Jingzhi Bao

jingzhibao03@gmail.com | jingzhi03.com | ZqlwMatt | (+86) 187-5877-9575

Education

University of Illinois Urbana-Champaign — Sep. 2025 – Dec. 2026 (Expected)
Master in Computer Science — Champaign, IL, USA

The Chinese University of Hong Kong, Shenzhen — Sep. 2021 – June 2025
BEng. in Computer Science and Engineering (First Class Honor) — Shenzhen, China

Nanyang Technological University — Aug. 2023 – Dec. 2023
Exchange Student of Computer Science and Engineering — Singapore

Experience

ByteDance — May 2026 – Aug. 2026
Research Scientist Intern - LLM/UMM Post-Training — San Jose, CA, USA

- Research topic: model merging in post-training, data construction for unified multi-modal model (UMM).
- **[Data]** Designed a hierarchical data construction framework to improve text rendering, decomposing text-rich visual content into a 4-level hierarchical supervision schema spanning global context, visual elements, inter-element relations, and OCR details. Curated a 70K+ multidisciplinary training corpus from infographics, articles, packaging designs.
- **[Model]** Conducted large-scale distributed training (GRPO/OPD) of Qwen3-based unified multimodal models (8B–235B) with Megatron-LM and vLLM. Optimized the FLUX.2 VAE for text-dense visual content, improving PSNR by ~3.

MeshyAI — Oct. 2025 - May 2026
Research Scientist Intern (Star Intern) - General 3D Generation — Sunnyvale, CA, USA

- Core contributor to Meshy 6/7 commercial models; Project lead for 8K PBR texture full-stack pipeline.
- **[Model]** Proposed UltraTex (TOG 2026), a VLM- and FLUX.2-guided 8K PBR texturing model that outperformed commercial methods by >83% in pairwise studies (results), achieving professional artist-level quality with fast inference speed (~20s).
  - Elicited PBR intrinsics from VLM and introduced photometric representation alignment to guide PBR generation.
  - Collaborated on VAE-DiT training with hybrid representations for a joint 3-stage geometry-texture generation (Meshy 7).
  - Investigated on joint text, image and 3D understanding and generation in a unified multi-modal architecture.
  - [Distillation] Developed distillation-guided VAE pruning and decoupled DMD, reducing inference time by 86%.
- **[Data]** Improved the data pipeline with K8s+Celery+AWS+Databricks, supporting 2M assets rendering within a day.
  - Trained a prompt enhancer (PE) for dense captioning and CoT-based data captioners for PBR, aesthetic, and detail quality.

Lightspeed Studios, Tencent IEG — Feb. 2025 – Aug. 2025
Research Intern - General 3D Generation (Outstanding Project) — Shenzhen, China

- Core contributor to 3D texturing model R&D, delivering product-ready quality surpassing Tripo2.5 and Meshy-5.
- **[Model]** We proposed LumiTex, CaliTex (ICLR2026, CVPR2026), a LoRA-finetuned FLUX.1 model with calibrated attention trained over 120 GPU days for multi-view generation, as well as a decoder-only transformer LVSM for novel-view completion via test-time training.
- **[Data]** Rebuild the 3D data cleaning and rendering workflow with `S3` and TencentCOS for the 3D generation team.
  - Rendered a 3D dataset of 92K objects (curated by aesthetic/clarity filters) into high-res images and G-buffers (20 TB).

Future Network of Intelligence Institute, CUHK-Shenzhen — April 2024 - Dec. 2024
Research Intern - Inverse Rendering — Shenzhen, China

- We proposed photometric inverse rendering (3DV2025), a framework integrates model self-shadows, and inter-reflections with DINO-distilled priors for GGX-based rendering, achieving SoTA PBR material reconstruction results.
- We proposed Tex4D, a novel method that takes an untextured, animated mesh sequence along with a text prompt as inputs, and generates temporally consistent 4D textures, enabling dynamic character generation with video diffusion model priors.

Publications

- **UltraTex: Aligning Physical Knowledge for High-Fidelity PBR Texture Generation**
  Jingzhi Bao, Xiaohe Ma, Hong Li, Yuanming Hu, Xiaoguang Han
  [Project Page] [Demo] SIGGRAPH Asia (TOG) 2026. (Meshy 8K PBR Texture Product)
- **CaliTex: Geometry-Calibrated Attention for View-Coherent 3D Texture Generation**
  Chenyu Liu, Hongze Chen, Jingzhi Bao, Lingting Zhu, Runze Zhang, Weikai Chen, Zeyu Hu, Yingda Yin, Keyang Luo, Xin Wang
  [arXiv] [Project Page] CVPR 2026.

- **LumiTex: Towards High-Fidelity PBR Texture Generation with Illumination Context**
  Jingzhi Bao*, Hongze Chen*, Lingting Zhu, Chenyu Liu, Runze Zhang, Keyang Luo, Zeyu Hu, Yingda Yin, Weikai Chen, Xin Wang, Zehong Lin, Jun Zhang, Xiaoguang Han.
  [arXiv] [Project Page] [Code] ICLR 2026.
- **PIR: Photometric Inverse Rendering with Shading Cues Modeling and Surface Reflectance Regularization**
  Jingzhi Bao, Guanying Chen, Shuguang Cui.
  [arXiv] [Project Page] [Code] 3DV 2025.
- **Tex4D: Zero-shot 4D Scene Texturing with Video Diffusion Models**
  Jingzhi Bao, Xueting Li, Ming-Hsuan Yang.
  [arXiv] [Project Page] *arXiv preprint*, 2024.

Honors and Awards

Dean's List, School of Data Science, CUHK-Shenzhen	2021, 2022, 2023, 2024
CUHK-Shenzhen Academic Performance Scholarship, Class A,B,C ($20,000, awarded to rank 1)	2022, 2023, 2024
China Undergraduate Mathematical Contest in Modeling, Leader, First Prize	2023
China Collegiate Programming Contest (CCPC), Leader, Silver Medal	2022

Skills

- **Programming Languages:** Python(PyTorch), C++, TypeScript.
- **Technologies**: Blender, Cocos Engine, Unity, UE, Mitsuba, COLMAP.
- **Language**: English (TOEFL 102), Chinese (Native).

Last Updated in August, 2026.