AI-Enabled Dance Aesthetics in Higher Education: Rationale, Model, and Course Governance

Chen Wang
School of Art & Design, Qingdao University of Technology, Qingdao, China

Abstract: Against the backdrop of New Humanities initiatives and national strategies for education digitalization, the objectives of dance aesthetics education in higher education are shifting from "program-oriented showcasing" to an integrated continuum of cultural understanding — embodied experience — artistic creation — public communication. Artificial intelligence (AI) introduces new methods for classroom feedback, learning analytics, and creative stimulation; however, its value depends on appropriate pedagogical positioning and robust course governance. From the perspectives of embodied learning and multimodal perception, this study delineates AI's functions in dance aesthetics as tutor, co-creator, and evidence base, and proposes a course model comprising goal architecture, content organization, classroom process, platform support, evaluation language, and data governance. Online activities provide interpretable kinematic cues and learning trajectories; offline activities generate meaning and judgement through oral explication, ensemble rehearsal, and public presentation. The two are mutually corroborated within a materials — vocabulary — context loop. Further arguments address cultural authenticity, algorithmic bias, privacy compliance, faculty development, and cross-campus sharing, and offer implementable governance procedures and scaling paths. We conclude that, when human-centered and evidence-oriented principles are upheld, AI can enhance the granularity of feedback, the imaginative range of creation, and the auditability of course archives—without altering the educational ontology of dance aesthetics.

Keywords: Artificial Intelligence; Dance Aesthetics; Learning Analytics; Pose Estimation; Music-conditioned Generation; Course Governance

1. Introduction

Within the current policy environment, dance aesthetics in higher education is entrusted with a stronger educational chain: students are expected not only to *perform*, but also to *articulate, revise,* and *create*, and to assume responsibilities of explanation and communication in public settings. Traditional courses, often driven by short-term stage deadlines, emphasize rehearsal speed and final polish; students may complete routines yet struggle to explicate the semantic, ritual, and regional underpinnings of movement. Feedback typically relies on teachers' real-time observation and demonstration; under large classes and compressed schedules, individualized support is scarce. Assessment has prioritized outcomes over process, leaving formative evidence and public feedback outside the quality-improvement loop.

Recent advances in generative AI and lightweight computer vision provide a technical basis for "visible feedback and timely diagnosis" on mobile and standard PCs (keypoint detection, beat alignment, slow-motion review). Music-conditioned dance generation and editable generative tools enable "sketch-based" experimentation in choreography at low cost. Learning analytics platforms can aggregate logs, clips, and peer review to produce comparable learning trajectories. Yet technology does not automatically yield educational uplift. Without a model aligned to the ontology of dance education and clear boundaries, classrooms can quickly regress into metric-centrism, importing recognition bias and privacy risk. Accordingly, this paper first clarifies AI's educational position; second, it explains how technology should operate *appropriately* across online and offline links [1]; and third, it proposes governance procedures designed for peer scrutiny and replication, avoiding "project-based bustle that resets to zero at term's end."

Our argument proceeds through theory—model—implementation—evaluation—

governance—scaling. The theoretical section draws on embodied learning, multimodal perception, and social constructivism to characterize knowledge in dance learning and to outline legitimate boundaries for AI. The model section offers a structured scheme around goals, content, process, and platform that instructors can use to plan courses, rehearsals, and archiving [2]. The implementation section details how online and offline elements cross-validate, weaving keypoints — clips — explication — works — public feedback into a traceable chain. The evaluation and governance section proposes an evidentiary language and compliance baseline that keep technology in a service-to-learning role. The scaling section considers cross-campus sharing and faculty development for low-overhead dissemination, while maintaining methodological stability. Rather than experimental/control statistics, we center the clarity of course logic and evidentiary chains as the primary criteria for value and reusability in peer contexts.

2 Theoretical Foundations and Value Positioning

2.1 Embodied Learning, Multimodal Perception, and Social Constructivism

Dance learning reconstructs movement—meaning—context. Codified body methods, rhythmic structures, and spatial organizations are not merely visible sequences; they carry communal knowledge, originating in ritual and labor, serving festivals and communities, and continuing through oral transmission. Embodied learning posits that knowledge is not solely propositional; sensations, intentions, and affect in motion are crucial to meaning-making. Multimodal perception shows that vision, audition, and kinesthesia in temporal coupling co-shape the aesthetics of rhythm—space—force. Social constructivism reminds us that judgement in dance is refined not only internally but also in peer gaze, audience response, and the discipline of public space [3]. Effective dance education therefore binds doing with saying, enabling students to elevate embodied experience into communicable conceptual expression and public engagement through a cycle of practice—reflection—re-creation.

2.2 AI's Roles and Limits within the Dance-Education Ontology

Within this ontology, AI's value has scope and limits. First, as a second pair of eyes, keypoint-based pose estimation translates felt deviations into interpretable joint angles, weight shifts, and beat-alignment errors, supporting short-cycle self-correction. Unlike expensive optical motion capture, lightweight models run in real time on common devices, aligning with constraints of cost and latency in university classrooms. Second, as a sketch engine for creation, music-conditioned and editable generation establish an editable mapping between musical structure and movement syntax; students treat model outputs as *material drafts* to be revised back into canonical body methods and rhythmic logic, expressing within the tension of preservation and innovation. Third, as an evidence amplifier, learning analytics can strand together logs, clips, peer review, and public feedback on a timeline—surfacing cohort bottlenecks, visualizing individual trajectories, and providing auditable archives for course governance. Crucially, none of these functions replaces teacher judgement, nor should scores overrule explication [4]. AI's central contribution is to make deviations visible earlier, materials more organized, and works more explicable—not to pronounce artistic value.

2.3 Cultural Authenticity Versus Educational Adaptation

Balancing cultural authenticity and educational adaptation is a continuing question for AI in dance aesthetics [5]. Generative tools can induce style drift and semantic rupture, especially in genres with strong regional or ritual identities; naïve collage can erode internal logic. Classes therefore require explicit norms for *percentage of generative intervention*, *secondary revision steps*, and *work statements*, asking students to specify model names, intervention stages, and human labor, and to justify what is retained or discarded. Once explanatory obligations enter the grading discourse, AI is more likely to become a learning resource rather than a shortcut [6].

3. Model Architecture: Goals, Content, Process, and Platform

3.1 Goal Architecture

The first layer is the goal architecture. Objectives shift from short-term staging to the continuum of cultural understanding — embodied experience — artistic creation — public communication. In cultural understanding, students situate movement semantics, stylistic lineage, and ritual context within historical and regional coordinates, recognizing how the same vocabulary points differently across settings. In embodied experience, the aim is not mere technique but a layered sense of rhythm, order, mutual aid, and improvisatory responsiveness in ensemble rehearsal — understanding how individual expression yields to collective order. In artistic creation, students carry out topic-driven expression grounded in canonical methods and rhythmic logic, producing works that have something to say. In public communication, students use academic and public "bilingual" discourse to guide, answer, and workshop with real audiences, learning to recalibrate stance and strategy under questioning. These layers are not parallel modules but mutually pulling links; every class meeting should manifest their connectivity.

3.2 Content Organization

The second layer is content organization. Courses are woven around theme — vocabulary — materials — context rather than isolated by tool or genre. For example, under the theme "Ritual and Order," online micro-materials preview movement features, meter, and the history of attire and props; offline training explores group rules through formations and transitions, translating the rhythmic skeleton — spatial pathways — mutual aid logic into embodied experience; class then re-contextualizes clips within festival or public-space footage, debating how the same vocabulary expresses or is constrained across contexts, with short written explications as conceptual bases for the next creative round. This design forces students to bind visible movement to articulable meaning, ensuring each learning item points to concrete performative choices.

3.3 Classroom Process

The third layer is classroom process, paced as see—do—say—create. A 10-minute introduction supplies concise historical–regional–ritual cues, avoiding lectural overload. Training decomposes vocabulary factors into observable units; meticulous attention to breath, center-of-mass, and force directions brings students into feelable and sayable embodied experience. Explication then requires students to chain movement pathways and rhythmic logic into defensible meaning, using key-frame screenshots and slow-motion to justify *why this transition*, *why this deceleration*. Creation is topic-led, producing small works rehearsed in ensemble and previewed via "open-box" showings for peer scrutiny. Every phase can be captured as learning clips, archived for cross-week review and cross-course comparison.

3.4 Platform Support and Metadata

The fourth layer is platform support. The platform is not an arbiter of "good" or "bad," but infrastructure for material aggregation and evidentiary trace. Core functions include key-frame clipping and slow-motion playback; visual prompts for beat-alignment and angular deviations; structured collection of logs and peer comments; paired archiving of works and explication texts; and weekly learning trajectories. It also manages metadata, indexing by *genre — vocabulary — rhythm — prop — context — region* to enable inter-course corroboration and inter-term traceability. The design emphasizes interpretability over mystique: prompts are delivered as graphics and prose, not opaque black-box scores.

4. Classroom Implementation: Cross - Validating Online and Offline

4.1 Online Activities: Lightweight, Immediate, Interpretable

Implementation ensures online data are the starting point for offline dialogue, not the endpoint. Online activities prioritize lightweight, immediate, interpretable support. Students record key movement clips on mobile; the system returns schematic prompts on joint-angle deviations, weight shifts, and beat error, and converts key frames into learning clips archived in personal dossiers [7]. Trajectories visualize improvement curves alongside micro self-assessments; instructors perform group-level diagnostics before class to set priorities. Recommendations obey a minimal-necessary principle to prevent over-consumption of materials at the expense of bodily work; each automated suggestion includes a brief why, linking how to adjust with why it matters [8].

4.2 Offline Activities: Presence of the Body and Ensemble Rehearsal

Offline, the body in presence remains central. After vocabulary drills, students give oral explications that bind body method, rhythm, and ritual/prop logic into defensible narratives. Ensemble rehearsal uses open-box sessions for semi-public viewing so students recalibrate expression under the tension of being watched. In creation, music-conditioned generation provides an editable sketch, which is then revised back into canonical methods and rhythmic logic; every retention and omission must be justified. The instructor, as editor–dialogue partner, redirects attention from "are the moves neat" to "is the expression apt, the narrative clear, the style pure." For genres with strong regional identity, classes use lineage charts and canonical footage comparisons to warn against drift, while slow-motion review and key-frame analysis locate sources of error.

4.3 Restrained Use of Computer Vision and Threshold–Review Mechanism

To minimize disruption, computer-vision use is deliberately restrained. Students receive a "capture quick guide" on viewpoint, lighting, attire; avoid backlighting/mirrors, keep full-body in frame with floor markers, and segment high-spin passages. For clips with high recognition error, slow-motion + human review is preferred over model scoring to prevent statistical noise being misread as personal deficit. The platform adopts a threshold-plus-review mechanism: when beat error exceeds a threshold, entries are flagged "human review required" for peer/TA verbal explanation and demonstration.

4.4 Text Work and Traceable Knowledge Blocks

Text work is equally important. Each clip pairs with a short movement–context note explaining *why a transition, why a tempo change, why a given prop*. Notes are archived alongside video, instructor comments, and public feedback, then cited or revised in subsequent creation. Over time, classes produce traceable knowledge blocks that support inter-term continuity and cross-course corroboration. In this sense, AI does not merely make class "high-tech"; it organizes scattered experience into discussable evidence, pushing teachers and students toward higher-quality dialogue on why [9].

5. Evaluation Language and Data Governance

5.1 Evidence-Centered Evaluation

Evaluation centers on evidence, not scores. Process evidence comes from logs, clips, and classroom dialogue, marking concrete *deviation—adjustment—reproduction* cycles. Synthetic evidence comes from in-class showings and work statements that articulate causal links among movement semantics, ritual context, and creative choices. Reflective evidence comes from field notes and oral-history entries documenting interactions with community members and artistic trade-offs. Social evidence comes from public feedback and media echoes, initiating the next iteration. Quantitative indices—beat alignment, trajectory smoothness, ensemble-order stability—are confined to formative feedback roles to structure review and discussion; decisive judgement returns to whether explanations hold, evidence suffices, and feedback is absorbed. To avoid vague adjectives, the evaluation language reframes generalities into observable entrances for discussion—e.g., "good coordination" is unpacked into formation stability, beat-error convergence, and transition fluency; "cultural understanding" into semantic accuracy, contextual appropriateness, and historical clarity.

5.2 Data Governance, Privacy, and Fairness

Data governance sets the baseline. Classes follow minimal necessity, purpose limitation, de-identification, consent logging, and scheduled deletion. All filming requires explicit consent; exports default to masking and metadata suppression. Generative creation must disclose model names, intervention points, and human labor, clarifying authenticity and authorship; syllabi specify an upper bound for generative intervention and the required secondary revision steps [10]. For fairness risks, classes run algorithm check-ups, sampling recognition error across attire, body types, and lighting; grade sheets include a standing reminder that "AI indices are reference evidence only," and provide appeal/re-assessment channels. Annual platform reports include usage, correction records, and deletion logs in school-level quality governance, curbing both technological overreach and data sedimentation.

5.3 Faculty Development through Evidence-based Reflection

The evaluation mechanism doubles as faculty development. Each term, instructors compile problem lists from archives, reflecting on *why cohort-wide deviations occur, which prompts worked best, which content organizations sparked the richest dialogue*. Schools then convene cross-course preparation to consolidate shared task scripts and explication templates, moving method from personal experience toward consensus and norms, reducing variability across instructors and clarifying pathways for new colleagues.

5.4 Evidence Synthesis and Narrative Judgment

Evaluation is framed as a process of building arguments from artifacts rather than extracting numbers from instruments. Each learning artifact—classroom dialogue, learning clips, movement–context notes, rehearsal photos, oral-history snippets, and audience comments—retains a stable identifier so that claims can be anchored to verifiable sources. A judgment thus reads as a short narrative that states a claim, cites the specific artifact(s) that warrant it, and explains the reasoning that connects evidence to conclusion. When the artifact is a clip, timecodes are quoted; when it is public feedback, the citation records date, venue, and the respondent's role. Assessors are encouraged to seek counterevidence in the same archive, asking whether an alternative clip or statement undermines the claim; acknowledging contrary traces raises the credibility of the final judgment and keeps it dialogic rather than accusatory.

To keep judgments comparable across instructors without forcing homogenization, the program adopts an "anatomy of a judgment" template: a one-paragraph statement that begins with the central observation, follows with the chain of evidence and counterevidence, and ends with the pedagogical implication—what should be adjusted next and why. This template travels well across genres and styles because it focuses on reasoning with evidence, not on style compliance. It also helps students read feedback productively: they learn to track how their own texts, clips, and public responses are being mobilized in argumentation, and to respond with revisions that point back to the same identifiers. Over time, the archive becomes a network of linked claims and artifacts, enabling longitudinal inquiry into what kinds of tasks, prompts, and rehearsal arrangements most often lead to defensible work. The language of evaluation thereby doubles as a method of research: teachers and students co-construct a corpus where artistic decisions can be traced, challenged, and improved without resorting to reductive scoring.

5.5 Data Operations, Risk Response, and Sustainability

Day-to-day operation follows a lightweight cycle that faculty and students can remember under pressure: capture, annotate, discuss, revise, and archive. Consent is explicit and renewable whenever task types change; de-identification is applied at the moment of export; every item is time-stamped with a planned deletion date so that end-of-term cleanup is procedural rather than heroic. When devices fail or models misbehave, the class falls back to analog means—demonstration, peer mirroring, and hand-drawn pathway sketches—so that learning remains embodied and continuous. Students may opt out of public sharing without penalty; alternative paths focus on private peer review, instructor conferencing, and reflective writing attached to their nonpublic artifacts.

Risk is managed through transparent versioning and fairness checks. Pose-estimation models are pinned to versions for a full term; updates are accompanied by plain-language change logs that clarify expected effects on accuracy and latency. Periodic "algorithm check-ups" sample recognition error across lighting, attire, and body types; findings are circulated to instructors with suggested mitigations in camera placement and pacing so that statistical artifacts do not become personal deficits. Appeals and second readings are resolved within a fixed turnaround, and anonymized case notes reenter the materials bank as cautionary examples. Sustainability depends on workload: the platform automates low-value actions—time-stamped clipping, template generation for movement–context notes, and batch feedback—so that instructors devote their energy to interpretation, coaching, and redesign. Teaching assistants are trained as evidence stewards who help students select clips that best display change through time and who monitor the integrity of identifiers across revisions. When consent permits, small pairs of clip + explication are curated as open micro-assets for other cohorts, allowing learning by comparison rather than repetition. Success is not measured by the quantity of captured data but by two practical questions: can students defend their artistic choices with evidence, and can teachers improve courses because that evidence exists.

6. Conclusion and Routes for Scaling

6.1 Conclusion

AI enters dance aesthetics not to replace teachers or students, but to make deviations visible earlier, materials more organizable, and works more explicable. With tutor, co-creator, and evidence as its core position, technology returns to its service-to-learning role: it supports correction with keypoints and trajectories, stimulates topic-led imagination with music-conditioned sketches, and undergirds auditable quality improvement with learning traces and archives. The model presented here offers a replicable structure for goals, content, process, and platform, and a discussable language and baseline for evaluation and governance—operable under large classes and limited contact hours without additional institutional cost.

6.2 Scaling Routes and Institutionalization

Scaling hinges on regional cross-campus networks. A workflow of collection—annotation—analysis—visualization—sharing can convert classroom materials into public assets, reduce duplication, and accelerate textbook renewal; metadata centered on vocabulary—rhythm—context enables interoperable course archives. Faculty development should braid AI literacy, joint course preparation, and action reflection into routine practice with a rhythm of tool colloquia—thematic didactics—evidence reviews, maintaining methodological stability while preserving critical distance from tools. Institutionally, fairness testing, privacy protection, and appeal mechanisms should be codified in course agreements and grading language, while platform operations and archive quality should enter annual quality reports — normalizing a evidence → improvement mechanism. Grounded in culture, problem-driven, evidence-based, and public-facing principles, AI can become a steady propellant for clarity and maturity in university dance aesthetics—without becoming the protagonist.

Acknowledgments

This work was supported by the Shandong Provincial Undergraduate Teaching Reform Research Project, *Culture-Led and Virtue-Oriented Aesthetic Education: Research and Practice on Integrating Fine Traditional Chinese Culture into the Construction of Aesthetic-Education Curricula* (M2022274), and by the Qingdao University of Technology Higher Education Research Special Project, *Integrating Excellent Traditional Dance Culture into University Aesthetic Education under the New Liberal Arts Initiative* (FG202415).

References

[1] Wang, Z. (2024). Artificial intelligence in dance education: Using immersive technologies for teaching dance skills. *Technology in Society, 77*, 102579.

[2] Kang, J., Kang, C., Yoon, J., Ji, H., Li, T., Moon, H., Han, J. (2023). Dancing on the inside: A qualitative study on online dance learning with teacher–AI cooperation. *Education and Information Technologies, 28*(9), 12111–12141.

[3] Söö t, A., & Viskus, E. (2013). Teaching dance in the 21st century: A literature review. *The European Journal of Social & Behavioural Sciences*.

[4] Xu, L. J., Wu, J., Zhu, J. D., & Chen, L. (2025). Effects of AI-assisted dance skills teaching, evaluation and visual feedback on dance students' learning performance, motivation and self-efficacy. *International Journal of Human-Computer Studies, 195*, 103410.

[5] Wang, Y., & Zheng, G. (2020). Application of artificial intelligence in college dance teaching and its performance analysis. *International Journal of Emerging Technologies in Learning (iJET), 15*(16), 178–190.

[6] Li, X. (2021, April). The art of dance from the perspective of artificial intelligence. *Journal of Physics: Conference Series, 1852*(4), 042011.

[7] Xu, S., Rahim, N., & Yahaya, W. A. J. W. (2025). The impact of hybrid learning integrating mobile interaction and AI timely feedback on students' dance performance, perceived motivation, and engagement. *An-Najah University Journal for Research – B (Humanities), 39*(10).

[8] Wang, X. (2024). Artificial intelligence-driven training and improvement methods for college students' line dancing. *International Journal of Advanced Computer Science and Applications, 15*(1).
[9] Miko, H., Frizen, R., & Steinberg, C. (2025). Using AI-based feedback in dance education: A literature review. *Research in Dance Education*. Advance online publication.
[10] Wang, F. (2025, September). Construction of street dance basic movement library based on artificial intelligence and analysis of teaching paths. In *Proceedings of the International Conference on Mechatronics and Artificial Intelligence (MAI 2025)* (Vol. 13795, pp. 426–433). SPIE.