergraduate lectures of eighty to one hundred and thirty students have become the norm. Faced with such rooms, teachers often feel that meaningful interaction is out of reach without technological support: smart response systems, live bullet-screen chat, or AI-mediated question platforms. This paper argues that the bottleneck in large lectures is design, not technology. It proposes that “interaction” in a large room should be redefined away from public speech by a few visible students and toward a ninety-second cognitive operation in which every student is engaged at the same time. Drawing on recent empirical work on active learning, retrieval practice, and large-class teaching, the paper identifies three structural constraints of the Chinese large lecture—the voice problem, the accountability problem, and the timing problem—and translates them into three reusable, no-technology interaction patterns: neighbour recount, folded-paper vote, and three-sentence exit. Holding the attention of a hundred or so students is a design task, and a smart classroom platform can support—but hardly replace—the teacher’s weekly practice of inserting three ninety-second cognitive contracts into a forty-minute lecture.

Keywords: Large Lecture; Classroom Interaction; Teaching Design; Undergraduate Teaching; Low-Tech Pedagogy; Chinese Higher Education

Published: Oct 22, 2026

DOI: <https://doi.org/10.62177/jetp.v3i4.1756>

Introduction

Step into a large lecture hall at a typical Chinese university and the scene is easy to recognise. A hundred or more students fill the tiered seats. The teacher stands at the front, moves through the slides at a steady pace, and from time to time asks, “Is that clear?” A few heads nod; nobody says anything. Ten minutes in, attention in the back rows starts to drift and phones come out. By the thirty-fifth minute the class has settled back into what it was at the beginning: one voice at the front, a hundred silent listeners.

The usual explanation for this state of affairs is that the class is simply too large to allow interaction. The usual remedy, increasingly funded by universities, is technological: clickers, live-polling apps, bullet-screen commenting, learning-analytics dashboards that flag inattentive students. Both the explanation and the remedy, this paper argues, tend to overlook the real bottleneck. A large Chinese lecture rarely lacks devices. What it often lacks is a workable answer to the question of what interaction should mean when a hundred people sit in one room.

The paper therefore asks a deliberately constrained question: what would it mean to design interaction in a large Chinese university lecture if no additional technology were available—no app allowed, only a podium, a slide deck, and a stack of paper? The answer begins with a redefinition. Interaction is reframed from “public speech by visible students” to “simultaneous cognitive engagement of every student in the room.” Once the term is reframed in this way, the large lecture stops looking like a broadcast medium. It becomes a structure in which the teacher’s job is to design, roughly every ten minutes, a short moment in which every student has something concrete to do.

Three claims follow. The large Chinese lecture imposes three structural constraints that technology alone does not remove—the voice problem, the accountability problem, and the timing problem. Each constraint can be met by a low-cost interaction pattern that needs no technology and assumes no acquaintance between neighbours. And the teacher’s role in such a room is better understood as that of a conductor than a performer: the aim is not to hold the floor but to hold the room.

The argument proceeds as follows. Section 1 examines the “large lecture equals lecture-only” assumption and the technological remedy built on it. Section 2 sets out the three constraints as an analytical framework, with an illustrative coding schema. Section 3 turns the framework into three reusable interaction patterns. Section 4 discusses implications for Chinese university teaching; Section 5 concludes.

1. Background

1.1 The “Large Lecture Equals Lecture-Only” Assumption

Massification has changed the material conditions of Chinese undergraduate teaching. General-education courses, large introductory classes, and second-major programmes routinely meet in rooms of eighty to one hundred and thirty seats. Staffing has not grown in proportion, so a single teacher is expected to transmit an entire discipline to a room in which anonymity is structural. A tacit division of labour follows: the teacher lectures, and interaction is treated as a luxury reserved for small seminars.

This division produces a predictable script. The teacher asks rhetorical questions, accepts silence as the expected answer, fills the gap personally, and moves on. Students learn that their role is to record rather than to think. Nor is the pattern uniquely Chinese. The largest observational study of undergraduate STEM teaching to date, covering more than two thousand classes in North America, found lecturing prominent throughout the curriculum—even in classrooms built to support active learning ^[1]. This suggests that scale and habit, rather than national culture, account for much of the pattern.

A common response is to reach for technology. Universities invest in clicker systems, classroom-response apps, and AI dashboards that flag inattentive students. These tools do raise the number of visible responses; what they do not necessarily raise is cognitive engagement. Students learn to tap a response button without thinking, just as they once learned to nod without listening. A silent lecture becomes a noisy one, but a passive lecture does not thereby become an active one.

1.2 What “Interaction” Actually Means in a Room of a Hundred

The confusion rests on an unexamined definition. In a large room, two very different activities go by the name of interaction. The first is public-speech interaction: a student stands, raises a hand, and addresses the room. In a lecture of a hundred, this activity is structurally expensive. It demands visible bravery, rewards the already confident, and admits at most three or four participants in a forty-minute period. Teachers who measure success by this kind of interaction can easily come to judge large lectures a failure.

The second is cognitive-engagement interaction: every student in the room completes, within ninety seconds, a bounded operation—explaining a concept to a neighbour, writing a single line of reasoning, voting on a claim. This activity is structurally cheap. It scales with room size, demands no bravery, and can be repeated three or four times in one period without a single visible hand. Its cognitive rationale is well established. In controlled experiments on vocabulary learning, once students could recall an item, further re-study had no measurable effect on their retention a week later, whereas continued retrieval produced a large gain ^[2]. In that setting, brief obligatory acts of retrieval, rather than extra exposure, drove the learning gains.

None of this is to dismiss public speech. The claim is narrower: in a large Chinese lecture, public speech may be better treated as the exception, with ninety-second cognitive engagement as the design goal—repeated often enough that few students

spend the forty minutes entirely in mental absence.

1.3 Why Technology Is Not the First Move

The objection to technology here is one of sequencing, not principle. A large room that has not learned to design ninety-second cognitive contracts will use a polling app to ask rhetorical questions and will produce, predictably, rhetorical votes. Technology amplifies the pedagogy that is already in place. If the pedagogy is broadcast, the app makes broadcast louder.

Recent evidence sharpens the point. A meta-analysis of 225 studies of undergraduate STEM courses found that active learning raises examination performance and lowers failure rates relative to traditional lecturing^[3]. Yet when students were randomly assigned to actively taught or traditionally taught sections of the same physics course, those in the active sections felt they had learned less while objectively learning more^[4]. The implication is uncomfortable. A teacher who waits for students to report that the new design feels better may wait indefinitely. The workable sequence is the reverse of a technology-first one: design the cognitive contract first, then ask whether any tool deserves to be purchased.

2. An Analytical Framework: Three Structural Constraints of the Large Lecture

Large lectures are hard for three observable, structural reasons, and installing more devices changes none of them.

2.1 The Voice Problem: One Mouth, a Hundred Ears

A teacher has one mouth; the room has a hundred pairs of ears. Direct two-way conversation between the teacher and every student is physically impossible, and this is geometry, not a technical defect. The standard response—"Who would like to share?"—asks for volunteers, and volunteers are by construction a small subset of the room. The structural inference is straightforward. In a large room, the teacher cannot be the conversation partner of every student, so the teacher must design moments in which students speak to each other.

2.2 The Accountability Problem: No One Is Watching You

In a room of a hundred, any individual student is effectively invisible. Taking notes, scrolling a phone, and thinking about dinner are indistinguishable from the front. A voluntary-contribution system therefore rests on students' motivation to perform in public—unevenly distributed anywhere, and in some Chinese classrooms further dampened by face-consciousness. Synthesis work on feedback in higher education is relevant here: impact lies in the nature of the information students receive about where they stand in their learning, not in the amount of evaluative comment delivered^[5]. Such information cannot be handed to each student individually in a large room, but it can be engineered into short structured tasks.

2.3 The Timing Problem: Forty Minutes Is Too Short for a Hundred Conversations

A forty-minute lecture holds about twenty-four hundred seconds. Give groups five minutes to discuss and then invite three groups to report for two minutes each, and eleven minutes are gone; three groups have been heard, and the other ninety-odd students have contributed nothing the room could register. Small-class techniques—group work followed by whole-class sharing—quickly overrun the time budget under this arithmetic. They were designed for a denominator of twenty.

Table 1 makes the framework operational. It lists the three constraints together with two pressures that intensify them—students' reluctance to speak publicly, and the decay of attention across a forty-minute period—and pairs each with a typical teacher response and a redesigned low-cost move. The middle columns summarise typical behaviour reported in large-scale observational studies of university teaching rather than measurements of one specific classroom^[1]; the schema is offered as an analytical baseline.

Table 1. An Illustrative Coding Schema for Large-Lecture Interaction: Structural Constraints, Typical Responses, and Redesigned Moves

Structural constraint	Typical teacher response	Why it fails in a large room	Redesigned low-cost move
Voice problem (one mouth, a hundred ears)	Whole-class questioning	Only 3–5 visible students speak; the rest observe	Structured peer talk: every student explains to a neighbour within 90 seconds
Accountability problem (no one watches)	Voluntary volunteering	Quiet students stay quiet; confident students dominate	Safe random invitation: one student reports a neighbour's answer
Timing problem (40 minutes, ~100 students)	"Share your group's conclusion"	A few groups report; most groups never speak	Written artefact: every student writes one line; the teacher samples two

Structural constraint	Typical teacher response	Why it fails in a large room	Redesigned low-cost move
Public-speaking pressure	“Who wants to share?”	Students perform rather than think	Anonymous channel: folded notes passed to the front
Attention decay over 40 minutes	Uninterrupted lecturing	Attention drifts well before the period ends	A 90-second cognitive task inserted about every 10 minutes

3. Designing Interaction: Low-Cost Moves for Large Chinese Lectures

Together, the rows of Table 1 point to three reusable patterns. None requires an app, acquaintance between neighbours, or extra class time.

3.1 The 90-Second Contract

The basic design unit is the ninety-second contract. About every ten minutes, the teacher interrupts the lecture and asks the room to perform one bounded cognitive operation within ninety seconds. A forty-minute lecture thus carries three such contracts—near minutes ten, twenty, and thirty. The teacher does not have to monitor every pair; the structure does the supervising.

The goal, it should be stressed, is not to maximise active-learning minutes. In a controlled study of a large-enrollment introductory biology course in the United States, three instructors taught with identical materials and exams to classes of 340 to 380 students each; they gave students significantly different amounts of time for thinking and peer discussion, yet these differences did not predict performance on the associated exam questions ^[6]. What matters is the design of the episode, not its duration—whether the ninety seconds genuinely oblige every student to retrieve, explain, or decide. The three patterns below are deliberately short and bounded, because stretching them would neither improve learning nor reward the teacher. In the authors’ own large-class teaching, they have so far proved workable.

3.2 Three Reusable Interaction Patterns

The first pattern is the neighbour recount. The teacher asks student A to explain, in about half a minute, a concept just introduced, while partner B listens and then says back what A said. The teacher then points at a random row and asks one B to report what their A said. The decisive move is that students report a neighbour’s view, not their own. Personal exposure drops; both partners must listen; and the teacher hears several voices without leaving the podium.

The second pattern is the folded-paper vote. The teacher poses a single evaluable claim—“The author’s argument here is wrong,” or “This formula applies to this case.” Every student writes “agree,” “disagree,” or “not sure” on a slip, folds it, and passes it forward along the row. The teacher draws three slips at random, reads them aloud, and asks the room to defend each position. Because the channel is anonymous, no face is at risk; because nobody knows whose slip was drawn, the votes are more likely to reflect genuine judgement. No polling app is required.

The third pattern is the three-sentence exit. In the last three minutes of the lecture, each student writes three sentences on a slip: one point still unclear, one claim that surprised them, one question worth asking. Slips are collected at the door. The teacher picks two or three and opens the next lecture with them. One small device does three jobs at once: it closes the lecture cognitively, gives every student a voice, and tells the teacher what the room actually understood.

3.3 The Teacher as Conductor, Not Performer

What these moves imply is a change in professional self-image. In a small seminar the teacher is a conversation partner. In a large lecture the teacher is closer to a conductor: the work is not to sing every note but to write a score in which a hundred people play at once. Conducting needs no technology. It needs the discipline to stop talking every ten minutes, to design the coming ninety seconds, and to trust that a room will engage when the invitation does not require anyone to expose themselves.

4. Discussion

The framework carries three implications for Chinese university teaching.

It reframes the bottleneck as one of design rather than equipment. When a large lecture fails to engage, the more useful first question is not “what platform should we buy?” but “what ninety-second contracts did the teacher design this week?” For

institutions, this suggests that support for teachers' instructional design—workshops, peer observation, preparation time—deserves at least the same priority as investment in smart-classroom hardware.

The three patterns are tailored to the specific conditions of the Chinese large lecture. None of them requires students to know their neighbours, because the task itself defines the interaction. None requires the teacher to patrol the aisles, because random sampling does the checking. Paper and fixed rows replace any shared app, so all three can be piloted in an ordinary semester without capital expenditure.

The argument also complements, rather than competes with, the broader active-learning literature. Meta-analysis says active learning improves outcomes on average ^[3]; experiments warn that students may feel worse in active settings even while learning more ^[4]; and recent controlled evidence suggests that the sheer time spent on active learning is not the decisive variable ^[6]. What the present paper adds is concrete: what active learning can look like inside a forty-minute Chinese lecture of a hundred students. Three ninety-second contracts, designed by the teacher at the front of the room—no clickers, no dashboard, no bullet-screen feed.

The limits of the paper should be stated plainly. It is a conceptual, teaching-focused paper rather than an empirical study. Table 1 offers illustrative baselines synthesised from the higher-education literature, not measurements of a specific Chinese classroom. The three patterns need to be tested across disciplines, class sizes, and student populations before stronger claims can be made. Future work should pair the framework with structured classroom observation and student-reported measures of cognitive engagement.

5. Conclusion and Recommendations

This paper has argued that the large Chinese university lecture is not structurally doomed to passivity. When such lectures fall flat, the cause is more often design than hardware, and the remedy lies in redefining interaction as ninety-second cognitive engagement rather than public speech by visible students. Holding a room of a hundred is a professional practice that any teacher can begin with the very next forty-minute period—without a clicker, without an app, and without waiting for a smart classroom.

Three recommendations follow from the argument. Teachers should prepare, for every forty-minute lecture, three ninety-second cognitive contracts alongside the slide deck. Departments should treat neighbour recount, folded-paper vote, and three-sentence exit as teachable methods and disseminate them through regular teaching workshops, instead of reaching first for more expensive technologies. Universities evaluating large lectures should count something other than raised hands: the proportion of students who, within a given forty minutes, were cognitively required to do something.

The back row is not a different room. It is the same room, and it is listening. Whether it is also thinking is a question dashboards are poorly placed to answer. The teacher answers it every ten minutes, in the decision to stop talking for ninety seconds.

Funding

No

Conflict of Interests

The authors declare that there is no conflict of interest regarding the publication of this paper.

Reference

- [1] Stains, M., Harshman, J., Barker, M. K., et al. (2018). Anatomy of STEM teaching in North American universities. *Science*, 359(6383), 1468–1470. <https://doi.org/10.1126/science.aap8892>
- [2] Karpicke, J. D., & Roediger, H. L. (2008). The critical importance of retrieval for learning. *Science*, 319(5865), 966–968. <https://doi.org/10.1126/science.1152408>
- [3] Freeman, S., Eddy, S. L., McDonough, M., et al. (2014). Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the National Academy of Sciences*, 111(23), 8410–8415. <https://doi.org/10.1073/pnas.1319030111>
- [4] Deslauriers, L., McCarty, L. S., Miller, K., et al. (2019). Measuring actual learning versus feeling of learning in response

to being actively engaged in the classroom. *Proceedings of the National Academy of Sciences*, 116(39), 19251–19257. <https://doi.org/10.1126/science.1821936116>

- [5] Hattie, J. (2015). The applicability of Visible Learning to higher education. *Scholarship of Teaching and Learning in Psychology*, 1(1), 79–91. <https://doi.org/10.1037/stl0000021>
- [6] Cen, X., Lee, R. J., Contreras, C., et al. (2024). Time spent on active learning activities does not necessarily correlate with student exam performance: A controlled case study. *Journal of Microbiology and Biology Education*, 25(3), e00073–24. <https://doi.org/10.1128/jmbe.00073-24>