

Zhen Yang

New Haven, CT | zhen.yang.zy372@yale.edu | Website | GitHub | LinkedIn

Summer 2027: Eligible to apply for CPT authorization through Yale

Education

Yale University

M.S. in Biostatistics, Data Science Pathway

Nanjing University of Posts and Telecommunications

B.Eng. in Data Science and Big Data Technology

New Haven, CT

Sep 2026 – May 2028 (Expected)

Nanjing, China

Sep 2022 – Jun 2026

Selected Research

Full publication list: personal website

Risk-Aware Structural Repair for Topology-Sensitive Image Segmentation *1st Author · WACV 2027 · Under Review*

- Proposed RASR for vessel/road segmentation, combining evidential uncertainty with structural-violation signals to construct risk-guided ROIs; froze the backbone and trained only a 20.5K-parameter shallow repair head that applies gated residual logits locally while strictly preserving predictions outside the ROI.
- Evaluated on DRIVE, CHASE DB1, STARE, DeepGlobe and Massachusetts Roads across UNet, Attention-UNet, UNet++ and Swin-UNet against cIDice, TopoPH, DMT, SRL and post-processing baselines, with strong overall overlap, boundary and topology performance plus cross-backbone transfer analysis.

Evidential U-KAN for Trustworthy Medical Image Segmentation *1st Author · Preprint + Patent*

- Designed Progressive Evidential Uncertainty-guided Attention (PEUA) with low-rank denoising to focus representation on ambiguous regions and Semantic-preserving Adaptive Evidence Learning (SAEL) to retain boundary semantics under uncertainty.
- Integrated PEUA/SAEL into U-KAN and evaluated across 4 medical-image datasets, improving segmentation accuracy and reliability; resulted in a patent on progressive uncertainty-guided trustworthy medical image segmentation.

Bi-level Meta-Policy Control for Dynamic Uncertainty Calibration in Evidential Deep Learning *1st Author · Preprint*

- Proposed a state-aware bi-level meta-policy that dynamically controls KL regularization and adaptive/class-specific Dirichlet priors, replacing fixed EDL hyperparameters with learned calibration decisions.
- Evaluated across classification, medical segmentation and OOD settings, reporting +12.1% accuracy and -24.7% MUE while studying the accuracy-calibration tradeoff under non-stationary uncertainty.

EviE: Evidential Exploration for Value-Based Reinforcement Learning *1st Author · NeurIPS 2026 · Under Review*

- Used evidential vacuity to allocate exploration adaptively by state, connecting uncertainty estimation to downstream decision-making with tabular convergence guarantees and validation across classic and pixel-based RL environments.

Industry Experience

Baidu

Algorithm Intern - Search Retrieval

Beijing, China

Jan 2026 – Jun 2026

- Owned post-training, data construction and evaluation for dual-encoder query-URL retrieval; diagnosed term omission/mismatch, semantic misinterpretation and relevant-but-not-answerable failures.
- Designed hard-negative mining, loss reweighting and staged optimization to strengthen retrieval boundaries, improving average Recall by 10.5% across multiple cutoffs and evaluation sets.
- Built four LLM-driven hard-case generation routes for answer-preserving paraphrase, condition perturbation, confusion-pair reconstruction and URL-to-query generation, adding +3.2% sequential average Recall.
- Added LLM+rule teacher scoring and listwise post-training/distillation for finer candidate supervision, delivering a further +6.0% sequential average Recall improvement.

HONOR

Algorithm Intern - Generative Engine Optimization

Nanjing, China

May 2025 – Dec 2025

- Independently built a closed-loop AI optimization system spanning social-demand mining, weekly keyword allocation, content generation, online publication handoff, automated posterior visibility measurement and weekly strategy iteration.
- Separated evaluation into source-level visibility and final product recommendation; automated target-prompt queries after publication so live connected-LLM behavior became the main feedback signal.
- Built an 8B LoRA+contrastive prior evaluator using 45 high-ranking SEO pages not selected by the LLM as negatives and 5 LLM-selected references as positives per query, enabling pre-publication scoring before live evaluation.
- Reached >70% product visibility on 14 core prompts with 67 initial articles across 13 online outlets; later scaled to 1,000+ articles while sustaining >65% visibility for six months.

Industrial and Commercial Bank of China (ICBC)

Algorithm Intern - Credit Risk Modeling

Changzhou, China

Jan 2024 – Feb 2024

- Cleaned, integrated and standardized multi-source customer records, resolved data-format inconsistencies and engineered features for an internal credit-risk model-comparison pipeline.
- Implemented KNN, Naive Bayes and MLP loan-risk baselines, with the best lightweight model reaching 72% accuracy.
- Developed an LSTM-Transformer over three months of behavioral data, reaching 0.89 AUC and identifying 85% of high-risk customers two weeks in advance as an experimental early-warning signal.

Technical

Python, PyTorch; statistical modeling, uncertainty quantification/calibration, medical image segmentation, machine learning, deep learning