HANZE JIN

scyhj5@nottingham.ac.uk | Nottingham, UK (Jubilee Campus) | jinhanze.com | github.com/KIM-ocean-HZ

PROFILE

Final-year BSc Computer Science with Artificial Intelligence student at the University of Nottingham who builds products that keep people in control of what they create: Spool, a released local-first desktop app that governs what AI assistants may read and write, and FLUX, a final-year project on how a pretrained music model can help musicians develop their own MIDI drafts. PyTorch research experience in computer vision; interested in controllable generative models and human–AI interaction.

EDUCATION

BSc (Hons) Computer Science with Artificial Intelligence, University of Nottingham *Sep 2023 – Jul 2027 (expected)*

2+2 programme (Ningbo → UK). Credit-weighted averages calculated from the transcript: 76.0% (2024–25, Ningbo) and 81.3% (2025–26, UK); the First-Class threshold is 70%. TOEFL iBT 104 (Sep 2025).

- **Relevant modules:** Developing Maintainable Software (93%), Fundamentals of AI (89%), Introduction to Formal Reasoning (88%), Systems & Architecture (88%), AI Methods (83%), Algorithms, Data Structures & Efficiency (83%), Operating Systems & Concurrency (78%), Image Processing (76%).

UI/UX Design Summer School, National University of Singapore *Jun – Jul 2024*

- Grade A; our team's final project received the highest mark in the class. Led feature planning and interaction design for a campus-map mobile prototype, built in Figma.

SELECTED PROJECTS

Case studies at jinhanze.com

Spool — Local-First Context Hub (desktop app) *May 2026 – Present*

Latest release v0.7.4 (26 Oct 2026), Developer ID–signed and Apple-notarised for Apple Silicon Macs; a Windows build is in testing · spoolapp.org · github.com/KIM-ocean-HZ/spool

- Sole developer of a desktop app (Tauri 2 / Rust, React/TypeScript, SQLite with FTS5 trigram search) that keeps each long-running project as a local, time-ordered log and hands it to AI assistants as plain-text briefings, through an MCP server, or via the user's own AI CLI.
- Kept the user in charge of AI access: MCP clients get no access until the user turns the server on, cannot edit or delete what the user wrote, and any change they propose to existing content waits for the user's approval; the main process links no HTTP or TLS library.
- Designed and adjudicated six evaluation rounds (557 real API calls, at least ¥29.37) before opt-in AI compression shipped in v0.6.2; fixed and re-tested the four release blockers they found, and dropped a price estimate that landed within ±30% for only 46% of 89 real bills, against a 90% bar set in advance.

FLUX — Symbolic Music Co-Creation *May 2026 – Present*

Individual final-year project · github.com/KIM-ocean-HZ/flux

- Built a browser MIDI workbench (React, FastAPI) for multi-track editing and playback, with chord detection and key inference that stay provisional until the musician confirms them (102 unit tests).
- Preliminary study (Oct 2026): with pinned MIDI-GPT weights, on a development set of 25 Lakh MIDI excerpts (332 bass requests, three seeds each), writing the chords into the context as a guide track met 41% of unusual harmonic requests, against under 1% for the built-in controls, as scored by a rule-based verifier rather than listeners.
- Dissertation (proposal submitted Oct 2026): compares structure-preserving variation with context-only rewriting of a musician's own draft, scoring only the notes the model writes against position-matched non-learned baselines, alongside a self-supervised part–accompaniment fit model and a study with musicians (subject to ethics approval).
- Implemented both modes with note-level provenance, the baselines and context interventions, and ran a first probe on development drafts; built a study version of the workbench that logs everything a musician hears, copies and changes (demonstrated Oct 2026).

SmartTravel — Travel Planning App (team project) *Oct 2025 – May 2026*

COMP2002 Software Engineering Group Project · eight-person team

- Proposed the app on top of the module's recommender brief, defined its structure and feature set, and designed the interface in Figma; built the swipe-card onboarding, sign-up and new-trip flows in Expo React Native, and rebuilt TripBag and Schedule.
- Built the behavioural event tracking the team's recommender learns from (views, dwell time, taps, TripBag changes), with batching and offline caching, plus the backend endpoint and schema that store the events.

MemoHead — On-Device Personal Voice Archive *Dec 2025; Oct 2026 – Present*

- Built a Mac prototype that transcribes and identifies speakers entirely on the device: faster-whisper on the CPU (int8), pyannote speaker embeddings matched by cosine distance against several voice templates per speaker, and ChromaDB as the memory store (Python, tkinter).
- Rebuilding it as a personal voice archive that aims to keep only its owner's speech: a Flutter shell (macOS first) over a local Python engine, with a harness comparing on-device speech-recognition models (faster-whisper, whisper.cpp, Qwen3-ASR) on error rate, hallucination on silence, speed and memory.

RESEARCH & INDUSTRY EXPERIENCE

R&D Intern, Libiao Robotics *Summer 2026*

Simulation-based multi-robot warehouse control (C#/.NET 8)

- Sole developer, to a brief set and reviewed by the R&D team: built a simulator and controller covering its 13 tasks, from grid maps, kinematics and Dijkstra/A* planning to multi-robot traffic control, sorting tasks and WES–RCS messaging over HTTP, UDP, TCP and gRPC.
- Found that the first version mixed traffic state into the map and that its cell locks deadlocked robots meeting head-on; rebuilt the core as six layers guarded by architecture tests that fail the build if a layer depends on a higher one, so replacing GDI+ with SkiaSharp changed only the rendering layer.
- Rebuilt traffic control around spatiotemporal reservations, with each robot booking its whole route in space and time before moving; 20 robots ran a seeded 30-minute simulation (90,000 ticks) without a collision.
- On my own initiative, added a PPO policy that sets each robot's speed along its A* path; in fixed-seed runs it cleared all conflicts and cut mean task time from 35.4 s to 26.6 s with four robots, though conflicts rose with six.

Remote Research Assistant to Dr Qinghua Li, then at Northwestern University *Jun 2025 – Jan 2026*

The Beauty Bias Effect in Scientific Careers

- Tested whether off-the-shelf tools (DeepFace, Face++, Amazon Rekognition) could supply the study's facial attributes; against self-reported labels for 562 Chicago Face Database faces, DeepFace's gender labels were right for almost all men but only about one woman in five.
- Trained a gender classifier in TensorFlow/Keras; when the from-scratch CNN did not transfer to everyday photos, moved to two-stage MobileNetV2 fine-tuning (1,207 CFD → 10,000 CelebA images) with augmentation and train/validation/test splits, reporting accuracy, F1 and confusion matrices, and accuracy by ethnic group.
- Broke down the tools' errors on 500 CelebA images by expression, head pose, mouth, age and apparent ethnicity, to judge whether they could bias the study's conclusions.

Research Assistant (FURP), University of Nottingham Ningbo China — Dr Boon Giin Lee *May – Aug 2025*

Monocular Depth Estimation with Deep Learning

- Ran four recent monocular depth models (UniDepth, ZoeDepth, Depth Anything V2, Depth Pro) in PyTorch, in local notebooks and on Colab.
- Tested Depth Pro and Depth Anything V2 on iBims-1 under Gaussian and shot noise, 14 style transfers and both combined, measuring drift from each model's own clean-image output; the two degraded very differently.
- Wrote a survey of depth-estimation methods across input types such as video and thermal images, organised by depth type, learning paradigm, architecture and dataset.

IT Intern, Industrial Bank Co., Ltd *Dec 2023 – Feb 2024*

- Refactored backend logic in a Python/Flask app, optimised its SQLite queries and fixed front-end rendering bugs.

MUSIC & LEADERSHIP

Composer, Producer & Performer (Independent) *2020 – Present*

- 5+ years writing, arranging, performing, producing, mixing and mastering original songs across post-rock, R&B, neo-soul, film-score ambient and electronic rock; 30+ original tracks released since 2023; two years performing in a band. Selected tracks: jinhanze.com/music.

Chairman, Modern Music Association, University of Nottingham Ningbo China *Sep 2024 – May 2025*

- Ran the association's day-to-day operations: recruiting members, scheduling rehearsals, planning events and liaising with sponsors.
- Organised campus performances, including the Campus Acoustic Music Festival and the UoN Ningbo Band Special Performance, and an off-campus commercial show at K11, liaising with venues and performers and running the day.

TECHNICAL SKILLS

Languages & Tools: Python (proficient), Rust, TypeScript/JavaScript, C#/.NET, C, Java, SQL, MATLAB, Lean 4; Git, GitLab CI/CD, Linux, Tauri 2, React, React Native (Expo), FastAPI, SQLite/FTS5, MCP, Figma, LaTeX.

Machine Learning: PyTorch (proficient), TensorFlow/Keras, Hugging Face Transformers; CNN, ResNet, LSTM, Transformer, reinforcement learning (PPO); OpenCV, MediaPipe, NumPy, Pandas, scikit-learn, Matplotlib.

Audio & Music: Librosa, torchaudio, faster-whisper, whisper.cpp, pyannote, sounddevice; MIDI-GPT, Anticipatory Music Transformer, muspy, MIDI tooling; Logic Pro, Ableton Live, Cubase, Pro Tools.