Article

From Canon to Classroom: A Framework for Integrating China's Excellent Traditional Culture into Digital Ideological and Political Theory Courses in the Age of Artificial Intelligence

Yiwen Zhang[1,*]

[1] Geely University of China, Chengdu, China
* Correspondence: Yiwen Zhang, Geely University of China, Chengdu, China

Abstract: Ideological and Political Theory Courses (IPTC) are the curricular backbone of moral and political formation in Chinese higher education, and China's Excellent Traditional Culture (CETC) has long been treated as one of their richest intellectual reservoirs. The arrival of artificial intelligence (AI)—large language models, multimodal generation and immersive media—promises to industrialise the digitisation of that reservoir. This article argues that the dominant techno-centric reading of this opportunity is mistaken. Drawing on the curriculum-theory distinction between the intended, enacted and experienced curriculum, and reinterpreting the author's earlier Chinese-language study within a critical, design-oriented frame, the article proposes a three-layer integration framework: a semantic layer (structured cultural knowledge with provenance and encoded contestation), an experiential layer (embodied and narrative encounters) and a reflexive layer (training in reading machine-mediated tradition). It specifies four design principles—interpretive transparency, cultural fidelity, cognitive apprenticeship and contested-space pedagogy—a four-dimensional evaluation model, and a staged readiness ladder that allows under-resourced departments to implement the framework's core without new infrastructure. It then examines risks specific to probabilistic mediation of a normative tradition: semantic and temporal flattening, appropriation of interpretive authority, the atrophy of interpretive stamina, the commodification of heritage affect, and the erosion of deliberative space. The conclusion is that the central question is not whether AI can deliver tradition to students, but whether students can still be taught to argue with tradition—and with the machines that now speak on its behalf.

Keywords: Excellent Traditional Culture; Ideological and Political Theory Courses; generative artificial intelligence; curriculum design; cultural semantics; digital humanities pedagogy

Received: 22 August 2026
Revised: 13 October 2026
Accepted: 28 October 2026
Published: 31 October 2026

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

1.1. A Familiar Aspiration, an Unfamiliar Technical Environment

Few educational projects in contemporary China carry as much institutional weight as the integration of China's Excellent Traditional Culture (CETC) into the Ideological and Political Theory Courses (IPTC) that every undergraduate is required to take. The policy lineage runs from the 2004 Central Document No. 16 on strengthening ideological and political education, through the guidelines of 2017 and 2019 on inheriting fine traditional culture, to the current emphasis on "cultivating virtue through education" (li de shu ren). The ambition is equally clear. CETC is not meant to be an ornament attached to political theory; it is meant to supply the moral vocabulary—ren (benevolence), yi (rightness), li (ritual propriety), zhi (wisdom), xin (trustworthiness), he (harmony)—in which contemporary citizenship can be articulated and felt.

What is new is not the aspiration but the technical environment. Early "digitisation" meant putting texts, images and video online. The current generation of tools—large

language models, diffusion-based generation, speech synthesis, learning analytics—does something categorically different: it generates cultural content, speaks in the voice of tradition, and adapts itself to the individual learner without a human authorising each utterance. The international literature on generative AI in higher education has moved rapidly from enthusiasm to a sober assessment of epistemic and ethical risk [1].

Chinese-language scholarship on CETC in IPTC was largely written for the earlier moment. It tends to describe digital integration as a matter of resource abundance, platform construction and teacher upskilling—important concerns, framed around the assumption that the digital artefact is authored by a human and consumed by a student. The author's earlier study exemplifies the genre: it identifies the necessity of integration, diagnoses three challenges (fragmented resources, uneven teacher capacity, imprecise cultural translation) and proposes three pathways (immersive multi-scenario presentation, intelligent personalised recommendation, multi-interactive smart classrooms). It was written within a single national conversation and organised its argument around delivery.

This article keeps the substantive commitment and rebuilds the argument for a different audience and a different technological moment. Three things change. The unit of analysis shifts from delivery to curriculum: instead of asking how to push resources more efficiently into classrooms, it asks what happens to a tradition when it passes through a generative system. The vocabulary shifts from resource integration to the curriculum-studies triad of intended, enacted and experienced curriculum, which exposes problems the delivery frame conceals. And the critical stance sharpens: where the earlier study treated algorithmic limitations as implementation difficulties to be solved by better technology, this article treats them as constitutive features of the medium. A probabilistic model trained on an indiscriminate corpus does not merely "fail to understand" the Analects; it redistributes interpretive authority away from teachers and traditions toward whatever the corpus has made statistically salient. That is a curricular problem, not a bug.

1.2. Why CETC Is a Hard Case for AI

CETC is an unusually demanding object for machine mediation—harder than a corpus of scientific facts or procedural law.

Commentarial rather than propositional. The canon is not a set of doctrines but a set of texts read, contested and rewritten for two and a half millennia. The Analects is not self-interpreting; its meaning is constituted by the layered tradition of commentary—Mencius, Xunzi, Zhu Xi, Wang Yangming. A model trained on a flat corpus reproduces the most frequent reading, usually the most recent textbook reading, usually a simplified moralism. The commentarial structure is precisely what a frequency-based model is worst equipped to represent.

Embodied and performative. Much of CETC lives in practices—ritual, calligraphy, music, festival, kinship etiquette—which are know-how traditions transmitted by apprenticeship. Text-based AI can describe them but cannot transmit them; multimodal AI can simulate them, but simulation is not apprenticeship.

Normatively load-bearing. Errors are not merely factual but formative. A model that hallucinates a plausible Confucian maxim has inserted a fabricated norm into a student's moral vocabulary, with the authority of tradition behind it.

Internally plural. "Excellent traditional culture" is a policy category, not a natural kind; it gathers Confucian, Daoist and Buddhist strands, literati and vernacular practices, majority and minority heritage. Any curriculum addressing CETC as a single object will privilege one strand, usually the Confucian literati one. This article focuses on the classical-Confucian textual tradition, as the strand most present in IPTC syllabi and most distorted in machine mediation.

1.3. Scope, Contribution and Structure

Three disclaimers are in order, because the argument is easily misread. The article does not argue that AI is unsuitable for cultural education—the tools are genuinely powerful for access, translation and multimodal representation, and the argument is about design, not adoption. It does not treat traditional culture and critical thinking as

opposed: the tradition is itself a tradition of argument, and teaching students to contest positions within it extends that dynamic rather than imposing something foreign; this is worth stating because Anglophone commentary often assumes Confucian heritage to be inherently authoritarian. And it is not a proposal for a specific platform: platforms will change, the design questions will not.

The contributions are four. Conceptually, the article reframes digital integration from resource delivery to curricular mediation, and supplies the vocabulary to discuss the latter. Critically, it identifies risks specific to generative mediation of a normative tradition that are invisible in generic discussions of AI in education. Practically, it specifies achievable designs, including a ladder for departments without new infrastructure. Programmatically, it sets out an empirical agenda, since the field's weakness is not a shortage of proposals but of evidence about what students learn. Section 2 situates the argument in the relevant literatures; Section 3 sets out the three-layer framework; Section 4 specifies four principles and their tensions; Section 5 addresses assessment; Section 6 examines risks; Section 7 discusses institutional conditions and research questions; Section 8 concludes.

2. Literature and Conceptual Grounding

2.1. Two Translation Choices

"Ideological and Political Theory Courses" renders sixiang zhengzhi lilun ke, but "ideological" carries connotations the Chinese term does not: sixiang refers to thought and formation of mind, not to ideology in the pejorative sense. "Excellent Traditional Culture" renders youxiu chuantong wenhua, and "excellent" is awkward in English, as are "fine" and "outstanding". The awkwardness is diagnostic: these are not neutral labels but terms embedded in a policy discourse, and smoothing them into "civic education" or "values education" loses the specificity that makes the case interesting. The official renderings are retained for referential accuracy, and used critically.

2.2. Curriculum as Intended, Enacted and Experienced

The most useful analytical resource predates the digital turn by decades. Curriculum scholars distinguish three versions of any curriculum: the intended curriculum (what syllabi and policy documents say should happen), the enacted curriculum (what teachers actually do) and the experienced curriculum (what students actually take away). The gaps between them are the ordinary substance of educational research. Any reform is judged by which version it actually changes, and reforms frequently change the first while leaving the third untouched.

Generative AI transforms these gaps in a specific way. It narrows the gap between intended and enacted: a well-designed platform can deliver a consistent, high-fidelity version of the intended curriculum to thousands of classrooms. But it widens the gap between enacted and experienced, because machine mediation inserts a layer of interpretation between the teacher's intention and the student's encounter that neither party controls. A student conversing with a chatbot about ren is not experiencing the teacher's enacted curriculum; they are experiencing a statistical composite of everything the model has read about ren, including Western secondary literature, popular summaries and mistranslations.

This asymmetry—more standardised delivery, less controllable experience—is the central curricular fact of the AI era, and it is invisible to a delivery-oriented framework. It reframes the design task: the goal is not to maximise the reach of the intended curriculum but to keep the experienced curriculum honest, by making the machine's interpretive layer visible, interrogable and accountable.

2.3. The Policy Context

Three features of the setting shape what any framework can achieve. IPTC is a required, credit-bearing, nationally specified component of every undergraduate programme; it is not an elective in moral or religious studies, and comparison with

Western general-education courses in ethics is misleading. It is centrally framed but locally enacted: guidelines specify aims and content domains, while institutions and teachers design syllabi and assessment—which is why the gap between intended and enacted curriculum is such a live issue here, and why teacher capacity is decisive. And integration of traditional culture into IPTC is an explicit policy objective articulated across documents over two decades, so the pedagogical question is never whether to include CETC, but how to include it seriously rather than ceremonially.

2.4. Values Education and the Chinese Debate

Zhu's account of moral education in Chinese curriculum reform describes a shift from transmission of prescribed virtues toward more experiential and child-centred approaches, while noting the persistent centrality of socialist core values and traditional culture in the official curriculum [2]. More recent work examines the tension between the revival of Confucian classical education and the demands of a modern, exam-oriented system, showing how parents and teachers mobilise traditional pedagogy for contemporary anxieties about moral formation [3].

This literature matters for two reasons. It prevents the crude reading that IPTC is simply indoctrination: the Chinese debate about how to teach values is a genuine pedagogical debate, whose positions on experiential learning and the relation between knowledge and conduct resonate with international work in character education. And it foregrounds a problem this article takes seriously—the tension between transmitting a tradition and teaching students to think within it. A tradition merely transmitted becomes inert; one only critiqued becomes unavailable as a resource for living. AI mediation pulls hard toward transmission, because generative systems are optimised for fluent reproduction rather than productive difficulty.

2.5. Heritage Pedagogy

Work on digital storytelling in cultural heritage management shows that the narrative structure of interpretation matters more than the fidelity of the medium: audiences engage when heritage is framed as a story with stakes, agents and unresolved tensions, not as an inventory of artefacts [4]. Research on immersive virtual reality for intangible heritage similarly finds that presence and emotional engagement, rather than visual realism, drive outcomes, and that poorly designed immersion can distract from rather than deepen historical understanding [5].

Two implications follow. The value of digital mediation lies in structuring interpretive encounters, not in storing or displaying materials. And the design question is always comparative: which encounter does this medium afford that the classroom could not otherwise provide, and what does it cost? A 360-degree tour of a Song academy may afford a felt sense of spatial order that a photograph cannot; it may also substitute spectacle for the slow labour of reading that the tradition itself regards as formative.

2.6. AI in Higher Education

Research on AI in higher education is the fastest-moving of these literatures. Early commentary emphasised opportunity—personalised tutoring, reduced administrative burden, scaffolded writing—alongside concerns about integrity. Multidisciplinary assessments have since shifted toward governance and equity, distinguishing sharply between using AI to produce content and teaching students to understand AI as a cultural and epistemic technology [1]. A related distinction separates tool proficiency from critical AI literacy: understanding what a model is and is not, and being able to interrogate the data and design choices behind it. For a course whose purpose is to form critical, value-committed citizens, that literacy is not an add-on but part of the content.

Teacher capacity is the persistent bottleneck. Professional development concentrates on tool operation rather than pedagogical reasoning, and teachers are rarely given frameworks for deciding when not to use a technology. IPTC teachers face a double demand: mastering the technology while remaining authoritative custodians of a tradition it represents imperfectly.

2.7. A Fourth Resource: The Chinese Theoretical Tradition Itself

Chinese educational thought contains a differentiated theory of exactly the question at issue. The Xueji grounds instruction in the relation between study and practice; Zhu Xi's reading method prescribes slow, repeated engagement with a text; Wang Yangming's unity of knowledge and action explains why propositional learning fails to change conduct. These are not decorative. Wang Yangming states precisely why the experiential and behavioural dimensions of assessment cannot be reduced to the cognitive one; Zhu Xi supplies a rationale for why AI-mediated shortcuts are corrosive, since the value of reading lies partly in the labour of reading. Using the tradition as a theoretical resource rather than merely as content makes the framework intelligible from within the tradition it concerns, and guards against a dependency in which the only vocabulary for criticising educational technology is imported from elsewhere.

2.8. The Gap

Across these literatures a gap is visible. Curriculum theory supplies analytic vocabulary but not cultural specificity; the China-focused values-education literature supplies context but rarely engages generative technology; heritage pedagogy supplies design insight but is oriented toward museums rather than required courses; AI research supplies the technical agenda but treats "content" as an interchangeable input. What is missing is an account of what happens when a normatively authoritative, commentarial, embodied tradition is mediated by probabilistic, generative, scalable systems inside a required curriculum with formative aims. The gap reduces to one question: what must a generative system be able to represent before it may be trusted to speak about a tradition in a course that exists to form values? No single technical fix will answer it, because the requirements are of different kinds.

3. A Three-Layer Integration Framework

3.1. Overview

The framework treats integration as the construction of three interdependent layers, each answering a different curricular question.

Table 1. Three-Layer Integration Framework for AI-Assisted Learning in a Tradition: Semantic, Experiential, and Reflexive.

Layer	Curricular question	Core artefacts	Dominant AI affordance	Primary risk
Semantic	What does the tradition say, and who says so?	Structured knowledge base, annotated primary texts, citation provenance	Retrieval, entity linking, summarisation	Flattening of commentarial depth; fabricated citation
Experiential	What does it feel like to live inside this tradition?	Immersive scenes, ritual simulations, student media	Multimodal generation, VR/AR, voice synthesis	Spectacle displacing understanding
Reflexive	How do I know what I know?	Prompt journals, model comparison, argument maps	Transparency tools, counterfactual generation, error injection	Reinforcing algorithmic authority

The layers are ordered by dependency, not sequence. "Layer" is used rather than "module" or "stage" because the three are simultaneous conditions of the same learning event: a student reading a translated passage (semantic), reacting to it (experiential), wondering how the translation was chosen (reflexive). What is proposed is not three new activities but a redistribution of attention across three dimensions always already in play. The semantic layer supplies the ground on which the experiential layer can be trusted; the

reflexive layer supplies the distance without which both become indoctrination by other means.

Two boundary conditions apply. The framework assumes a teacher with discretion over materials and either a class small enough for discussion or one divisible into sections; it does not describe a fully automated massive course, where only the semantic layer is achievable. And it is described in terms of functions rather than products, so that it can be implemented with modest tools: a department with no budget can implement the reflexive layer with two free chatbots and a packet of primary texts.

3.2. The Semantic Layer: From Corpus to Knowledge Structure

The foundational move is to stop treating cultural resources as a corpus and start treating them as a structure. A corpus is a pile of texts; a structure encodes relations—between a concept and its classical locus, between a locus and its commentaries, between a commentary and the controversy it intervened in. Four properties matter.

Provenance. Every retrievable proposition is linked to a primary source, a specific edition and, where relevant, a commentary tradition. Asked what junzi means, the system returns not a definition but a chain: which passages, which translation, which commentator. Provenance is what allows a generative system to be used in a course whose subject matter is the authority of texts.

Contestation. The structure should encode disagreement, not consensus. Xing (human nature) is not a concept but a site of a famous dispute between Mencius and Xunzi; a good base represents that dispute as a first-class object rather than averaging it away. This runs against the grain of how language models are trained, and it is the most important decision for preserving the commentarial character of the tradition.

Granularity. Concepts must be represented at several levels—term, passage, chapter, school, era—so that a student can move between close reading and panorama. Generative systems are good at panorama and bad at close reading; a structured base compensates.

Bilingual anchoring. Ren is not "benevolence"; li is not "ritual". These glosses are useful entry points and dangerous endpoints, and encoding the translation problem in the structure makes it teachable rather than invisible.

The AI affordance is retrieval-augmented generation: the model answers from the curated structure, with citations, and surfaces disagreement where the structure records it. This does not eliminate hallucination but converts an unbounded risk into an auditable one. A micro-example clarifies the difference. Asked "Did Confucius believe human nature is good?", a general model produces a fluent synthesis—"Confucius emphasised moral cultivation and the inherent potential for goodness"—which is not false but collapses a two-thousand-year dispute into a slogan. A structured system returns the relevant Analects passages and what they do and do not say; the note that the explicit doctrine belongs to Mencius rather than Confucius; Xunzi's contrary position and its basis; the Song settlement of the question through the Four Books; and a prompt asking the student to judge whether the textbook formulation is defensible. Both answers contain true statements, but only one is curricularly honest: the first teaches a slogan, the second a discipline.

Two cautions follow. Curated bases are expensive and age quickly; they need the treatment given to a textbook, with editions, review cycles and named editors. And over-curation is a real risk: a base encoding only the editor's preferred reading becomes a more efficient instrument of indoctrination than any textbook. The remedy is procedural—multiple editors, published disagreement, and a record of what was excluded and why.

3.3. The Experiential Layer: Designing for Productive Disorientation

If the semantic layer addresses truth, the experiential layer addresses presence—the felt sense that this tradition is a world one could inhabit rather than a list of doctrines. Four design patterns carry very different pedagogical weight.

Reconstruction. Immersive recreation of a historical scene: a Song academy, a lineage hall, a xiangyue community covenant ceremony. The value is spatial and social—students grasp the relational order within which classical norms made sense. The risk is

antiquarian spectacle, so reconstruction should always be paired with a question: what is absent here, and who decided?

Simulation of practice. Structured role-play of a ritual or a moral dilemma, with AI in supporting roles—a sceptical interlocutor, a historical figure with a documented position, a contemporary student with objections. The AI should generate resistance, not agreement. A simulation in which the Confucian interlocutor always wins teaches nothing; one in which the interlocutor has a coherent but contestable position teaches argument.

Counterfactual interrogation. Ask the system how a practice would have been understood by actors usually absent from textbooks—a peasant household rather than a literatus, a woman in an inner chamber rather than a magistrate, a merchant rather than a scholar-official. The model will fabricate testimony, and the fabrication must be treated as an object of study rather than evidence. Used well, it teaches that any account of a tradition has a social position, and that the canonical account is a selection.

Student-authored narrative. Students produce short digital stories—a family ritual, a local craft, a temple festival—using generative tools for scripting, image and audio, and present them with an account of what the tools changed. Digital storytelling is well supported in the heritage literature as a mode of interpretation rather than presentation [4], and it moves students from consumers of cultural content to producers of cultural argument, which is where internalisation occurs.

In all four, the yield depends on a choice the earlier literature overlooks: the framing task surrounding the experience. Presence without a subsequent analytic task produces a memorable hour and no durable learning. The experiential layer must be bookended by a semantic question before and a reflexive task after.

A module design illustrates the point. A two-week unit on li (ritual propriety) opens with a close reading of a short passage from the Analects and a passage from the Xunzi on the origin of ritual—slow, teacher-led, no devices. Students then enter a reconstructed lineage-hall ceremony in specified roles; the AI supplies the officiant's responses and, when a student deviates, the historically documented correction rather than an improvised one. In the debrief, students receive two machine-generated accounts of the ceremony—one from a general model, one from the curated system—and identify what each gets wrong and why. The unit closes with an essay on whether ritual propriety can survive without the social hierarchy that sustained it. The sequence runs encounter, embodied experience, machine interrogation, argued position; the technology occupies the third position, not the first.

3.4. The Reflexive Layer: Teaching Students to Read the Machine

This layer is the one most often missing from proposals for digital integration, and in the AI era it is the most important. In a course whose content is a tradition of interpretation, the interpretive apparatus becomes part of the content: if students will encounter CETC primarily through machines—and they will, whether or not the course provides it—then the course must teach them how that mediation works. Five practices operationalise this.

Model comparison. Students put the same question to two or three systems (including a retrieval-augmented one built on the curated base) and to a primary text, then account for the differences—cheap, robust, and revealing about how training data and prompting shape representation.

Error archaeology. Instructors seed plausible but fabricated citations, mistranslations or anachronistic attributions; students detect and correct them from primary sources, so that the machine becomes an object of study rather than an oracle.

Provenance auditing and rewriting. For any AI-assisted claim the student supplies a source chain, and given a machine-generated paragraph they rewrite it so that every claim carries a traceable source, deleting or flagging whatever cannot be sourced. This is stricter than conventional citation, because it requires understanding why a claim might be true rather than where a similar sentence appears; the exercise is humbling, since much of what sounds authoritative has no origin they can locate, including some of their own prior knowledge.

Prompt archaeology. Students examine prompts written by previous cohorts and analyse how each reshaped the output: what the framing presupposed, what it made invisible, which assumptions the model accommodated without complaint.

Arguing with the tradition. The hardest practice: students identify a position they find genuinely wrong or inadequate, reconstruct its strongest defence, and argue against it using resources the tradition itself supplies. This is not critique imposed from outside but the tradition's own dynamic—Confucians argued with Mohists, Neo-Confucians with Buddhists, Qing evidential scholars with Song metaphysicians. Teaching students to participate in that argument is the opposite of transmitting a canon, and it is what makes cultural confidence substantive rather than declarative.

3.5. How the Layers Interact

The three layers form a loop rather than a stack. Reflexive work exposes gaps in the semantic layer—students find that the base has no answer to their question, or that its answer conceals a dispute. Those gaps generate new experiential tasks, such as a simulation built around the dispute rather than the consensus. The experiential encounter in turn produces the investment that makes reflexive work feel worth doing rather than like debunking.

Two failure modes follow, and both are common. The semantic-only course produces students who can define ren accurately and have never felt its pull—the failure mode of much examination-driven IPTC teaching, which AI makes cheaper and more scalable. The experiential-only course produces students who remember a spectacular VR session and can articulate nothing about it—the failure mode of much edtech innovation, which generative media makes more seductive. The reflexive layer prevents both, because it forces the student to move between the two and notice what each leaves out.

4. Four Design Principles

4.1. Interpretive Transparency

Every AI-mediated encounter with the tradition should make visible the chain of mediation—source, translation, model, prompt, and the choices embedded in each. The distinctive epistemic hazard of generative systems is not error but unmarked error: a textbook's simplification is at least authored and therefore criticisable, whereas a model's arrives without an author, in fluent prose, with the apparent authority of consensus. Transparency is the only available substitute for authorship.

Implementation means citation-on-demand interfaces, visible translation variants, a "how this answer was produced" panel, and assignments requiring students to reconstruct the chain—designed as a learning object, not a compliance feature. Two qualifications matter. Transparency is not explanation: a system showing which documents an answer drew on has made its inputs visible, not its reasoning; input provenance is nevertheless the more useful of the two, since it lets a student check a claim against a source. And transparency has a cognitive price, since a student who sees five competing translations before reading any of them may be paralysed rather than enlightened. It is a curriculum, not a switch.

4.2. Cultural Fidelity

The system should represent the tradition's internal complexity—its disputes, its development, its untranslatability—rather than a smoothed composite. Fidelity is often misunderstood as accuracy of detail; for a commentarial tradition it is primarily fidelity to structure, to the fact that the tradition is an argument extended over time. A model presenting Confucianism as a unified doctrine is less faithful than one presenting it as unresolved disagreements, even if less fluent.

Implementation: encode contestation as a first-class data structure; flag anachronism; treat translation as content rather than interface; involve specialists in curation. Three tests distinguish fidelity from its imitations. The dispute test: does the representation include at least one significant dissenting position? If not, it presents a settlement as a description.

The anachronism test: does it preserve the distance between text and present? The absence test: does it say what it is not covering? A faithful representation is necessarily partial, and marking the partiality is part of the fidelity.

4.3. Cognitive Apprenticeship

AI should model expert practice and then withdraw, rather than substituting for practice. Formation in a discipline happens through supervised practice with gradually reduced scaffolding; a system that always produces the answer short-circuits that formation, while one that demonstrates how an expert reads a passage and then asks the student to do it supports it. Implementation: think-aloud exemplars, worked examples with fading prompts, AI as a respondent to student interpretation rather than a source of it, and deliberate restriction of assistance in the phases where the learning occurs.

4.4. Contested-Space Pedagogy

The curriculum should preserve and design for genuine disagreement—between traditions, within traditions, and between the tradition and contemporary life. Value formation requires a space in which positions are genuinely at stake, and generative systems are consensus machines trained to produce the most acceptable response. In a course whose purpose is normative formation this is a serious liability, because the dispositions being formed are precisely those required to hold a position under pressure.

Implementation: structured debate with AI in adversarial roles; assignments requiring students to argue a position they do not hold; explicit treatment of cases where CETC and contemporary norms are in tension, without resolving the tension prematurely. Three such tensions are where a course either becomes real or becomes ceremonial. Filial obligation and geographical mobility: the account of xiao presupposes co-residence and lineage continuity, while contemporary labour markets presuppose dispersal; treating filial piety as an unqualified virtue teaches a slogan, treating it as obsolete teaches contempt, and the teachable question is what the structure of the obligation might look like under dispersal. Hierarchical harmony and equality of voice: he is not agreement but harmony through differentiated roles, and the useful work is distinguishing it from tong (uniformity). Ritual propriety and authenticity: li trains the outward forms of respect on the theory that practice shapes disposition, while contemporary culture suspects form as hypocrisy. Each becomes a stronger classroom episode when an AI asked for a ruling returns a smooth, hedged answer that satisfies no one; the student then has to work out why, which is exactly the analytical task the course exists to set.

4.5. Tensions between Principles

The sharpest conflict is between cultural fidelity and contested space. Fidelity requires the tradition be presented on its own terms, with its authority structures intact; contested space invites students to treat its claims as candidates for rejection. Push the first too far and the course becomes a heritage display; push the second too far and it becomes debunking, in which the tradition is only ever an object of critique. Three resolutions are available: sequencing, establishing the strongest self-account before inviting critique; symmetry, applying the same scrutiny to contemporary normative frameworks—consumer choice, therapeutic self-realisation, procedural neutrality—as to classical ones; and immanence, preferring internal critique, which uses the tradition's resources against some of its inherited positions, to external critique that never engages its reasons. A milder tension runs between transparency and apprenticeship: full disclosure can overwhelm novices, who need to work with a text before they can interrogate a pipeline. The resolution is graduated disclosure—source chain first, then model and prompt, then the design assumptions behind retrieval.

5. Assessment

5.1. Why Conventional Assessment Fails

The default instruments in IPTC—closed-book examinations, short-answer recall, attendance scores—measure the semantic layer at best, and reward exactly the flattening

the framework exists to prevent. A student who writes that "Confucianism emphasises harmony and benevolence" scores well and has demonstrated nothing about cultural understanding. Generative AI also makes recall-based assessment trivially gameable while making gaming harder to detect; the common response—banning the tools or reverting to handwritten examinations—solves the integrity problem by abandoning the curriculum problem. The target can be stated without jargon: a graduate of a well-integrated course should be able to read an unfamiliar passage with a commentary and form an independent judgement; recognise a contemporary projection, an anachronism or a mistranslation; give an account of their commitments that engages the tradition's strongest counter-arguments; and explain how a machine-generated statement came to be, including what it left out.

5.2. A Four-Dimensional Model

Cultural understanding—not recall of doctrines but the ability to read a primary text in context, identify its commentarial tradition, and explain why a concept resists translation; assessed through annotated translation, commentary tracing and source-chain audits. Normative reasoning—the ability to construct, defend and criticise a position using the tradition's resources and to recognise their limits; assessed through argumentative essays with a mandatory counter-argument and through oral defence. Affective identification—the extent to which students experience the tradition as theirs rather than as material to be mastered; the hardest dimension to measure honestly, since self-report is weak, so reflective portfolios assessed for depth of engagement rather than sentiment are preferable, and it should be treated cautiously in summative judgements. Behavioural disposition—whether students act differently in their treatment of family, strangers, public goods and disagreement; direct measurement is impossible within a course, so proxy measures should be used with explicit acknowledgement of their limits.

5.3. Design for an AI-saturated Environment

Four adaptations follow. Process artefacts over products: prompt logs, revision histories and reflection on tool use support integrity and directly assess the reflexive layer. Oral defence as standard: a short viva on any AI-assisted work is the most reliable integrity measure available and forces the student to own the argument. Alternation of AI-resistant and AI-collaborative tasks: some tasks should be designed so that assistance is unhelpful (close reading of a contested passage; interpretation of a ritual in a specific local context), others should require documented critical reflection on AI use; the alternation itself teaches judgement about appropriate reliance. Assessment of the class, not only the individual: where the experiential layer involves collective simulation, weight should attach to the quality of deliberation—whether the group surfaced disagreement, took up minority positions, and reached a defensible rather than a comfortable conclusion.

A further point concerns form. CETC has its own traditions of assessment—the imperial examination, but also commentarial practice, the biji record of reflections, learning by exemplary imitation. A course that asks students to keep a running commentary on a classical text, or to construct a lineage of interpretations with their own position at the end, uses forms the tradition recognises as marks of understanding. This is not antiquarianism but a case for variety: a single instrument measures a single dimension, and a course about cultural inheritance should show something of that inheritance in its own evaluative practice.

5.4. Two Design Vignettes

A general course, no custom technology. A compulsory course for 120 second-year engineering students, one lecturer, four assistants, no budget. Six classical passages are paired with a modern commentary and a contemporary case. Weekly, students obtain an answer to a fixed question from two free chatbots and one primary source, then submit a one-page "source reconciliation" note: where the accounts diverge, which is better evidenced, what they conclude. Assessment: the notes (30 per cent), an argumentative

essay with a mandatory counter-argument (40 per cent), a ten-minute oral defence (20 per cent), structured disagreement (10 per cent). The reflexive layer is fully present, the semantic partly, the experiential minimally; the only extra cost is the lecturer's preparation.

A specialised elective with infrastructure. Sixteen weeks on ritual and civic life, forty students, a virtual-reality laboratory and a departmental base covering two classical texts and their commentaries. Six weeks are seminar, four immersive reconstruction with structured roles, four laboratory sessions in which students construct and defend entries for a contested concept, and two a public-facing exhibit with an accompanying argument. All three layers are present, and the semantic layer is co-produced by students. The risk is that the infrastructure becomes the point; the argumentative requirement is the safeguard. The first design may be pedagogically stronger for its purpose, since the reconciliation note is a more demanding intellectual task than participating in a simulation; both are defensible instances of the same logic.

6. Risks and Their Mitigation

6.1. Semantic and Temporal Flattening

Probabilistic models represent a tradition as a distribution over its textual traces, and the mode of that distribution is almost by definition the most conventional reading. At scale this produces a curriculum that is fluent, agreeable and thin. A subtler version is temporal flattening: models do not sharply distinguish a Han-dynasty text, a Song commentary and a recent conference paper; all are simply "Chinese philosophy". Historical development disappears, everything becomes contemporaneous, and students encounter the tradition as a timeless set of ideas rather than a history of arguments. Mitigation: structured bases with explicit contestation; specialist curation; inclusion of marginal and dissenting texts; periodic audits against scholarly consensus; explicit dating of every source.

6.2. Appropriation of Interpretive Authority

When a system answers questions about ren in fluent, confident prose, it occupies the position of the teacher—patient, available at 2 a.m., never impatient. The risk is not that it is wrong but that it becomes authoritative by default, so that the tradition's interpretive authority is transferred to a system with no standing to hold it. Mitigation: transparency design; explicit teaching about the difference between a source and a synthesis; norms positioning AI within a relationship rather than as its replacement. This is a question of pedagogical authority and cannot be solved technically.

6.3. Over-Reliance and the Atrophy of Interpretive Stamina

A risk invisible in any single lesson: the gradual loss of tolerance for difficulty. Reading a commentary is slow, the first encounter confusing, and understanding arrives, if it arrives, after several passes. A generative system offers a route around this experience, and students reasonably take it. At cohort level, the capacity to sit with an unreadable text—interpretive stamina—is not exercised and therefore not developed; the tool removes the difficulty through which the skill was acquired. The remedy cannot be exhortation but must be structural: device-free reading time, assessment rewarding close attention rather than speedy synthesis, and an explicit account of why the slow route is being taken. Students accept difficulty when they understand its purpose, and rightly reject it when they suspect it is merely traditional.

6.4. Commodification of Heritage Affect

Immersive heritage media is good at producing awe. Awe is not understanding, and it is highly marketable; there is a real risk that CETC becomes an experience commodity, satisfying and memorable and disconnected from the commitments the course exists to form. The heritage literature offers a cautionary parallel, since immersive experiences that prioritise spectacle can produce engagement without comprehension [5]. Mitigation: mandatory analytic framing; assessment rewarding interpretation rather than enthusiasm;

explicit discussion of the difference between consuming cultural affect and inheriting a cultural argument.

6.5. Erosion of Deliberative Space

Generative systems are optimised for smoothness: they defuse conflict, hedge disagreement, and produce answers designed to be acceptable. A course that outsources deliberation to such systems produces students who have never had to hold a position against resistance—the capacity normative education exists to build. Mitigation: contested-space pedagogy; AI in adversarial rather than conciliatory roles; human-only deliberation where the cognitive work is the disagreement itself.

6.6. Equity and Data Sovereignty

Access to high-quality tools, bandwidth and immersive hardware is unevenly distributed; a curriculum depending on them advantages well-resourced universities and reproduces the inequality public education is meant to counter. Tool-centred design also creates a subtler divide within classrooms, between students who can interrogate the technology and those who can only consume it. Separately, community-held heritage carries collective interests, and large-scale ingestion into commercial models raises questions of consent, attribution and benefit-sharing that institutional licences do not resolve. Mitigation: low-bandwidth fallbacks; the reflexive layer treated as core and the immersive as enrichment; AI literacy made universal; institutional data governance and locally hosted models for sensitive material; and discussion of these questions with students, which makes cultural sovereignty a live curricular topic.

7. Institutional Conditions

7.1. The Teacher as the Binding Constraint

No framework survives contact with a teacher who has neither the time nor the mandate to implement it. Programmes emphasising tool operation while neglecting pedagogical reasoning leave teachers able to use a technology but unable to decide when it is appropriate [1]. IPTC teachers carry an additional burden: they are expected to be authoritative interpreters of a tradition the technology represents imperfectly, and asking them also to become critics of that technology without support is unrealistic. Three conditions matter. Time and recognition: building a curated base, designing a simulation and developing an assessment scheme are substantial academic labour, and institutions treating this as an add-on will get add-on quality. A community of practice: shared prompt libraries, peer review of AI-mediated assignments and joint curation convert individual effort into institutional capacity. Permission to restrict: teachers need cover to say "not here, not for this task", especially where institutional enthusiasm for AI pushes the other way.

7.2. A Readiness Ladder

A framework assuming a built knowledge base and an immersive laboratory will be ignored by most departments, so a four-stage ladder is more useful than a single model. Stage 1, critical use of general tools: no custom infrastructure, the reflexive layer implemented through model comparison and error archaeology; minimal cost, addressing naive trust. Stage 2, curated reading packs with constrained retrieval: a departmental base for two or three modules, with source chains and recorded disputes; cost is editorial time, addressing flattening within those modules. Stage 3, designed experience: one or two immersive or simulation-based modules per course, bookended by analytic tasks; cost is technical support, addressing inert knowledge. Stage 4, integrated programme with evaluation: all three layers, four-dimensional assessment, longitudinal evidence.

The point of the ladder is that the reflexive layer is available at Stage 1 and remains the framework's core at Stage 4. Departments that cannot afford the upper stages are not excluded from the argument; they are best placed to show that the pedagogical work does not depend on the technology.

7.3. Curricular Sequencing and the Temptation of Efficiency

The framework implies a sequencing logic different from conventional design: not theory then application, but encounter (a text or practice, read slowly), complication (the commentarial dispute or contemporary tension), mediation (how the machine represents this, and why), and position (the student's argued judgement, defended orally). The machine sits in the middle, as an object of study encountered after the student has something to compare it against.

One further dynamic decides whether any of this happens. AI adoption is often justified on efficiency grounds—larger classes, automated feedback, reduced marking—each legitimate in isolation, but together creating pressure against the framework, because its most valuable components are labour-intensive. If departments count only completion rates, satisfaction and cost per student, drift toward the automatable layer is inevitable. They need to count something else: the share of assessment involving oral defence, the number of primary texts read in full, whether students can reconstruct a source chain, whether graduates two years on can still articulate a position they argued in the course.

7.4. Research Questions

Four questions should be priorities. Does the reflexive layer improve normative reasoning, and at what cost in coverage? Which experiential designs produce durable rather than ephemeral engagement? How do outcomes vary with prior cultural background, and does AI mediation narrow or widen those differences? Do any effects persist beyond the course? Three measurement cautions apply: self-report of cultural identity is highly susceptible to social desirability here; comparison groups are hard to construct honestly when sections are not randomly assigned; and since the framework predicts that the character of learning changes rather than its quantity, qualitative evidence should carry equal weight with scores. A comparative dimension is also worth pursuing—Islamic seminaries confronting automated Qur'anic commentary, Jewish institutions confronting machine-generated Talmud study aids and Catholic universities confronting AI-assisted catechesis face structurally similar questions about who has standing to interpret—since examining how traditions set boundaries around machine interpretation would make the Chinese case intelligible as one instance of a global problem.

8. Conclusion

The integration of CETC into IPTC is usually discussed as a problem of resources and delivery: how to make cultural materials abundant, accessible and engaging. That framing was adequate to an earlier technological moment. It is no longer adequate, because the technologies now available do not merely deliver cultural content—they generate it, interpret it and speak in its voice at scale.

The appropriate response is a curricular one. The framework proposed here—semantic, experiential and reflexive layers, governed by principles of interpretive transparency, cultural fidelity, cognitive apprenticeship and contested space—aims to keep the experienced curriculum honest when the mediating layer is a probabilistic system no one fully authored. Its central claim is that the reflexive layer is not an accessory to cultural education but its contemporary core. In an environment where students will encounter the tradition through machines regardless of what their courses do, teaching them to interrogate those machines is inseparable from teaching them to inherit the tradition.

The argument reduces to three propositions. First, the medium matters more than the message: a generative system does not transmit a tradition unchanged, it re-weights it, and the re-weighting is invisible unless it is designed against. Second, cultural education is a matter of structure rather than volume: adding resources to a course that lacks an interpretive architecture produces abundance without understanding. Third, in a curriculum that exists to form values, the critical layer is not the sceptical enemy of the affirmative one; it is what makes affirmation credible.

Two limitations should be stated. The framework is conceptual and has not been tested empirically; its value will be determined by whether it can be operationalised into measurable designs. And it is written from within a particular institutional and national context, and its claims presuppose a relationship between curriculum and the state that readers elsewhere may reasonably contest. A third is more personal: the author's earlier study, written in Chinese for a domestic readership, organised its recommendations around three pathways—immersive multi-scenario presentation, intelligent personalised recommendation, and multi-interactive smart classrooms with learning-process tracking. Those recommendations remain sound as far as they go, and this article does not repudiate them. It relocates them. Immersive presentation becomes one pattern within the experiential layer rather than a solution in itself; personalised recommendation becomes a delivery mechanism whose value depends on what is recommended and on whether the student can see why; learning-process tracking becomes a data source whose ethical status must be argued rather than assumed. The relocation is not a rejection but a recognition that the question "how do we deliver this better?" was never the most important question, and in the age of generative systems is no longer even the most interesting one.

The question that remains is not whether artificial intelligence can deliver tradition to students. It can, more efficiently than any previous technology. The question is whether students can still be taught to argue with tradition—and with the machines that now speak on its behalf. A curriculum that answers yes is doing something no amount of technological capacity can do for it. A curriculum that does not answer is not really a curriculum of cultural inheritance at all, but a very sophisticated delivery service.

Funding: This paper is a phased research achievement of the General Project of the "Three Campuses" Teaching Reform Research Program of Geely University: Research on the Digital Path of Integrating China's Excellent Traditional Culture into Ideological and Political Theory Courses in the "AIGC+" Era, Project No. 2025SGXYJG045.

References

1. Y. K. Dwivedi et al., "Opinion paper: 'So what if ChatGPT wrote it?' Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy," *Int. J. Inf. Manage.*, vol. 71, p. 102642, 2023, doi: 10.1016/j.ijinfomgt.2023.102642.
2. X. Zhu, "Moral education and values education in curriculum reform in China," *Front. Educ. China*, vol. 1, no. 2, pp. 191–200, 2006, doi: 10.1007/s11516-006-0002-7.
3. C. Wang, "Resurgence of Confucian education in contemporary China: Parental involvement, moral anxiety, and the pedagogy of memorisation," *J. Moral Educ.*, vol. 52, no. 3, pp. 325–342, 2023, doi: 10.1080/03057240.2022.2066639.
4. O. I. Psomadaki, C. A. Dimoulas, G. M. Kalliris, and G. Paschalidis, "Digital storytelling and audience engagement in cultural heritage management: A collaborative model based on the Digital City of Thessaloniki," *J. Cult. Herit.*, vol. 36, pp. 12–22, 2019, doi: 10.1016/j.culher.2018.07.016.
5. J. Barbara, "Re-live history: An immersive virtual reality learning experience of prehistoric intangible cultural heritage," *Front. Educ.*, vol. 7, p. 1032108, 2022, doi: 10.3389/feduc.2022.1032108.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of Publisher and/or the editor(s). Publisher and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.