

Statistical Foundations of AI

Conference Program

November 13–14, 2026 • UC Santa Barbara

Corwin Pavilion and University Center meeting rooms. All times are Pacific Standard Time.

Talks appear in presentation order. Each room's two-hour talk session is divided equally among its speakers. Panel discussions are highlighted.

Time and room	Friday November 13
Parallel Session 1	
1C Corwin 8:30–10:30 am Statistical computation and structured modeling Chair: Grace Yi	Hani Doss (University of Florida) Theoretical and Empirical Evaluation of a Grouped Gibbs Sampler for Parallel Computation in the Latent Dirichlet Allocation Model Hao Helen Zhang (University of Arizona) Dynamic Supervised Principal Component Analysis for High-Dimensional Classification Yazhen Wang (University of Wisconsin-Madison) Learning Relationship between Quantum Walks and Underdamped Langevin Dynamics Ying Nian Wu (UCLA) Quantum Field Theory as Generative Models, Emergent Gravity as Latent Variable Model Hongzhe Li (University of Pennsylvania) Generation-Powered Inference for Distribution-Valued Outcomes Bertrand Clarke (U. Nebraska-Lincoln) Uncertainty Quantification via the Posterior Predictive Variance
1F Flying-A 8:30–10:30 am AI agents and adaptive evaluation Chair: Jin Zhou	Jingshen Wang (UC Berkeley) Agentifying AI Evaluation: What to Evaluate, When to Stop, and How to Estimate Linjun Zhang (Rutgers University) Statistical Evaluation of LLMs David Talby (John Snow Labs) Statistical design for medical AI safety benchmarks: precision targets, joint balance, and perturbation validity Jason Klusowski (Princeton University) Statistical Attribute Alignment for Black-Box Generative AI Jie Ding (University of Minnesota) The Future of AI Scientists: Emerging Directions and Fundamental Challenges Ruoqing Zhu (University of Illinois at Urbana-Champaign) Building AI Agents for Better Human-Agent Interaction
1H Harbor 8:30–10:30 am Statistical learning for biomedical data Chair: Yang Feng	Jingjing Zou (UC San Diego) Structured Self-Supervised Learning for Heterogeneous Wearable Data Andrew Holbrook (UCLA Biostatistics) Decoding Aging Neurons Michael Daniels (Department of Statistics, University of Florida) A case study of causal mediation using Bayesian nonparametrics and semiparametric corrections Sam Behseta (Kaiser Permanente Bernard J. Tyson School of Medicine) The Specification Gap: What Artificial Intelligence Leaves Unresolved in Statistical Inference Lan Xue (Oregon State University) Cluster Analysis of Longitudinal Associations Between Microbial Abundance and Environmental Factors Sandra Safo (University of Minnesota) Scalable Randomized Kernel Methods for Multiview Data Integration and Prediction
1S State Street 8:30–10:30 am Generative models and learning dynamics Chair: Jacob Bien	Ning Ning (Texas A&M University) Beyond Prediction: World Models for Understanding, Planning, and Action Yuchen Zhou (University of Illinois Urbana-Champaign) Optimal Convergence Analysis of DDPM for General Distributions Shihao Yang (Georgia Institute of Technology) Statistical Interpretation of AI Architecture and AI Agents for Statistical Research Zhonghua Liu (Columbia University) Gradient-Flow Optimization as Dynamic Random-Effects Inference with Applications to Deep Learning Mingzhao Hu (Mayo Clinic) ComposeCFM: Flow-Matching Synthetic Generation and Harmonization Weining Shen (UC Irvine) SportD: How do VLMs physically strategize?
10:30–11:00 am	Break

Time and room	Friday November 13
11:00–11:05 am Corwin	OPENING REMARKS Shelly Gable Dean of Mathematical, Life, and Physical Sciences at UCSB
11:05 am–12:05 pm Corwin	PLENARY TALK Susan Murphy (Harvard University) Reinforcement Learning for Digital Health Interventions in the Dyadic Setting Chair: Annie Qu
12:05–1:00 pm	Lunch
Parallel Session 2	
2C Corwin 1:00–3:00 pm Trustworthy AI and structured inference Chair: Jin Zhou	Weijie Su (University of Pennsylvania) Alignment in Large Language Models: Statistical and Game-Theoretic Perspectives Aaditya Ramdas (Stanford University) Controlling the Probability of Unknown Unknowns in Pre-Deployment Testing of AI Systems Padhraic Smyth (UC Irvine) Prediction, Confidence, and Calibration in the Age of Deep Learning Runze Li (Penn State University) Federated Feature Selection with False Discovery Rate Control Ji Zhu (University of Michigan) Bias-Corrected Inference for Edge-Covariate Effects in Latent Space Network Models
2F Flying-A 1:00–3:00 pm Regression and predictive inference Chair: Yazhen Wang	Oscar Hernan Madrid Padilla (UCLA) Risk Bounds For Distributional Regression Grace Yi (University of Western Ontario) Residual Correction and Representation Learning for Robust Regression with Continuous Noisy Labels Boxiang Wang (University of Iowa) ABC: Certified Selective Evaluation with a Panel of LLM Judges Matteo Sesia (University of Southern California) Noise-Adaptive Conformal Classification with Marginal Coverage Hao Ho/Ker-Chau Li (UCLA) ATCG: Learning Structural Traits for Auditable Deep Classification Min Xu (Rutgers University) Estimation of the sub-Gaussian Parameter
2H Harbor 1:00–3:00 pm Statistical learning and assisted inference Chair: Ying Nian Wu	Kun Chen (University of Connecticut) Tree-Guided Feature Aggregation as Structured Representation Learning Emma Jingfei Zhang (Emory University) Beyond Point Predictions: Distribution-Informed Prediction-Powered Inference Yang Feng (New York University) Federated Transfer Learning with Differential Privacy Kai Zhang (UNC Chapel Hill) BELIEF in Dependence Qiyang Han (Rutgers University) Algorithmic inference via nonconvex gradient descent Arun Kuchibhotla (Carnegie Mellon University) Mean-Median gaps in real line and Inference in Infinite-dimensions
2S State Street 1:00–3:00 pm Network models, representations, and experiments Chair: Bertrand Clarke	Xiwei Tang (University of Texas at Dallas) HODGEDEPTH: SPATIAL REPRESENTATION LEARNING THROUGH GLOBAL GEOMETRY Keith Levin (University of Wisconsin—Madison) Testing for correlation between network structure and high-dimensional node covariates Jacob Bien (University of Southern California) Hierarchical Clustering With Confidence Xiaoyue Maggie Niu (Pennsylvania State University) A Dyadic Latent Space Model for Directed Networks Jesús Arroyo (Texas A&M University) Identifiability and Spectral Estimation in Multilayer Degree-Corrected Mixed Membership Models Alexander Volfovsky (Duke University) How to Experiment When You Don't Yet Know What Matters
3:00–3:30 pm	Break

Time and room	Friday November 13
Parallel Session 3	
3C Corwin 3:30–5:30 pm PANEL DISCUSSION Industry perspectives on trustworthy AI Moderator: Hao Helen Zhang	PANEL DISCUSSION Industry perspectives on trustworthy AI Moderator: Hao Helen Zhang (University of Arizona) David Talby (John Snow Labs) Sam Behseta (Kaiser Permanente Bernard J. Tyson School of Medicine) Guannan Zhai (Amazon) Yujia Deng (Zoom Communications Inc.) Haoda Fu (Amgen) Louis Ehwerhemuepha (CHOC)
3F Flying-A 3:30–5:30 pm Online learning and decision making Chair: Lan Xue	Xuan Bi (University of Minnesota) Optimal Auditing of Sequential Decision Systems under Dynamical Constraints Chengchun Shi (London School of Economics and Political Science) Demystifying Group Relative Policy Optimization: Its Policy Gradient is a U-Statistic Yuhua Zhu (UCLA) Reinforcement Learning in the Physical World Stefan Wager (Stanford University) Estimating Dynamic Marginal Policy Effects under Sequential Unconfoundedness Yufeng Liu (University of Michigan) Structured Learning of Individualized Decision Rules with Complex Options Arash Amini (UCLA) Two Sample Sizes in Test-Time LLM Ensembling: Buffered ERM and Minimax Theory
3H Harbor 3:30–5:30 pm Statistical inference after learning Chair: Hongzhe Li	Lili Zheng (University of Illinois Urbana - Champaign) LOCO-AdaMP: Built-in LOCO Inference for Adaptive Minipatch Ensembles with Enhanced Prediction Armeen Taeb (University of Washington) Consensus Tree Estimation with False Discovery Control via Partially Ordered Sets Dan Kessler (UNC Chapel Hill) Inference After Exploration: A Grammar for Adaptive Data Science Tate Jacobson (Oregon State University) Global Inference for High-Dimensional Quantile Regression with Heterogeneity Qi Xu (University of Minnesota) Towards Efficient Inference under Nonmonotone Missingness with General Imputation Joshua Cape (University of Wisconsin-Madison) Network Signflip Parallel Analysis for Selecting the Embedding Dimension
3S State Street 3:30–5:30 pm Causal discovery and treatment effects Chair: Hani Doss	Yuqi Gu (Columbia University) Discrete Causal Representation Learning Carlos Cinelli (University of Washington) Challenges in Causality: Sensitivity Analysis and Automation Matteo Bonvini (Rutgers University) On regression with estimated covariates and conditional effects given the propensity score Yubai Yuan (The Pennsylvania State University) Causal inference under interference with incomplete exposure mappings Didong Li (UNC Chapel Hill) When Can We Trust Synthetic Data for Statistical Inference? Ying Zhou (University of Connecticut) Learning from the Unseen: Offline Reinforcement Learning with Hidden Actions
5:30–6:30 pm Corwin	Poster session All posters share this session. See the poster directory.
7:00–8:30 pm Corwin	BANQUET TALK Sudipto Banerjee (UCLA) Carbs, Calories, and Computation: Serving Up Statistical Inference in the AI Ecosystem.

Time and room	Saturday November 14
Parallel Session 4	
4C Corwin 8:30–10:30 am Biomedical modeling and AI-assisted inference Chair: Hua Zhou	Lily Wang (George Mason University) Synthetic Surrogate Functional Regression for Neuroimaging Analysis Lexin Li (UC Berkeley) Integrating Statistics and AI for Neuroscience Research Ying Guo (Emory University) Task-Aware Statistical AI for Learning Directed Brain Connectomes from Neuroimaging Data Fengzhu Sun (University of Southern California) Statistical and AI Approaches for Investigating Viruses and Virus-host Interactions Gang Li (Department of Biostatistics, UCLA) SAPA: Safe and Adaptive Prediction Augmentation for Survival Analysis Xinping Cui (UC Riverside) Scalable Joint Modeling of Multiple Longitudinal Biomarkers and Competing Risks Time-to-Event Data: Applications to Mega-Scale Health Research
4F Flying-A 8:30–10:30 am Evaluation and inference with AI predictions Chair: Veronica Berrocal	Carey Priebe (Johns Hopkins University) Performance guarantees for LLMs Junwei Lu (Harvard University) Language Model Preference Inference via Fisher Random Walk Will Wei Sun (Purdue University) Low Rank for Rank: Uncertainty-Aware Task-Specific LLM Ranking from Sparse Pairwise Comparisons Jiwei Zhao (University of Wisconsin-Madison) SADA: Safe and Adaptive Aggregation of Multiple Black-Box Predictions in Semi-Supervised Learning Peng Wang (University of Cincinnati) Building Statistical AI Tools for Comparative Analysis Hongquan Xu (UCLA) Prompt engineering via order-of-addition experiments and screening
4H Harbor 8:30–10:30 am Structured representations and clustering Chair: Ali Shojaie	Shujie Ma (UC Riverside) Learning What to Share: Adaptive Spectral Integration and Clustering of Multimodal Data Mladen Kolar (University of Southern California & MBZUAI) Borrowing Strength Across Studies: Structure-Aware Transfer Learning for High-Dimensional Models Gongjun Xu (University of Michigan) Beyond Vintage Rotation: Bias-Free Sparse Representation Learning with Oracle Inference Jason Xu (UCLA) Majorization-Minimization Algorithms in Clustering Xiaohui Chen (University of Southern California) GEM-KMeans: Memory-Efficient and Accurate Clustering on Massive Scale with GPU Optimization Marianthi Markatou (University at Buffalo) A study on kernels for discrete random variables
4S State Street 8:30–10:30 am Reinforcement learning and adaptive decisions Chair: Ming Xie	Feifang Hu (George Washington University) AI-Powered Adaptive Experiments: From Sequential Analysis to Intelligent Clinical Trials. Lan Wang (University of Miami) Quantile Temporal Difference Algorithm for Distributional Reinforcement Learning Linglong Kong (University of Alberta) Principled Fast and Meta Knowledge Learners for Continual Reinforcement Learning Babak Shahbaba (UCI) Uncertainty-Guided Control: Statistical Foundations for Brain-Inspired Decision-Making Dehan Kong (University of Toronto) Successive classification learning for estimating quantile optimal treatment regimes Walter Dempsey (University of Michigan) Future-Path Incremental Excursion Effects in Sequentially Randomized Trials
10:30–11:00 am	Break
11:00 am–12:00 pm Corwin	PLENARY TALK Xihong Lin (Harvard University) From Generative AI to Statistical Inference: Harnessing Synthetic Data for Validity and Power Chair: Xiaowu Dai
12:00–1:00 pm	Lunch

Time and room	Saturday November 14
Parallel Session 5	
5C Corwin 1:00–3:00 pm PANEL DISCUSSION AI influence and the next generation of research Moderator: Regina Liu	PANEL DISCUSSION AI influence and the next generation of research Moderator: Regina Liu (Rutgers University) George Michailidis (UCLA) Bertrand Clarke (U. Nebraska-Lincoln) Heping Zhang (Yale University) Peter X. K. Song (University of Michigan) Radu Craiu (University of Toronto) Michael Daniels (Department of Statistics, University of Florida) Xiao Wang (Purdue University) Ying Wei (Columbia University) Michele Guindani (UCLA)
5F Flying-A 1:00–3:00 pm Statistical learning and computational inference Chair: Maranthi Markatou	Minge Xie (Rutgers University) Model-Free Inference for High-Dimensional Binary Classification Using Repro Samples Faming Liang (Purdue University) An Introduction to Stochastic Deep Learning Weixin Yao (UC Riverside) Online Kernel-Based Mode Learning Ryan Giordano (UC Berkeley) Black Box Variational Inference with a Deterministic Objective: Faster, More Accurate, and Even More Black Box Debdeep Pati (University of Wisconsin-Madison/Boston University) Probabilistic Symbolic Regression for Equation Discovery via Operator-induced and Regularized Symbolic Forests Ali Shojaie (University of Washington) GLM Inference with AI-Generated Synthetic Data
5H Harbor 1:00–3:00 pm Multi-view learning and data integration Chair: Hua Zhou	Rong Ma (Harvard University) Finding Shared Signals Across Many Datasets: Statistical Insights from RMT Cong Ma (University of Chicago) Bias-Corrected Subspace Intersection: Minimax-Optimal Shared Subspace Estimation in Multi-View Data Tiffany Tang (University of Notre Dame) Consensus dimension reduction via multi-view learning Ziyi Liang (UC Santa Barbara) Meta Fusion: A Unified Framework For Multimodal Fusion with Adaptive Mutual Learning Xinzhou Guo (HKUST) Stein-Encoder: A White-Box Supervised Encoder via Stein Identities in Multi-Modal Studies Tian Gu (Columbia University) Whom to trust and what to keep: adaptive Bayesian transfer learning from pretrained models
5S State Street 1:00–3:00 pm Statistical learning with heterogeneous data Chair: Hongtu Zhu	Zhezhen Jin (Department of Biostatistics, Columbia University) Subsampling for the analysis of large data Gourab Mukherjee (University of Southern California) Kernelized Stein Discrepancy for Learning from Heterogeneous Data Peijun Sang (University of Waterloo) Model positive and unlabeled data with a generalized additive density ratio model Bei Jiang (University of Alberta) fDA: A data-augmented predictive framework for generating fair synthetic labels with controlled fairness-faithfulness trade-offs Lu Tang (University of Pittsburgh) Fairness within Clinical Context: Stratified Tilted Risk Minimization for Medical Prediction Veronica Berrocal (UC Irvine) A Bayesian scalar-on-function kernel regression model for crop yield
3:00–3:30 pm	Break

Time and room	Saturday November 14
Parallel Session 6	
6C Corwin 3:30–5:30 pm Reliable learning and structured inference Chair: Shujie Ma	Hongtu Zhu (UNC Chapel Hill) MLFM: Manifold-Based Latent Factor Discovery in Functional Brain Networks Jian Kang (University of Michigan) BROWNIE: Bayesian Random Weight Neural Networks as a Simple Alternative to Deep Learning George Michailidis (UCLA) Beyond Flat Minima: Sharpness-Aware Bilevel and Minimax Optimization. Jinchi Lv (University of Southern California) LIDS: LLM Summary Inference Under the Layered Lens Tracy Ke (Harvard University) The Gold Rush in AI4Math: Structural Signatures of AI Involvement Qing Zhou (Department of Statistics & Data Science, UCLA) Causal Discovery on Dependent Data
6F Flying-A 3:30–5:30 pm Foundations of AI and distributed learning Chair: Faming Liang	Xiaoqian Liu (UC Riverside) An Interpretable and Scalable Framework for Evaluating Large Language Models Nabaran Deb (University of Chicago Booth School of Business) Sample complexity bounds for categorical Markov random fields via Discrete Diffusions Wenqing He (University of Western Ontario) Beyond Marginal Coverage: Efficient Localized Conformal Prediction via Residual Rank Calibration Rui Miao (University of Texas at Dallas) TOPPO: Rethinking PPO for Multi-Task Reinforcement Learning with Critic Balancing Jeffrey Zhang (UC Riverside) Enhancing Federated Learning Methods using Third Order Optimization Fei Xue (Purdue University) Federated Learning for Feature Mismatch
6H Harbor 3:30–5:30 pm Learning from distributions and complex measurements Chair: Lily Wang	Linbo Wang (University of Toronto) Estimating Wasserstein Barycenters from Sparse Deep Sequencing Irina Gaynanova (University of Michigan) Learning from multivariate distributions with nonparanormal transport Hyunkeun Ryan Cho (UCSD) Scalable four-level nested functional principal component analysis for large dense functional data Soham Jana (University of Notre Dame) Modeling with speckle noise in high-dimensional image sensing Rui Duan (Harvard T.H. Chan School of Public Health) Unsupervised Aggregation of Multiple Learning Algorithms Wen Zhou (New York University) A Latent Semiparametric Mixture Model for Robust Heterogeneity Learning in High Dimensions
6S State Street 3:30–5:30 pm Learning from sequential and process data Chair: Radu Craiu	Xiucan Ding (UC Davis) Transformers for Time Series: Statistical Optimality, Structural Exploration and Gradient Descent Consistency Robert Lund (UC Santa Cruz) Multiple Changepoint Detection in Non-Gaussian Time Series Hernando Ombao (Statistics, Bioengineering, and Applied Mathematics, KAUST) Dependence in Networks: Extremal Vulnerability and Granger Causality Beyond Linear Models Chunming Zhang (University of Wisconsin-Madison) Finite-Step Training of Deep ReLU Networks Zhiliang Ying (Columbia University) Process data and educational measurement Yan Yu (University of Cincinnati) (Mis)Understanding Benign Overfitting in Equity Return Prediction

Poster Session

Friday November 13 • 5:30–6:30 pm • Corwin

All posters are presented in this single session. Listings are ordered by the first presenter's surname.

- P01 Susmit Bhattacharyya** (Colorado School of Mines) Marginal Distribution of the LatticeKrig coefficients under Spatial Autoregressive model with Spatial Extremes
- P02 Stevenson Bolivar** (California Polytechnic State University (Cal Poly)) Time-Varying Factor Models for Dynamic High-Dimensional Data
- P03 Keivan Bolouri** (UC Irvine) Aggregation Frequency and Active-Set Agreement in Federated Lasso
- P04 Baiting Chen** (UCLA) Common-agency games for multi-objective test-time alignment
- P05 Xingran Chen** (University of Michigan) Title to be confirmed *Topic: inference with machine learning imputations under general missingness.*
- P06 Yifan Chen** (UC Santa Barbara) Statistical Guarantees for Preference Learning with Convex Neural Networks
- P07 Yassine Chouikh and Shannon Li** (UC San Francisco) Verify-or-Abstain Evidence Pipelines: When Cryptographic Custody Is Necessary but Not Sufficient
- P08 Chengyu Cui** (Univ of Michigan - Ann Arbor, MI) Convexity in Disguise: A Theoretical Framework for Nonconvex Low-Rank Matrix Estimation
- P09 Dongyu Cui** (University of Minnesota) Observation-Level Watermarking and Detection for Tabular Data
- P10 Ying Cui** (Stanford University) Learning Interpretable Point-Based Clinical Risk Scores via Direct Optimization
- P11 Annan Deng** (New York University) Geometry-Aware Representation Learning for Functional Connectivity Matrices
- P12 Prasanjit Dubey** (Georgia Institute of Technology) Matching Rates and Optimal Allocation for Federated Probe-Logit Distillation under Heterogeneous Bandwidth Budgets
- P13 Zijun Gao** (USC Marshall) Adaptive sample splitting for randomization tests
- P14 Daeyoung Ham** (University of Texas at San Antonio) Black-Box Knowledge Transfer across Distinct Feature Sets
- P15 Junchao Han** (UC Santa Barbara) Off-Policy Evaluation for Linear Episodic Heterogeneous MDPs by Hierarchical Bellman Residual Minimization
- P16 Rohan Hore** (Stanford University) Conformal Changepoint Localization
- P17 Pingbo Hu** (Columbia University Irving Medical Center) Deep Learning for Continuous-Time Dynamic Survival Prediction With Time-Varying Covariates
- P18 Yichen Hu** (UC Davis) Adaptive Sparsity-Aware Projection-Assisted Spectral Embedding for High-Dimensional Noisy Data with Nonlinear Structures
- P19 Yue Huang** (University of Southern California) iLV-HD: Reconstructing microbial dynamics from high-dimensional compositional time series
- P20 Hyeong Jin Hyun** (University of Cincinnati) Transport-induced Approximate Bayesian Computation
- P21 Ali Kaazempur-Mofrad** (UCLA) Beyond a Single Answer: Proper Scoring of LLM Uncertainty
- P22 Changwoo Lee** (University of Florida) Transformers Can Learn Posterior Predictive Distributions In-Context
- P23 Julianna Lee** (Westat) Using Large Language Models to extract productivity measures from CVs for Program Evaluation
- P24 Chunyin Lei** (UC Santa Barbara) Risk-Aware Quantile Learning for Personalized Dynamic Treatment Regimes
- P25 Dongyang Li** (Department of Mathematics, UC Santa Barbara.) t-Muon: Tensorizing Muon with the t-SVD for Language Model Pretraining
- P26 Pengtao Li** (University of Southern California) Smoothed estimation of Wasserstein barycenters
- P27 Xiaoke Li** (Department of Statistics and Applied Probability, UC Santa Barbara) History-Aware Conformal Prediction for Irregularly Observed Longitudinal Data
- P28 Ling Liang** (The University of Tennessee, Knoxville) A Provably Convergent and Practical Algorithm for Gromov-Wasserstein Optimal Transport
- P29 Qifeng Lin** (UC Davis) Learning common structures from high-dimensional noisy biviues using cross-integral operator
- P30 Yangyang Liu** (Department of Quantitative and Computational Biology, University of Southern California, Los Angeles, CA, USA) MACONi: a marker- and composition-aware framework for meta-Hi-C contact normalization
- P31 Jiecheng Lu** (Georgia Institute of Technology) Rethinking Sequence Models: Autoregressive Alignment and Expressive Operator Design
- P32 Muzhi Ma** (UCLA) Mechanism Design for Heterogeneous LLM Advertising

- P33 Ziling Ma** (King Abdullah University of Science and Technology (KAUST)) Adaptive Multi-Scale Forecasting and Gate-Localized Conformal Prediction for Multivariate Nonstationary Time Series
- P34 Elizabeth W. Miller** (University of Virginia) Diagnostics for Individual-Level Prediction Instability in Clinical Prediction Models
- P35 Simon Doan Nguyen** (University of Washington, Seattle) REALITrees: Rashomon Ensemble Active Learning for Interpretable Trees
- P36 Ronan Perry** (University of Washington) Inference after selection on randomized p-values
- P37 Samya Praharaj** (Rutgers University) Bandit Simulation for Average Reward Inference
- P38 Chen Qian** (UC Davis) Structural Classification of Locally Stationary Time Series Based on Second-order Characteristics
- P39 Anik Roy** (Department of Biostatistics and Bioinformatics, Emory University) Regression Analysis of Amplitude and Phase Components in Functional Data-Based Biomarkers
- P40 Xinrui Ruan** (UC Berkeley) Learning When to Seek Human Guidance for Best LLM Judge Identification
- P41 Matias Shedden** (University of Florida) Machine-learning-based subgroup identification and adaptive enrichment in a seamless Phase II/III trial
- P42 Zhaiming Shen** (Utah State University) Understanding In-Context Learning with Kernel Methods
- P43 Zhongjie Shi** (Georgia Institute of Technology) In-Context Learning Across Tasks with Transformers: Approximation and Generalization under Low-Dimensional Structure
- P44 Shreyash Sinha** (Department of Quantitative and Computational Biology, University of Southern California) Comparative exploration of gene-based and k-mer based features for predicting antimicrobial resistance
- P45 Yang Sui** (UCLA) Statistical Inference for LLM Evaluations: From Multiple Auxiliary Predictors to Valid Uncertainty Quantification
- P46 Ao Sun** (USC Marshall) LLM-Powered Predictive Inference with Online Text Time Series
- P47 Pablo Suárez-Serrato** (UC Santa Barbara) EuLearn: a 3D database for learning Euler characteristics
- P48 Ye Tian** (Penn State University) Semi-Supervised Representation Learning with Shared Factor Structure
- P49 Yidan Tian** (Michigan State University) Minimax-Optimal Transformer Classification for Functional Data with a Dense-Sparse Phase Transition
- P50 Kaylee Vo** (Harvard University) Variational Boosting for Physics-Informed Neural Networks
- P51 Patrick Vossler** (UC San Francisco) LLMs Judging LLMs: A Simplex Perspective
- P52 Anela Leyan Wang** (UCLA, Department of Biostatistics) Nonlinear subspace factor analysis via latent bottlenecks for integrating data across studies
- P53 Kevin Wang** (Univ of North Carolina At Chapel Hill) Can Generative Artificial Intelligence Survive Data Contamination? Theoretical Guarantees under Contaminated Recursive Training
- P54 Kexin Wang** (University of Southern California) DeepLINK-T: deep learning inference for time series data using knockoffs and LSTM with applications to metagenomic data
- P55 Yizheng Wang** (Hong Kong University of Science and Technology) Unsupervised Federated Multi-task Learning for Heterogeneous Tasks
- P56 Joong-Ho Won** (Seoul National University) Pareto Variational Autoencoder
- P57 Runxiong Wu** (University of Wisconsin-Madison) VarScope-MVAE: Learning Physically Scoped Variations in Multimodal Data
- P58 Xiliu Wu** (Indiana University Bloomington) Evaluating Tabular Foundation Models for Risk Prediction: A Benchmark against Traditional Machine Learning
- P59 Qingyan Xiang** (Department of Biostatistics, Vanderbilt University Medical Center) De-biased Machine Learning for Count Data: a Partially Linear Poisson Model Based on Neural Networks
- P60 David Xie** (Department of Statistics and Applied Probability, UC Santa Barbara) Analysis of Psychological Stress Response to Physiological Data from Wearable Devices: A Study on College Student Cohort.
- P61 Tianmin Xie** (University of Southern California) Conformal Inference for Open-Set and Imbalanced Classification
- P62 Zhaoyu Xing** (University of Notre Dame) SORHA: Semi-parametric Ordered Random Hypernetwork with Attributes
- P63 Kathy Xu** (UCSB PSTAT Department) Latent Restorative Patterns Linking Physical Activity to Acute Cardiovascular Protection in Wearable Health Data
- P64 Minxing Xu** (Northwestern University) Conflict in Medical Multi-Agent Reasoning: From Disagreement to Reliable Communication
- P65 Wenlu Xu** (UCLA) Physics-Informed Kernel Learning

- P66 Haoran Yan** (Georgia Institute of Technology) Neural Scaling Laws of Transformer Operator Network
- P67 Mingsong Yan** (UC Santa Barbara) In-Context Learning for Gaussian Kernel Ridge Regression with Transformers
- P68 Tiantian Yang** (University of Idaho) Directed Weighted Graph Neural Networks with Tensor Fusion for Multi-Omics Cancer Subtype Classification
- P69 Yahong Yang** (Binghamton University) Score Estimation and Approximation of the Diffusion Model
- P70 Duoduo Danny Ying** (Department of Statistics and Data Science, UCLA) Prompt engineering using order-of-addition experiments: An application to generating two-level fractional factorial designs
- P71 Miaoqing Yu** (UC Santa Barbara) BART for multi-outcome CATE estimation with Unmeasured Confounding
- P72 Keming Zhang** (Brown University) Interpretable and Uncertainty-Calibrated Bayesian Machine Learning for Survival Outcomes with Evolving Patient Histories
- P73 Niuniu Zhang** (UCLA Anderson School of Management) Joint Calibration of LLM Agents
- P74 Sherry Zhang** (University of Texas at Austin) Code availability is not enough: schemas for documenting and extending data analysis
- P75 Yijiao Zhang** (University of Pennsylvania) Causal Network Discovery from Interventional Count Data with Latent Linear DAGs
- P76 Yikun Zhang** (University of Chicago) Emputation: Identification-Guided Neural Imputation Framework
- P77 Yuexia Zhang** (The University of Texas at San Antonio) Extracting Causal Relational Rules for Medical Question-Answering Tasks using Large Language Models
- P78 Shangcheng Zhao** (Psychology & Brain Sciences, UCSB) Moral Shadowbanning: Is Conversational AI Less Helpful, Honest, or Harmless to Users Deemed to Have Poor Moral Character?
- P79 Jialu Zhou** (University of Maryland, College Park) Source-Fair Transport Calibration for Heterogeneous AI Evaluation
- P80 Bicheng Zhu** (Department of Statistics, UC Davis) Adaptive Clustering for Heterogeneous Multiplex Networks
- P81 Jungang Zou** (Department of Biostatistics at Columbia University) AI4BayesCode: From Natural Language Descriptions to Validated Modular Stateful Bayesian Samplers

Conference information: statfoundations-ai.github.io