

YUXUAN YIN

✉ y_yin@ucsb.edu | [in](#) LinkedIn | [🏠](#) Homepage | [G](#) Google Scholar | ☎ (+1) 805-637-1800

EDUCATION

University of California, Santa Barbara, USA

Doctor of Philosophy in Electrical and Computer Engineering

Sep. 2021 –

Advisor: Prof. Peng Li [[G](#)]

Tsinghua University, Beijing, China

Bachelor of Engineering in Electronic Engineering & Bachelor of Science in Mathematics

Aug. 2016 – Jun. 2020

INTERNSHIP EXPERIENCE

Software Engineering Intern, Google, Sunnyvale, California

Research Topic: Large language models for cloud jobs' resource and runtime prediction

Jun. 2025 – Sep. 2025

Mentor: Shengke Zhou [[in](#)]

AI Research Intern, NXP Automotive Processing, Austin, Texas

Research Topic: Conformal prediction for time series forecasting of PCB board aging

Jun. 2024 – Sep. 2024

Mentor: Chen He [[in](#)]

AI Research Intern, NXP Automotive Processing, Austin, Texas

Research Topic: Reliable interval prediction for chip performance

Jun. 2023 – Sep. 2023

Mentor: Chen He [[in](#)]

SUMMARY

- Ph.D. candidate in ML and agentic AI with hands-on experience building and evaluating LLM systems for real-world tasks, including self-evolving from feedback, long-context modeling, and uncertainty quantification.
- Strong foundation in reliable machine learning under data scarcity, with published work in semi-supervised learning, Bayesian optimization, and conformal prediction at venues such as ICML, AAAI, DAC, and ICCAD.

SELECTED PROJECTS

I. Agentic AI

1. LLMs for Cloud Jobs' Resource and Runtime Prediction

[Under Review]

- Built a LoRA-based supervised fine-tuning pipeline for predicting cloud-job resource and runtime requirements from job metadata, and improved numeric reliability with scientific-notation formatting and constrained decoding.

2. Rotary Structural Embedding for Long-Context Understanding

[Ongoing]

- Designed a training-free extension to RoPE that learns per-chunk rotary angle adjustments from inter-chunk correlations, improving long-context reasoning quality on RAG and multi-turn dialogue settings.

3. Analog Design Bayesian Optimization with In-Context Learning of Large Language Models

[ICCAD' 24]

- Developed a self-evolving LLM+BO agentic framework for analog circuit parameter optimization, combining LLM prior knowledge and in-context adaptation with the sample efficiency and rigor of Bayesian optimization.

4. Graph-based Uncertainty Quantification of LLMs

[Under Review]

- Proposed a graph-based alignment method for uncertainty estimation in long-form LLM outputs by representing generations as fact graphs and quantifying semantic and structural disagreement.

5. High-Dimensional Bayesian Optimization with SSL and Unlabeled Sampling

[ICML' 24]

- Developed TSBO, a semi-supervised Bayesian optimization framework for data-efficient sequential design that selects high-value unlabeled samples and generates calibrated pseudo-labels via teacher-student feedback, achieving SOTA with up to 364.2x fewer labeled samples.

II. Uncertainty Quantification

1. Graph-based Uncertainty Quantification of LLMs

[Under Review]

- Introduced a graph-based alignment framework for quantifying uncertainty in long-form LLM generations by modeling atomic facts and measuring semantic and structural inconsistencies.

2. Data-Efficient Interval Prediction for Chip Performance Under Distribution Shifts

[DAC' 24]

- Proposed Calibrated Weighted CQR (CWCQR), a distribution-free interval prediction method with coverage guarantees under shift; on a 16nm automotive chip task, reduced interval length by 60% at 90% target coverage versus sigma-based baselines.

3. Reliable Interval Prediction of Chip Performance via Conformalized Quantile Regression

[DATE' 24]

- Applied Conformalized Quantile Regression with a CatBoost predictor and on-chip monitors for reliable interval estimation in semiconductor manufacturing, achieving SOTA accuracy and data efficiency on a 5nm chip task.

PUBLICATIONS

[ITC' 25] **Yuxuan Yin**, Rebecca Chen, Boxun Xu, Chen He, Peng Li. Transfer Learning for Minimum Operating Voltage Prediction in Advanced Technology Nodes: Leveraging Legacy Data and Silicon Odometer Sensing [\[paper\]](#)

[VTS' 25] **Yuxuan Yin**, Rebecca Chen, Chen He, Peng Li. Reliable Board-Level Degradation Prediction with Monotonic Segmented Regression under Noisy Measurement [\[paper\]](#)

[VTS' 25] **Yuxuan Yin**, Rebecca Chen, Chen He, Peng Li. Data-Efficient Prediction of Minimum Operating Voltage via Inter- and Intra-Wafer Variation Alignment [\[paper\]](#)

[ISCA' 25] Boxun Xu, **Yuxuan Yin**, Vikram Iyer, Peng Li. Bishop: Sparsified Bundling Spiking Transformers on Heterogeneous Cores with Error-Constrained Pruning [\[paper\]](#)

[ICCAD' 24] **Yuxuan Yin***, Yu Wang*, Boxun Xu, Peng Li. ADO-LLM: Analog Design Bayesian Optimization with In-Context Learning of Large Language Models [\[paper\]](#)

[ICML' 24] **Yuxuan Yin**, Yu Wang, Peng Li. High-dimensional Bayesian Optimization via Semi-supervised Learning with Optimized Unlabeled Data Sampling (**Spotlight**) [\[paper\]](#)

[DAC' 24] **Yuxuan Yin**, Rebecca Chen, Chen He, Peng Li. Data-Efficient Conformalized Interval Prediction of Minimum Operating Voltage Capturing Process Variations [\[paper\]](#)

[DATE' 24] **Yuxuan Yin**, Xiaoxiao Wang, Rebecca Chen, Chen He, Peng Li. Reliable Interval Prediction of Minimum Operating Voltage Based on On-chip Monitors via Conformalized Quantile Regression [\[paper\]](#)

[AAAI' 24] Yu Wang, **Yuxuan Yin**, Karthik Suryanarayana, Jan Drgona, Malachi Schram, Mahantesh Halappanavar, Frank Liu, Peng Li. Semi-Supervised Learning of Dynamical Systems with Neural Ordinary Differential Equations: A Teacher-Student Model Approach [\[paper\]](#)

[ITC' 23] **Yuxuan Yin**, Rebecca Chen, Chen He, Peng Li. Domain-Specific Machine Learning based Minimum Operating Voltage Prediction using On-Chip Monitor Data [\[paper\]](#)

[IMWUT' 21] Habiba Farrukh, Tinghan Yang, Hanwen Xu, **Yuxuan Yin**, He Wang, Z Berkay Celik. S3: Side-Channel Attack on Stylus Pencil through Sensors [\[paper\]](#)

PREPRINTS

[Arxiv] **Yuxuan Yin**, Shengke Zhou, Yunjie Zhang, Ajay Mohindra, Boxun Xu, Peng Li. Large Language Models for EDA Cloud Job Resource and Lifetime Prediction [\[paper\]](#)

[Arxiv] **Yuxuan Yin***, Yu Wang*, Peng Li. Towards the Mitigation of Confirmation Bias in Semi-supervised Learning: a Debaised Training Perspective [\[paper\]](#)