Guangbo Yu

guangboy@uci.edu ◇ (562) 523-6459 ◇ Github ◇ Google Scholar ◇ Irvine, CA ◇ English (Fluent), Mandarin (Native)

Production NLP engineer (Tencent: **ALBERT** NER at **500K+ DAU**) turned ML researcher: **machine perception**, **multimodal foundation models**, **GRPO/RLVR** VLM post-training; **24** peer-reviewed papers (11 first/co-first), incl. *ACL*.

Education

University of California, Irvine — Irvine, CA
Ph.D., Biomedical Engineering — Expected 12/2027

University of Southern California — Los Angeles, CA
M.S., Computer Science — 2017

University of Electronic Science and Technology of China — Chengdu, China
B.S., Software Engineering — 2015

Core Skills

- **Machine Learning Research**: Computer Vision/Machine Perception, Multimodal Foundation Models, Self-Supervised Learning, Post-Training (SFT/RLVR), Natural Language Understanding
- **ML/DL Frameworks**: PyTorch (DDP), Hugging Face (Transformers/TRL/PEFT), Scikit-learn, FAISS, ONNX
- **Architectures**: VLMs, LLMs, ViTs, DINOv3, GRPO/RLVR, On-Policy Distillation, LoRA/QLoRA
- **Engineering & Tools**: Python, C++, Java, SQL, Linux, Docker, Git, CI/CD, Jenkins, Flask/Gunicorn, NGINX

Work Experience

University of California, Irvine — Irvine, CA
ML Research Engineer — 2022 – Present

- Led machine learning research for a **$2M+ NIH-funded** program pre-training, fine-tuning, and adapting **foundation models** for cancer-treatment prognosis; authored **24 peer-reviewed papers** (11 first/co-first) and 12 under review.
- Adapted frozen foundation models (**DINOv3, MedSAM**) to volumetric imaging via spatial patch-token pooling + lightweight heads, beating end-to-end fine-tuning and domain-specific 3D FMs at **AUROC up to 0.96**.
- Ran **continued self-supervised pretraining** of DINOv3 on **2.8M+** CT/MRI slices; a **frozen-teacher geometric anchor** beat stock DINOv3 on pancreas-MRI tasks with KiTS23 retention intact.
- Post-trained **vision-language foundation models** (2B–8B) with **LoRA/QLoRA SFT + GRPO/RLVR**; showed the reward keeps reasoning only where it pays: $\sim$18$\times$ shorter answers on perception VQA, **+10 points** on charts.
- Designed **multi-task volumetric deep learning** (U-Net++/EfficientNet) for joint **segmentation and outcome prediction**, outperforming **5 CNN/ViT backbones** with predictions validated against gold-standard histology.

Tencent QTrade — Shenzhen, China
NLP Engineer — 2020 – 2022

- Raised NER Macro-F1 by **11 points (77% $\rightarrow$ 88%)** by replacing LSTM with **ALBERT** plus domain-specific continued pretraining (**PyTorch DDP, 8$\times$V100**), re-annotating data, and tuning post-processing.
- Architected and scaled an NLP microservice (Flask/Gunicorn, NGINX) to **500k+ DAU** at **sub-100ms P95**; canary rollout (5% $\rightarrow$ 100%) with online/offline monitoring and Jenkins model rollback.
- Built the production data pipeline: streaming ingestion, rule-based preprocessing, domain corpus normalization (entity alias resolution), and modular post-processing for output boundary repair at **99%+ specificity**.
- Redesigned the system for extensibility, rewriting post-processing from scratch to support varied data patterns and eliminate rule conflicts; led a team of **4** (2 MLEs, 2 annotators) through the migration.

Projects

Foundation-Model Adaptation for Low-Data Imaging 2026

Tech Stack: Python, PyTorch, DINOv3, MedSAM, Scikit-learn, Optuna, Grouped CV

- Adapted frozen **DINOv3/MedSAM embeddings** with spatial patch-token pooling and lightweight heads; beat end-to-end fine-tuning (peak 0.63/0.70) and domain-specific 3D FMs, reaching AUROC **0.94–0.96** mouse / **0.78–0.86** human.
- Validated across **2 independent imaging domains** with leave-one-subject-out/patient-grouped CV and a zero-transfer control; human imaging score improved adjusted survival concordance from **0.611 to 0.726** (**HR 1.87/SD, p<0.001**).
- **Sub-second** inference per scan with lightweight heads; first-author paper **in press** at *eBioMedicine*.
- **AbdoDINO**: domain-adaptive continued DINOv3 pretraining on **2.85M** CT/MRI slices (16.7K volumes, 10 public datasets, contamination-audited); a **frozen-teacher geometric anchor** preserved patch structure while learning the new domain, **+0.025–0.049 AUROC** on pancreas-MRI over stock DINOv3 with retention held.

Post-Training Vision-Language Models with RLVR 2026

Tech Stack: PyTorch, Hugging Face TRL/PEFT, LoRA/QLoRA, GRPO, On-Policy Distillation, vLLM, Slurm

- Post-trained **Qwen3-VL (2B/8B)** with LoRA/QLoRA SFT, GRPO from verifiable rewards and on-policy distillation over vLLM rollouts; the correctness reward kept reasoning only where it paid: cut **148 → 8 tokens** on perception VQA, retained on charts to lift **ChartQA 69.7 → 80.1** (+5.4 audited), above every short-answer model.
- SFT followed by GRPO was the top model on **SLAKE** (76.2 at 2B, **78.2** at 8B) and on held-out **OmniMedVQA** collections (85.6); a black-image control showed RL's gains need the image, while SFT gained up to 15.4 points blind.
- Built a reproducible evaluation harness: 9 test sets, official scorers matching the authors' code on **626K predictions**, an image-audited LLM-as-judge (GRPO lifted reasoning-mode VQAv2 **+4.4**), and a pre-registered 14-dataset screen that rejected **DeepTumorVQA**, answerable largely without the image (text-only rule 60.7% vs fine-tuned 61.4%).

Multilingual NER Framework 2022

ACL 2022, SemEval-2022 Task 11 (**Top-3 in Multilingual**)

- Co-developed a **multilingual NER foundation model**: **XLM-RoBERTa-large** fine-tuned across **11 languages** with augmentation and adversarial training (PGD/FreeLB), plus track-specific heads (lexicon GCN, R-Drop shared/private heads).
- **7-model** tag-level majority-vote ensemble refit with **2 rounds of pseudo-labeling** (no new annotation): **Top-3 Multilingual** (Macro-F1 77.66), Top-4 Code-Mixed, Top-7 Chinese at SemEval-2022 Task 11.
- Every component entered through a dev-set ablation (augmentation + adversarial training **+2.3 F1** on code-mixed); another product team adopted the domain-adapted backbone and improved its intent-detection and NER F1.

Selected Publications (6 of 24 peer-reviewed; 12 under review)

- Yu, G. *et al.* **Foundation-Model** Imaging Biomarkers for Pancreatic Cancer: Cross-Species Discovery of Conserved Tumour Biology in Mouse MRI and Human CT. *eBioMedicine* (*The Lancet* Discovery Science), **in press**, 2026.
- Yu, G. *et al.* A Systematic Review of **Reward Design** and Evaluation in **Reinforcement-Learning Post-Training** of Medical Imaging Models. *npj Digit. Med.*, **under review**, 2026.
- Yu, G. *et al.* Qtrade AI at SemEval-2022 Task 11: An Unified Framework for **Multilingual NER** Task. *ACL*, 2022.
- Yu, G. *et al.* **Multi-Task Deep Learning** on MRI for Tumor Segmentation and Treatment Response Prediction in an Experimental Model of Hepatocellular Carcinoma. *Diagnostics*, 2025.
- Yu, G. *et al.* AI-Derived MRI Biomarkers Using **Vision Transformer** for Predicting Combination Immunotherapy Outcomes in Liver Cancer. *ISMRM*, 2025.
- Yu, G. *et al.* **Deep Learning** Prediction of MRI Data for Evaluating Combined Transcatheter NK Cell Local Delivery and Sorafenib Therapy in Hepatocellular Carcinoma. *JVIR*, 2025.
- Yu, G. *et al.* **MRI Radiomics** to Monitor Therapeutic Outcome of Sorafenib plus IHA Transcatheter NK Cell Combination Therapy in Hepatocellular Carcinoma. *J. Transl. Med.*, 2024.