Dai Xudong

Personal Statement

Master of Arts in Arts and Machine Creativity
The Hong Kong University of Science and Technology

daixudong.art | github.com/xudongdai-ai

I am a composition student interested in how computational systems can respond to musical intention rather than merely generate more material. I am completing a five-year undergraduate programme at Xi'an Conservatory of Music, with graduation expected in June 2027. Studying harmony, counterpoint, orchestration and form has taught me to hear structure as relationships between musical events. In my orchestral work *The Vanished Kingdom*, I shape these relationships through register, instrumental colour and articulation. Yet a text prompt is often too coarse to specify a motif's return or revise a particular passage. I want composers to be able to inspect and control such decisions directly.

Alongside my five-year composition degree, I am pursuing a second undergraduate course of study in Computer Science and Technology through Sichuan University's Higher Education Self-taught Examination pathway (2025-present). This sustained study in programming and software development supports my work with Python, TypeScript, FastAPI and Electron: I want to implement and evaluate musician-facing systems, not only use existing AI tools. I also completed Xi'an Conservatory's first-cohort Micro-major in Music Industry Planning and Management, connecting technical design with copyright and distribution. At Stanford CCRMA's online summer workshop, *Perception-Based Orchestration for Composer-Performers*, I considered how auditory perception informs compositional choices. My assessed technical certificates include Agent Engineering through Datawhale and Ant Group, LLM Fine-tuning through Datawhale and iFLYTEK, and LLM Development Engineering through Datawhale and Inspur. Additional training with Datawhale partners covered AI-assisted programming (MarsCode) and AI data analysis (SenseTime). Together, my second undergraduate pathway, music-industry studies and technical training connect musical judgement with the software skills needed to build, test and refine creative systems.

Algorithmic Awakening made this connection tangible. I wrote programs to generate twelve-tone and random-sequence material, then developed selected fragments into the sonic narrative of a system awakening. Early sketches were rigid and sparse. I separated registers, added sustained layers and varied onset density; when a later version became harsh in the upper frequencies, I reduced arpeggiator density and high-frequency accumulation. In Studio One, I reshaped the buildup and expanded sustained tones to clarify the final arrival. The work received the Best AI Music Production Award at the 2nd HKUST AI Film Festival x MiniMax in 2026. This process taught me that the decisive creative act often lies in selecting, revising and structuring generated possibilities.

I founded Xunpu Music Creation Technology Co., Ltd. to develop AI music software, with *Xunpu Studio* as a central music-learning, analysis and composition platform. As Founder and Product Research Lead, I translate composition pedagogy into MusicXML/MIDI and OMR workflows, harmony and counterpoint checks, form suggestions, score annotations and explanations that musicians can review. An eight-week closed pilot involved 50 participants and recorded 1,247 core-function calls. The project asks how a musician can locate a system's evidence and revise or reject an annotation inside the score. Leading *Xunpu Studio* also brought a Shaanxi Silver Award in the China International College Students' Innovation Competition and a provincial-level student innovation project. Building and evaluating these tools has made computing an integral part of my musical research rather than a separate technical interest.

Within this venture I am also developing *Xunpu AI Synth*. Its playable browser prototype translates Chinese or English timbre descriptions into editable synthesis parameters, with A/B comparison and WAV export. Integration as a DAW plug-in is still in development. The gap between a verbal request and a sound that can be played, compared and refined has made controllable timbre an especially compelling interface problem for me.

My independent *Musicology-Guided Structural Attention* (MGSA) research examines whether motivic, harmonic, cadential, repetition and metric priors can inform measure-level boundary detection in symbolic music. I introduced a cross-boundary penalty into attention logits and evaluated a five-seed Mozart Fold 0 pilot. Exact F1 was 0.3332 versus 0.3297 for a matched control, while within-one-measure F1 was 0.4569 versus 0.4588. These small, mixed differences do not establish a consistent advantage; full grouped evaluation is the next step. I value this result because it requires me to distinguish a musically plausible mechanism from demonstrated improvement.

That emphasis on inspectable evidence also informs my *Music Segment Boundary Toolkit*, which combines AI references, waveform-energy cues and lyric timing for human confirmation, and *Xunpu GraphOMR*, which uses typed symbol graphs and rule-based decoding. As a co-author of *ReCoP*, a Shaanbei folk-song lyrics retrieval manuscript submitted to TISMIR, and a contributor to the HKUST AI Audio Annotation Project, I have worked with the demands of data representation, quality review and evaluation. I want this discipline to shape creative systems, not remain separate from them.

My second venture, Xi'an Huijing Shengchang Intelligent Technology Co., Ltd. (Huijing Soundscape), extended these questions into AIGC short-form drama. I founded and led a 25-person team coordinating scripts, storyboards, generated audiovisual material, music, editing and human review. At that scale, creative intent can disappear between stages unless choices remain traceable and revisable. Internships at Taihe Music and Xinjiang Cultural Tourism Investment Group, including a separate commission to compose *Kucha*, added experience of real music-industry and production settings. Leading the Rhythm Intelligence project to a provincial Second Prize in the Challenge Cup further developed my ability to turn an interdisciplinary idea into a coordinated project.

At HKUST, I propose to develop an interactive composition prototype and use it to create an original electroacoustic work. A composer would select symbolic material, alter its register or rhythmic density, audition timbral alternatives and compare versions while retaining a visible edit history. I would build first on editable material and reviewable, rule-based annotations, with any MGSA-derived suggestions treated as experimental. *Xunpu AI Synth* offers a practical starting point for linking verbal timbre intentions with audible, adjustable outcomes alongside symbolic editing. I would evaluate the prototype through task-based creator feedback: can composers understand suggestions, preserve intended details and make the changes they actually want?

HKUST's MA in Arts and Machine Creativity offers a setting where technical prototypes can be judged as artistic work as well as computational systems. In AMCC 5220 Technology in Music and Sound Art, I hope to investigate synthesis and processing as ways to render symbolic revisions audible. AMCC 5410 would help me examine authorship when composers select, transform and reject generated material. AMCC 5330 could guide a credible presentation and distribution plan, including copyright and production costs. Subject to the Programme Director's approval, AMCC 6100 could provide a setting to combine the prototype, composition and reflective documentation. These connections between making, critique and public presentation are central to the education I seek.

After the MA, I intend to continue developing musician-centred AI software through Xunpu Music Creation and to pursue research on musical intention, controllable generation and, eventually, adaptive accompaniment. My immediate aim at HKUST is a small but convincing creative work showing that computational assistance can remain responsive to a composer's judgement.