

ARIA 2023 | EXHIBITOR LIST

Industrial Exhibitors

Booth #1
Walmart Global Tech

Booth #2 & #3
SOTI

Booth #4
AMD

Departmental Exhibitors

Booth #5
Master of Mathematical Finance, U of T

Booth #8
Dept. of Computer Science, U of T

Booth #6
Masters of Financial Insurance, U of T

Booth #9
Master of Science in Applied Computing, U of T

Booth #7
Dept. of Statistical Sciences, U of T

Demonstrations

Demo #1 | Mustafa Haiderbhai, Radian Gondokaryono, Lueder A. Kahrs
Multi-platform Autonomous Surgical Robotics From Simulation Using Reinforcement Learning
Our project aims to perform autonomous surgical tasks using Reinforcement Learning

Demo #2 | Andrii Lenyshyn
“WaM” Cognitive Assessment in VR
I’m working on a project that focuses on using gamification techniques for cognitive assessments in Virtual Reality

Demo #3 | Fenil Patel, Chrisee Zhu, Elliot Chen
Triple Point
Award-winning game at Ontario’s Level Up Student Showcase 2023

Demo #4 | Shivanshi Gupta
The Moon Man
A captivating first-person exploration puzzle game set in a mesmerizing Virtual Reality landscape

Demo #5 | Soso (Zhifei) Song
3D Perspectives in Gaming: Exploring its Mechanics and Level Design Applications
Illucid originated from a game design course and introduces a unique first-person puzzle mechanism

Demo #6 | Brenna Li, Anna Kirk, Jiaqi Guo, Liam Baker, Khai Truong, Alex Mariakakis
Pre-consultation Medical Chatbot
Chatbots supported by Large Language Models are effective at helping doctors gather information

Demo #7 | Pritish Mishra, Brian Ramprasad
Hierarchical Data Management for Edge-Cloud Infrastructure
Data processing layer and a file-storage system built on top of a hierarchical computing infrastructure

Demo #8 | Sophie Liu
DynaCOMP: Bringing ML to the Transplantation Clinic
Created a machine learning model which stratifies patients into different risk categories

Demo #9 | Jasper (Jan-Philip) Gergik
Playing Various Strategies in Dominion Using Deep Reinforcement Learning
Reinforcement learning based agent with a novel multiset-based architecture

Demo #10 | Bingjian Huang, Mason Cai, Conner Ren, Jerry Cheng, Abel Sang
NebulaSuit: An Upper-Body High-Density Vibrotactile Suit Built With Scalable Hardware Solutions
VibroStar, a high-resolution vibrotactile suit that covers the upper body surface with dense arrays of actuators

Demo #11 | Aida Ramezani
Moral Norm Variation in Large Language Models
Computational tools to analyze language change in both temporal and cultural levels



Dept. of Computer Science, Research Stream

Dynamic Graphics Project Lab of the University of Toronto

Presentation #15 | Chuyi (Sky) Hou
Headshift

Enhancing VR immersion by addressing headset slippage ensures consistent visuals, reduces discomfort, and elevates the overall user experience.

IATSL: Intelligent Assistive Technology and Systems Lab

Presentation #16 | Soroush Mehraban, Vida Adeli, Babak Taati
MotionAGFormer: Enhancing 3D Human Pose Estimation with a Transformer-GCNFormer Network

A model with less parameters and computation compared to SOTA models while being more accurate in 3D Human pose estimation.

Presentation #17 | Vida Adeli
Vision-Based Analysis of Gait in Older Adults with Dementia

Unlocking new insights into gait patterns in dementia with computer vision, paving the way for personalized fall prevention and tailored interventions.

U of T Systems and Networks Group (SysNet)

Presentation #18 | Christina Christodoulakis
Pytheas: Pattern-Based Table Discovery in CSV Files

We present Pytheas, our supervised learning approach for accurately discovering tables in Open CSV files, with a confidence metric for efficiently labeling data.

Presentation #19 | Alexey Khrabrov
JITServer: Disaggregated JIT Compiler for the JVM

Improving JVM performance and reducing CPU and memory costs in the cloud using remote JIT compilation and caching compiled code and profiling data.



Master of Science in Applied Computing (MScAC)

16 Bit Inc.

Presentation #20 | Yuxuan Wang
Broca

Leveraging multimodal prompt tuning and large language model for Automated Radiology Reporting from X-ray Imagery.

Presentation #21 | Abdur Rahman, Sarthak Narayan
Revolutionizing Healthcare Application Deployment: A Kubernetes-Driven Approach

A seamless and reliable deployment and management solution for healthcare applications on-premise and cloud using Kubernetes.

Adobe

[Presentation #22 | Lukas Jurisica](#)

Vision-Based Analysis of Gait in Older Adults with Dementia

By utilizing an alternative coordinate system based on Integer Barycentric Coordinates for representing meshes, we are able to robustly and efficiently apply certain geometry processing algorithms.

Agentnoon

[Presentation #23 | Sarah Rozaidi](#)

Multi-Regional Salary Prediction Model

Empowering strategic decisions with comprehensive insights into compensation trends worldwide

AMD

[Presentation #24 | Sanjukta De](#)

Learning GPU Code Structure using Large Language Models

Revolutionizing GPU Kernel Optimization: Harnessing Transformer-Based Models to Improve Performance and Efficiency

[Presentation #25 | Vignesh Edithal](#)

Automatic Gameplay Highlight Reel Generation using Multimodal Learning

Detect and merge interesting events in first person shooting games using multimodal video transformer model to generate highlight clips

[Presentation #26 | Geoffrey Harrison](#)

Task Grouping for Multi-Task Reinforcement Learning

Measuring relationships between tasks enables disentangling knowledge transfer in multi-task RL.

[Presentation #27 | Yi Cheng \(Ethan\) Zhu](#)

3D Image Representation for Gameplay Automation

Self-Supervised image representation encoding methods do not generalize to novel exploration and other Reinforcement Learning tasks.

[Presentation #28 | Zichuan Guan](#)

Machine Learning for Pattern-Matching Optimization in GPU Compilers

Advancing GPU performance through AI-Driven compiler optimizations

[Presentation #29 | Mahesh Jayasankar](#)

Real-time Camera Tracking for Gameplay Automation in Video Games

A navigational assistance module for video games that works in real-time using only a single camera's footage

[Presentation #30 | Akash Haridas](#)

High-Fidelity Human Avatars from Laptop Webcams Using Edge Compute

Personalized photorealistic avatars for video calls and virtual reality, generated automatically using your webcam and running locally.

[Presentation #31 | Donghui \(Alex\) Xu](#)

Content-Aware Video Encoder Optimization

Elevating Conventional Video Encoder with Deep Learning Models for Enhanced Video Quality and Bitrate Efficiency

[Presentation #32 | Adeem Aryn Jassani](#)

AI-Based Saliency Prediction For Better Perceptual Quality in Display and Video Compression

Ultimate visual experience, from natural scenes to gaming contents and beyond, with our real-time saliency prediction-based quality improvement

[Presentation #33 | Amandeep Singh](#)

Automatically Parallelizing C/C++ Code with Large Language Models

Using Large Language Model to automatically classify and specify OpenMP directives for the Sequential C/C++ code.

Apollo.io

[Presentation #34 | Anurag Bajpai](#)

Prospect Recommendations with Integrated LLM Features

Leveraging LLM-generated Ideal Customer Profiles and Text Embeddings to Recommend Prospects to Sales Teams

Arctic Wolf Networks

Presentation #35 | Cijie Xia

Query Broker: An Efficient Cybersecurity Data Access Platform

Enabling Efficient, Secure Data Retrieval for Enhanced Security and AI Services

BenchSci

Presentation #36 | Tirth Patel

Efficiently Training Large Language Models with Advanced Distributed Training Techniques

Unlocking Peak Efficiency: Scaling Large Language Model Training with DeepSpeed's ZeRO Optimization Strategies and Vertex AI's Distributed Training.

Bluecat Networks

Presentation #37 | Rajesh Marudhachalam, Ruixuan Zhang

Detecting Malicious DNS Traffic Using Neural Networks

Enhancing Network Security Using Weakly-Supervised Learning and Pairwise Relation Prediction Network

BlueDot Inc.

Presentation #38 | Naveen Thangavelu

A Global Surveillance System for Infectious Diseases and Bio-Security Threats

Revolutionizing Outbreak Analysis Using Large Language Models for Comprehensive Disease Event Extraction

CPP Investments

Presentation #39 | Dixin Mou

Enhancing Forecasting Performance of LightGBM through Feature Screening

A systematic process to mitigate the impact of high feature dimensions while retaining informative features to produce forecasts with better out-of-sample ranking accuracy and volatility

Presentation #40 | Jie (Alisa) Yang

Exploring Graph Neural Network Methods on the Analyst Network Graph

Both Node2Vec and Graph Convolutional Networks can be used to exploit the topology of the Analyst Network to explain companies returns

Centre for Global Health Research (CGHR), Unity Health Toronto

Presentation #41 | Andy Lee

A Machine Learning Framework for Exploring Mortality in Developing Countries with Verbal Autopsies

Harmonizing verbal autopsy research into developing a conceptual framework which provides a roadmap to automating cause-of-death prediction using machine learning

Crossing Minds

Presentation #42 | Jason Tang

Image Captions Are Worth a Thousand Words: Enhancing Search in Conversational Recommendations

Both Node2Vec and Graph Convolutional Networks can be used to exploit the topology of the Analyst Network to explain companies returns

Presentation #43 | Anusha Prabhudev

DataStats: An Automated Framework Supporting the Design and Development of Recommender Systems

An innovative framework empowering recommender systems development with automated analysis, improved decision-making support, and domain-focused insights for enhanced personalization and efficiency.

Double Negative Ltd.

Presentation #44 | Kartikaeya Kumar

Automating Rotoscoping in Videos with NeRF Priors

Are NeRFs the new greenscreen?

Presentation #45 | Yihan (Nick) Ni

Latent DeepFakes: Fast High-Resolution Face Swap

Autoencoder based DeepFake face swap trained in Stable Diffusion latent space with reduced identity leakage using Vector Quantization layers.

Ecomtent

Presentation #46 | Zakaria Patel

Enhancing Image Composition Control with Loss-Guided Diffusion Models

Leveraging the synergy between loss guidance and attention injection to achieve precise layout and composition control.

Presentation #47 | Hantang Li

Generative AI for Product Listing and Conditional Infographic Generation

Design reliable, scalable, and maintainable machine learning system for startup and exploration of conditional layout generation models for product infographic generation.

Geotab Inc.

Presentation #48 | Xinyu Zhang

Data Quality Detection and Improvement for AI Dashcam Models

Algorithms and Experiments

Presentation #49 | Zhifan Wang

From ML Model to Production: Streaming Data Analytics Infrastructure for IoT Big Data

A modern approach for deploying machine learning applications on cloud-based infrastructure integrating distributed computing and streaming capabilities for IoT Big Data use cases

Presentation #50 | Xinrong Zhou

Identifying Risk Factors for Hazardous Driving and Accident Propensity for Safer Fleets and Smart Cities

Smarter cities and safer roads start here: Unveiling the impact of collective vehicular conduct within an intersection on accident propensity

Presentation #51 | Anh Tuan Tran

Automatic Detection of Service Center Locations from Vehicle Data

Clustering a large number of vehicle maintenance events and applying an anomaly detection model to detect service centers

Google

Presentation #71 | Yu-Chieh (Sunny) Wu

Image Caption Generation with Machine Learning for AR Wearable Devices

Elevate AR wearable devices with AI-generated image captions: Optimizing user focus, experience, and productivity.

Presentation #72 | Jingtao Zhou

Visual Blocks for ML

A visual programming framework for prototyping machine learning-based application pipelines

Huawei

Presentation #52 | Tao Zhang

LayoutDM: A Discrete Diffusion Approach for Optimal Smartphone Widget Layout Generation

Revolutionizing Smartphone UI Design with LayoutDM: A Discrete Diffusion Approach for Optimal Widget Layouts

Presentation #53 | Samarendra Dash

Leveraging LLMs for Operations Research Problems

Can LLMs reason about mathematical modelling from natural language descriptions?

Presentation #54 | Benjamin Correia

OmniRuntime: Squeezing Every Last Drop of Performance for In-Memory Data Analytics

Hand optimizing an SQL analytics engine through low level analysis of operators and integrating hardware accelerators.

Presentation #55 | Deyu Liu, Duy Tran

IOMemTable — Memory-Optimized Index-Organized Table for LSM-tree Storage

Toward a better trade-off of index performance for Hybrid Transactional/Analytical Processing (HTAP) workloads

Intel

Presentation #56 | Yuhang Liu

Enhancement of IP Wizard in Power and Thermal Calculator

Enrich the functionality of IP Wizard and improve the user experience

Iteration Matrix

Presentation #57 | Jiujiu (Jojo) Duan

AI-driven Web Payment Automation

Efficiently Automating High-Volume, Time-Sensitive Payments with Google reCAPTCHA Overcoming Solutions.

J-Squared Technologies

Presentation #58 | Jaskaran Singh, Gurman Bhullar

Revitalizing Retail: Edge-Driven Computer Vision for Real-Time Analytics

Infusing AI Excellence into Retail: Harnessing State-of-the-Art Deep Learning for Instantaneous Insights, Robust Security and Advanced Customer Analytics.

Kindred, Part of Ocado Group

Presentation #59 | Jiachen (Charlie) Xu

Efficient Embedding-Based Item Identification for Robotic Pick and Place Applications

Transforming Item Identification in Robot-Assisted Grocery Fulfillment: An Embedding-Driven Solution for Precision and Efficiency

Presentation #60 | Junbo Huang

Object Tracking for High-Speed Pick-and-Place Robot

Enabling Faster Robotic Operations with Real-Time, Detection-Based Item Tracking and Prediction Strategies

MDA

Presentation #61 | Val Kobilaski

CAPP: Conditional Variational Auto-Encoder Accelerated Path Planner

A machine-learning approach to path-planning for a 7-DOF robotic manipulator

Metabob

Presentation #62 | Anthony Rinaldi

Abstract Syntax Tree Coarsening via Deep Attention-Based Node Pooling Auto-Encoders

Attention-based graph coarsening allows for significant reduction of the size of graphs, while retaining sufficient information for graph reconstruction, as well as downstream classification tasks versus traditional algorithmic coarsening methods.

ModiFace

Presentation #63 | Kejia Yin

SCE-MAE: Selective Correspondence Enhancement with MAE for Self-Supervised Landmark Estimation

Learning dense discriminative representations from unlabelled data

Presentation #64 | Xudong Liu

Detailed Cross-Attention Control in Image Editing

Refining Image Editing: Achieve Personalized, Pinpoint Adjustments While Preserving Identity and Quality.

Presentation #65 | Yuqing (Sophia) Zeng

Real-Time 3D Landmark Regression from Monocular Image Input on Mobile GPUs with Synthetic 3D Face Data

Describing human face shapes and positions from face images remains a popular topic at the intersection of computer vision and deep learning.

Layer 6 AI

Presentation #67 | Ilan Gofman

Enhancing Document Retrieval with Large Language Models and Human Feedback

Using generative AI to make retrieval with human feedback more efficient

Presentation #68 | Linfeng Du

MultiResFormer: Transformer with Adaptive Multi-Resolution Analysis for Long-Term Time Series Forecasting

The first Transformer-based model with Fourier-guided time series segmentation which surpasses naive patch-based Transformer and CNN-based methods.

Presentation #69 | Apoorv Dankar

TabPFGen – Tabular Data Generation with TabPFN

TabPFGen creates an energy-based generative model for tabular data using the discriminative, in-context learning capabilities of pre-trained TabPFN transformer

Presentation #70 | Hamid Kamkari

Explaining the Out-of-Distribution Detection Paradox Through Likelihood Peaks

TabPFGen creates an energy-based generative model for tabular data using the discriminative, in-context learning capabilities of pre-trained TabPFN transformer

National Research Council Canada

Presentation #74 | David Guzmán

Speech Generation for Indigenous Language Education

Improving Speech Generation for Low Resource Languages via Spectrogram Denoising

Nureva

Presentation #75 | Miaoqi (Frank) Zhang

Sound Event Localization and Detection

Leveraging CRNNs in SELD: Elevating Precision in Sound Source Localization by using linear arrays of microphones.

Presentation #76 | Wenzhe Xu

Towards Broadcast-Quality Audio: Enhancement of Conference Microphones and Exploration of Objective Evaluation Techniques

Merging Deep Learning and Established Metrics for Accurate Audio Quality Assessment

Ocado Advanced Technology

Presentation #77 | Sumant Bagri

Hardware Accelerated Collision Checking

Improve collision detection latency for robotic manipulators by leveraging SIMD hardware

Pearson Canada Inc.

Presentation #78 | Hongbo (Eleanor) Zhou

Anomaly Detection for Automated English Language Test Scoring Improvement

Enhancing automated English language test scoring precision with multi-speaker audio data and anomaly detection.

Pelmorex

Presentation #79 | Charie Brady

A Machine Learning Approach to Retail Sales Forecasting Using Localized Weather Features

Weather-driven consumer demand: building explainable machine learning models that adapt to changing local weather conditions for precise sales forecasting

Princess Margaret Cancer Center

Presentation #80 | Jean Yaacoub

GNNs with Protein Dynamics for Enhanced Drug Targeting

Analysis of binding affinity prediction models using Graph Neural Networks with modified protein features.

QEYnet

Presentation #81 | Michael Luciuk

Advancing the Security of Commercial Quantum-Key-Distribution Systems

A Proactive Approach to Vulnerability Assessment and Resolution

Roche

Presentation #82 | Mingyi (Horus) Li

DORA: DevOps Research and Assessment

Track code development efficiency and progress in one platform

Samsung AI Center – Montreal

Presentation #83 | Rui Heng Yang

Accurate and Precise Force Estimation Algorithm for Robot Manipulation

Combining inverse finite element methods with neural networks to provide a robust and data-efficient algorithm for force estimation.

Samsung AI Center – Toronto

Presentation #84 | Tianshu (Chris) Kuai

Unsupervised Real Image Denoising

Using a pre-trained off-the-shelf diffusion model as an image prior for unsupervised real-world image denoising.

Presentation #85 | Ke Zhao

Towards Intelligent Edge Computing

From Research Proof-of-Concept Server-based Pipeline to on-Device Deployment: A Case Study

Presentation #86 | Vahid Zehtab

LUTFlow: Conditionally Bijective Implicit 3D Lookup Tables

Embed infinite photo styles in your pictures; if a style doesn't suit, simply roll it back or craft your very own!

Sanofi

Presentation #87 | Chaorui (Tom) Zhang

Closeness to Pathway: Clustering Patient Electronic Medical Records for Drug Repurposing

Evaluating and comparing drug repositioning approaches that use unsupervised learning and patient data from electronic medical records against supervised learning with known drug-disease relationship.

Scotiabank

Presentation #88 | Tsz Fung Yau

Two-Stage Deep-Learning Framework for Understanding Scanned Documents

Using a deep OCR model to identify texts and a transformer language model to extract relevant texts

Shopify

Presentation #89 | Tianquan (Andy) Di

Towards Optimized E-Commerce: Examining Performance-affecting Factors for Online Retail Platforms

Through A/B testing and user studies, we analyzed the impact of three experience enhancements on the conversion rates of e-commerce platforms.

Signal 1

Presentation #73 | Jingtao Zhou

Visualizing AI-Powered Patient Insights

Enhancing Machine Learning Interpretability and Explainability in Healthcare

Presentation #90 | Shixuan (Nadia) Li

Using Multi-Modal Data and Self-Supervised Approaches for Deep Learning in Healthcare

Transformers that simultaneously learn from time series and text data via self-supervised learning, thereby improving clinical predictions

SOTI

Presentation #91 | Yichen Shen

Cross-Account and Cross-Platform Cloud Cost Optimization: An Integrated Management for AWS and Azure

Harmonizing Costs: Streamlining Management Across Multiple Cloud Platforms and Accounts

Presentation #92 | Yao-Chen Chi

Benchmark Performance of Remote Control Operations Across Varied Mobile Device in Different Environments

Quantifying Latency in Remote Android Device Management: Emulation, Measurement, and Scaling Insights.

Presentation #93 | Xinman Liu

Simulation of Device Location

Optimizing the testing of location-based services and location-aware applications with a configurable location simulation tool.

Presentation #94 | Jazib Ahmad

Reactive-Proactive Obstacle Avoidance for Autonomous UAVs in Dynamic Indoor Environments

Drones Coexisting Peacefully with Humans & Gracefully Evading Unplanned Encounters

Presentation #95 | Hosuk (Jimmy) Woo

Smartphone Indoor Localization: A Sensor Fusion Approach

Seamless indoor localization of smartphones without use of external hardware devices

Presentation #96 | Jaspreet Singh Dhani, Simrat Bains

Cracking the Code: Gherkin Test Prioritization for Improved CI/CD Efficiency

Expediting regression and acceptance testing for incoming code changes by using software metrics and NLP methods to pinpoint an optimal Gherkin test subset.

Presentation #97 | Yuxuan Zhang

Model Compression for Unmanned Aerial Vehicle's Companion Computer

Accuracy & Latency Optimization and Tradeoff for Segmentation Models

Presentation #98 | Devashish Pradeep Khairnar

MABS: Make All Batteries Smart

Enhancing Battery Performance and Health Prediction: Smart Solutions for Smart and Non-Smart Lithium-ion Batteries

Presentation #99 | Xuduo Gu

Network Traffic Anomaly Detection by Clustering with Autoencoder Ensembling

An expectation-maximization approach to detect network traffic anomalies in real-time by clustering network packets with autoencoders.

Presentation #100 | Ruoqi Shu

Data Communication Optimization Between Mobile Devices and Servers

Insight of Enhancing Protocol Performance Through Bandwidth Analysis and Predictive Modeling

Presentation #101 | Yulin (Patrick) Wang

Driver Behavior Analysis Using Sensor Data On Android Devices

Software solution capable of collecting and analyzing sensor data from android devices and reporting road safety and driving behavior metrics for moving vehicles

SpassMed Inc.

Presentation #102 | Chen Dan, Bingjie (Lily) Shen, Yuwei Liu

Early Septic Shock Detection Using Vital Signs Forecasting and Classification

Utilizing deep learning, we forecast vital signs and employ different methods to detect sepsis and septic shock in time series data.

Tesla

Presentation #66 | Ilan Gofman

Machine Learning for Vehicle Diagnostics Automation

Investigating Time Series Classification Methods for Vehicle Failure Mode Detection

The Hospital for Sick Children

Presentation #103 | Adrian Yung

Analysis of Single Cells in Pediatric Brain Tumours with Transfer Learning and a Pre-trained Transformer Model

Transfer learning and a context-aware, attention-based, deep learning transformer model is applied to analyze individual cells from pediatric brain tumours, simulating gene overexpression and deletion to investigate contributors to metastasis

Presentation #104 | Sarah Hindawi

Prediction of Genetic Mutations in Pediatric Brain Tumors Using Pathology Images

Multimodal Genetic Mutation Prediction: Leveraging Convolutional Neural Networks for Advanced Histology Analysis and Evaluating the Impact of Clinical Data Integration

The Vanguard Group

Presentation #105 | Keerat Guliani

Enabling Large Language Models at Vanguard

Efficient Training and Usage of Domain Specific LLMs

Ubisoft

Presentation #106 | Tuan Nguyen

Facial Landmark Detection with Synthetic Data

Detecting facial landmarks on videos, by leveraging a synthetic dataset of rendered 3D faces

Presentation #107 | Yan Ma

Learning Unrestricted Facial Geometry from Lightstage Images

Push the quality of high-fidelity 3D face avatar generation to reach AAA production-level needs.

Presentation #108 | Yuchen Chen

Video Based Animation Retrieval

Finding similar animations using human motions in video sources

Unilever

Presentation #110 | Xiaolan (Landy) Xu

Advancements in Multi-Objective Optimization Algorithms to Trade Promotion Strategy

Employing stochastic search methods and the temperature cooling schedule for Multi-Objective Optimization Algorithms

Presentation #111 | Jingbo Yang

Analyzing Sales Cannibalization Through Utility Theory to Improve Accuracy of Predicting POS

Quantifying the Impact of Product Attributes on Consumer Purchase Decisions

Presentation #112 | Yu Pang

Project Sonata: AI-Driven Recommendation & Consumer Segmentation Platform of Unilever NA

A Recommendation Engine to Enable Multiple Related Downstream Tasks for Unilever NA

University Health Network

Presentation #109 | Haoran Sheng

Unsupervised Data Extraction from Radiology Reports with LLMs

Building an unsupervised NLP pipeline to extract nuggets of data from a vast 20 year archive of unstructured radiology reports.

Presentation #113 | Haining Tan

Visual Perception of Walking Environments Using Deep Learning and Large Language Models

Stair obstacle detection using AI-powered smart glasses and image captioning of walking environments using generative vision-language models to provide assistance to patients with visual impairments

University of Toronto Schools

Presentation #114 | Jingyu Liu

Transforming Learning Through AI-Enhanced Online Discussions: Collaboration, Equity, and Assessment

Helping support 'sober second thought' before the online harm begins

Unmodal Research

Presentation #115 | Cheng (Jamie) Xu

Project Li-Tooth: Scalable, High-Bandwidth Optical Information Broadcast for Spatial Computing

Space aware, line-of-sight data broadcast using VLC/OCC hybrid

Veeva Systems

Presentation #116 | Haoteng Wang

Explore New Initiatives on RIM Submissions Archive's Support of eSTAR FDA Medical Device Submission Format

With a focus on memory optimization for the importing and extraction of eSTAR submissions