

Thursday February 27	Author	Title	Session Chair
Morning			Bistra Dilkina
9:30-10:45am	Zheyuan Ryan Shi	Scalable, Sustainable, Generalizable, and Responsible AI for Public Sector	
	Elizabeth Bondi-Kelly	Realizing AI for Impact: Towards Participatory Human-AI Collaboration for Water Conservation and Reproductive Health	
Morning			Bistra Dilkina
11:15am-12:30pm	Jun Wu	Towards Trustworthy Machine Learning under Distribution Shifts	
	Xiaorui Liu	Harnessing Robust Statistics for Trustworthy AI	
	Linyi Li	Certified Trustworthiness in the Era of Large Language Models	
Afternoon			Anuj Karpatne
2:30-3:45pm	Daniel Brown	Leveraging Human Input to Enable Robust, Interactive, and Aligned AI Systems	
	Huazheng Wang	Efficient and Robust Reinforcement Learning from Human Feedback	
Friday February 28			
Morning			Sriraam Natarajan
9:30-10:45am	Roberto Martín-Martín	Learning Robots with Physical Interactive Intelligence	
	Erdem Biyik	Efficient Robot Learning via Interaction with Humans	
	Ruohan Gao	Multisensory Machine Intelligence	
Morning			Fei Fang
11:15am-12:30pm	Shengpu Tang	Transforming Healthcare Decision Making Using Artificial Intelligence	
	Tianlong Chen	Breaking the Resource Monopoly from Industries: Sustainable and Reliable LLM Serving By Recycling Outdated and Resource-Constrained GPUs	
	Kaiqing Zhang	Foundations of Multi-Agent Learning in Dynamic Environments: Where Reinforcement Learning Meets Strategic Decision-Making	
Afternoon			
2:30-3:45pm	Dongkuan Xu	Advancing Reliable Edge Intelligence	Biplav Srivastava
	Mahdi Khalili	Mitigating Bias in Machine Learning: A Comprehensive Review and Novel Approaches	
	Xiang Li	Every Opinion Matters: Evaluating and Building Models with Pluralistic Views	
Afternoon			Matt Taylor
4:15-5:30pm	Jiaqi Ma	Data-Centric Methods for Trustworthy AI Development	
	Evi Micha	Axioms for AI Alignment from Human Feedback	
	Serena Booth	Lessons Learned as an AI Policy Advisor in the United States Senate	
Saturday March 1			
Morning			Xiaoli Fern
9:30-10:45am	Aryan Deshwal	Adaptive Experimental Design to Accelerate Scientific Discovery and Engineering Design	
	Wuyang Chen	Automated, Interpretable, and Scalable Scientific Machine Learning	
	Vivek Gupta	Advancements in AI for Reasoning with Complex Data	
Morning			Parisa Kordjamshidi
11:15am-12:30pm	Freda Shi	Learning Language Structures through Grounding	
	Yunzhu Li	Learning Structured World Models From and For Physical Interactions	
	Manling Li	From Large Language Models to Large Action Models: Planning and Reasoning with Physical World Knowledge	
Sunday March 2			

Morning			Bo Liu
9:30-10:45am	Ruqi Zhang	Scalable and Efficient Probabilistic Inference for Bayesian Deep Learning and Generative Modeling	
	Zhaozhuo Xu	Compression-Aware Computing for Scalable and Sustainable AI	
Morning			Pradeep Varakantham
11:15am-12:30pm	Marlos C. Machado	Representation-driven Option Discovery in Reinforcement Learning	
	Yixin Wang	Representation Learning: A Causal Perspective	
Afternoon			Jana Doppa
2:00-3:45pm	Weiyang Shi	Persuasion for Social Good: How to Build and Break AI	
	Tian Li	Scalable and Trustworthy Learning in Heterogeneous Networks	
	Aviral Kumar	Scaling up Offline Reinforcement Learning	