

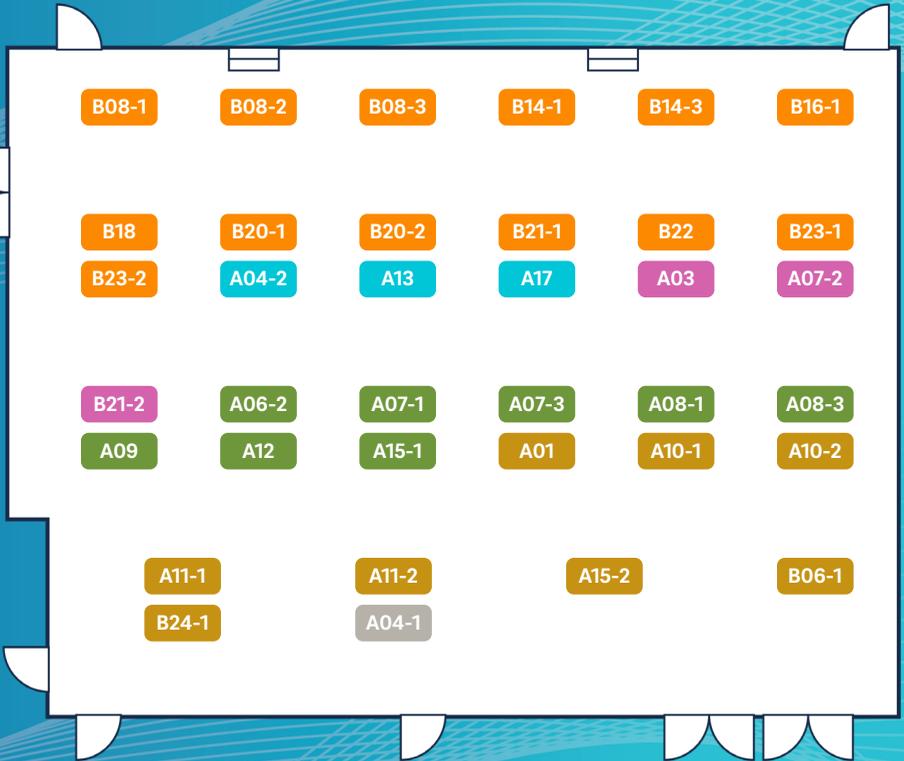
2026 CAPSTONE SHOWCASE

BLOCK 1 | 11:00 AM - 12:00 PM

Poster session for the undergraduate data science major's senior capstone program

Friday, March 13, 2026 | 10AM-3:00PM

Price Center East Ballroom,
UC San Diego Campus



AI and Health

- B08-1** Multi-Facility Patient Flow Optimization: A Simulation Approach
- B08-2** Large Language Models as Information Extractors in Healthcare
- B08-3** Large Language Models for Trauma Care Quality
- B14-1** The Influence of IQ on the Brain's Auditory Cortex Activation in Schizophrenia
- B14-3** Pre-Screening Motor-Imagery BCI Literacy Using Low-Sample EEG-Derived Features
- B16-1** Benchmarking Graph Algorithms for the Rediscovery of Mechanistic Drug-Repurposing Pathways
- B18** Absolute vs. Relative Microbial Abundances: Predictive Performance on EHR Data and Transformability
- B20-1** Multimorbidity-Aware Early Sepsis Prediction Using Patient-Specific Deterioration Trajectories
- B20-2** Multimorbidity-Aware Early Sepsis Prediction Using Cluster-Specific Deterioration Patterns
- B21-1** Classification of Pulmonary Edema: Combining Physician Assessment and Machine Learning Inference, Integrating Neural Networks with Large Language Models
- B22** Graph-Based Phenotyping of Morphine Dose Effects and Recovery in Mice
- B23-1** Neuro-Symbolic Framework for predicting Off-Target effects using Gene and Phenotype Ontologies
- B23-2** Neuro-Symbolic Off-Target Prediction for Pralsetinib Using Ontology-Grounded Knowledge Graphs

Algorithmic Foundations

- A04-2** Causal Calibration: Precision Tuning Causal Analysis Algorithms for Domain Performance
- A13** Communication lower bounds for multi-party graph traversal
- A17** Simulation Coding Exercises for Teaching Probability Theory

Explainability and Fairness

- A03** Opening the Bottleneck: Steering LLMs via Concept Intervention
- A07-2** Judging-the-Judges
- B21-2** Explaining Black box predictions in Medical AI

LLMs & Foundation Models

- A06-2** Research Agent Mobile: A Mobile-First Agentic Interface for Monitoring and Steering Machine Learning Experiments
- A07-1** Iterative Refinement of LLM as judge Framework
- A07-3** ContextEval: Evaluating LLM Agent Context Policies for ML Experiment Design
- A08-1** Think, Search, Correct: A Self-Reflection Mechanism for Adaptive Tool-Use Agents
- A08-3** Benchmarking LLM-Generated Biological Hypotheses for Scientific Discovery
- A09** Human-Centered Algorithmic Bias Auditing of Large Language Models
- A12** Uncertainty Quantification in Large Language Models
- A15-1** Motif-preserving Graph Tokenization for Biological Structure Generation using Sequence Transformer

ML Theory

- A01** Curating Fine-Tuning Data with KAIROS-Based Valuation
- A10-1** TeleSpeak: Modeling Early Language Acquisition with Reinforcement Learning
- A10-2** Sleep-Consolidated Learning & Plasticity Decay
- A11-1** What Enables In-Context Learning in Transformers? An Analysis of Attention Mechanisms
- A11-2** Investigation into In-Context Learning Capabilities of Transformers
- A15-2** Hypergraph Learning for General 3D Point Cloud Alignment
- B06-1** "Turbocharging Data Transformation: A Preconditioning Approach to Optimal Transport"
- B24-1** Generative AI for Good: Detecting and Mitigating Misinformation

NLP & IR

- A04-1** Market Foundry: From documents to intelligence — built your way.

2026 CAPSTONE SHOWCASE

BLOCK 2 | 1:00 - 3:00 PM

Poster session for the undergraduate data science major's senior capstone program

Friday, March 13, 2026 | 10AM-3:00PM

Price Center East Ballroom,
UC San Diego Campus



Applied AI

- A06-1** Toward Accelerated LLM Inference: Porting and Evaluating Diffusion-Based Speculative Decoding on TPU
- B01-1** Utilizing Cashflow Underwriting for Fairer Credit Assessment
- B01-2** Evaluating Credit Risk Through Transaction-Level Cash Flows: A Data-Driven Alternative to Traditional Scoring
- B01-3** Beyond Credit Scores: Transaction-Level Modeling for Credit Risk
- B02** Agentic Privacy Control Center
- B03-1** Integrating Wearable Devices with Digital Twin Models for Personalized Health Insights
- B03-2** Patient-Specific Cardiovascular Digital Twin for Personalized Cardiac Monitoring
- B04** DataLingo
- B05-1** Quantifying the Economic Impacts of Potential Sea-Level Rise
- B05-2** Predicting HAB Risk Through Climate Modeling and Machine Learning
- B06-2** Modeling NBA Team Dynamics via Wasserstein Gradient Flows
- B06-3** Body region matching with high density point clouds
- B07** Quality of Service Anomaly Detection in Wildfire Monitoring Networks
- B09** Hunting for Ghost Particles
- B12** Alternative Input Interfaces for Audio Consumption
- B13-1** Wildfire Property Intelligence: Finding Outliers Before Fire Finds Them First
- B13-2** Wildfire Severity and Structural Damage Prediction in Wildland-Urban Interface (WUI) Areas
- B13-3** Wildfire Frequency and Damage Prediction for Real Estate Applications
- B14-2** Structural vs. Chemical: A Statistical Approach to the Multimodal Covariance Structure of the Human Brain
- B15-1** Balancing Privacy and Utility: Differentially Private Synthetic Data Generation for Intel Telemetry
- B16-2** Personalized Supplement Recommender Using Agentic AI and Knowledge Graphs
- B19-1** PREVAIL: Predictive Response for Emergency Volume Assessment in Incident Locations

- B19-2** Seeing the System: A World Model for SDG&E Scheduling
- B19-3** ROAM: From Data to Dispatch
- B24-2** GenAI For Good: Detecting Misinformation Online
- B24-3** Multi-Agent Misinformation Analysis System
- B24-4** LiarDB
- B25** Workforce Optimization on Intuit's Virtual Expert Platform
- B26-1** Agentic and Fractal Reasoning for Article-Level Automated Fact-Checking
- B26-2** Misinformation/Disinformation Detection Using Agentic AI

Blockchain

- B10** Vera: Blockchain-Backed Trust

Optimization & Training Dynamics

- A05** Beyond Retraining: investigation on how early stopping and sensitivity of different unlearning methods help mitigate the adversarial attack of a simple model
- A14-1** Shortening the Loss Plateau
- A14-2** Investigating Optimizer Behavior Under Varying Training Regimes
- B15-2** Training-free versus training-based differential privacy for data synthesis

Systems & Hardware

- A02-1** Hardware Acceleration for Tensor Multiplication in Neural Networks
- A02-2** Hardware Acceleration for Tensor Multiplication in Neural Networks
- A08-2** Inference-Time Scaling for GUI Agents with Process Reward Models and Internal World Models