

| Thursday February 27      | Author                | Title  | Session Chair       |
|---------------------------|-----------------------|--|---------------------|
| <b>Morning</b>            |                       |  | Bistra Dilkina      |
| <b>9:30-10:45am</b>       | 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                     |                     |
|                           |                       |  |                     |
| <b>Morning</b>            |                       |  | Bistra Dilkina      |
| <b>11:15am-12:30pm</b>    | 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  |                     |
|                           |                       |  |                     |
| <b>Afternoon</b>          |                       |  | Anuj Karpatne       |
| <b>2:30-3:45pm</b>        | Daniel Brown          | Leveraging Human Input to Enable Robust, Interactive, and Aligned AI Systems   |                     |
|                           | Huazheng Wang         | Efficient and Robust Reinforcement Learning from Human Feedback  |                     |
|                           |                       |  |                     |
| <b>Friday February 28</b> |                       |  |                     |
| <b>Morning</b>            |                       |  | Sriram Natarajan    |
| <b>9:30-10:45am</b>       | Roberto Martín-Martín | Learning Robots with Physical Interactive Intelligence   |                     |
|                           | Erdem Biyik           | Efficient Robot Learning via Interaction with Humans   |                     |
|                           | Ruohan Gao            | Multisensory Machine Intelligence  |                     |
|                           |                       |  |                     |
| <b>Morning</b>            |                       |  | Fei Fang            |
| <b>11:15am-12:30pm</b>    | 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                |                     |
|                           |                       |  |                     |
| <b>Afternoon</b>          |                       |  |                     |
| <b>2:30-3:45pm</b>        | 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   |                     |
|                           |                       |  |                     |
| <b>Afternoon</b>          |                       |  | Matt Taylor         |
| <b>4:15-5:30pm</b>        | 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  |                     |
|                           |                       |  |                     |
| <b>Saturday March 1</b>   |                       |  |                     |
| <b>Morning</b>            |                       |  | Xiaoli Fern         |
| <b>9:30-10:45am</b>       | 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   |                     |
|                           |                       |  |                     |
| <b>Morning</b>            |                       |  | Parisa Kordjamshidi |
| <b>11:15am-12:30pm</b>    | 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                                  |                     |
|                           |                       |  |                     |
| <b>Sunday March 2</b>     |                       |  |                     |

|                        |                   |   |                     |
|------------------------|-------------------|---|---------------------|
| <b>Morning</b>         |                   |   | Bo Liu              |
| <b>9:30-10:45am</b>    | Ruqi Zhang        | Scalable and Efficient Probabilistic Inference for Bayesian Deep Learning and Generative Modeling |                     |
|                        | Zhaozhuo Xu       | Compression-Aware Computing for Scalable and Sustainable AI                                       |                     |
|                        |                   |   |                     |
| <b>Morning</b>         |                   |   | Pradeep Varakantham |
| <b>11:15am-12:30pm</b> | Marlos C. Machado | Representation-driven Option Discovery in Reinforcement Learning                                  |                     |
|                        | Yixin Wang        | Representation Learning: A Causal Perspective   |                     |
|                        |                   |   |                     |
| <b>Afternoon</b>       |                   |   | Jana Doppa          |
| <b>2:30-3:45pm</b>     | 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   |                     |