Agenda

Edge AI Deep Dive Day™
Tuesday May 21

8:00 am - 5:00 pm
Registration

Separate Registration Required ($25/Session)

Qualcomm

1:00 pm - 4:00 pm
Great America Ballroom

Accelerating Model Deployment with Qualcomm AI Hub

In this workshop we address the common challenges faced by developers migrating AI workloads from workstations to edge devices. Qualcomm aims to democratize AI at the edge, easing the transition to the edge by supporting familiar frameworks and data types. Additionally, we offer a range of tools to assist advanced developers in optimizing performance and power consumption for their AI applications.

In this session, we empower developers with knowledge and tools to efficiently deploy optimized models on real devices using the Qualcomm AI Hub. We'll begin with an introduction to the Qualcomm AI Hub. We'll then walk through examples showing how to optimize models using the AI Hub. Attendees will also have opportunities to speak with Qualcomm's AI model deployment experts.

Join us to learn how to leverage the AI Hub's capabilities to streamline the deployment process, enabling you to bring your AI applications to life in a matter of minutes.

This session is intended for machine learning engineers and AI/ML application developers, including those creating Android, Windows and IoT/AIoT applications.

Summit
Wednesday May 22

7:30 am - 7:00 pm **Registration**	**7:30 am - 9:00 am** **Coffee and Pastries**

9:00 am - 11:10 am
Introduction, Keynote and General Session
Mission City Ballroom—B1-B5

INTRODUCTION
A View from the Summit (Part 1)
Jeff Bier, Founder, Edge AI and Vision Alliance

KEYNOTE
Learning to Understand Our Multimodal World with Minimal Supervision
Yong Jae Lee, Associate Professor, Department of Computer Sciences, University of Wisconsin-Madison

GENERAL SESSION
Scaling Vision-Based Edge AI Solutions: From Prototype to Global Deployment
Maurits Kaptein, Chief Data Scientist, Network Optix and Professor, University of Eindhoven

12:30 pm - 7:30 pm Technology Exhibits Exhibit Hall

11:10 am - 11:25 am Break	
11:25 am - 12:30 pm	**Morning Sessions**
Technical Insights 1	Technical Insights 2
Fundamentals	Business Insights

12:30 pm - 1:30 pm Lunch Exhibit Hall	
1:30 pm - 3:10 pm	**Afternoon Sessions**
Technical Insights 1	Technical Insights 2
Fundamentals	Business Insights
Enabling Technologies 1	
Enabling Technologies 2	
Enabling Technologies 3	

3:10 pm - 4:15 pm Break	
4:15 pm - 5:55 pm	**Afternoon Sessions**
Technical Insights 1	Technical Insights 2
Fundamentals	Business Insights
Enabling Technologies 1	
Enabling Technologies 2	
Enabling Technologies 3	

6:00 pm - 7:30 pm Technology Exhibits Reception Exhibit Hall

6:00 pm - 7:30 pm	**Evening Events**

6:30 pm - 7:30 pm
Women in Vision Reception
Exhibit Hall—ET-3

Summit
Thursday May 23

7:30 am - 6:00 pm **Registration**	**7:30 am - 9:00 am** **Coffee and Pastries**

9:00 am - 11:10 am
Introduction, Awards Presentation, General Session and Panel
Mission City Ballroom—B1-B5

INTRODUCTION
A View from the Summit (Part 2)
Jeff Bier, Founder, Edge AI and Vision Alliance

GENERAL SESSION
What's Next in On-Device Generative AI
Jilei Hou, Vice President of Engineering and Head of AI Research, Qualcomm Technologies

Panel: Multimodal LLMs at the Edge: Are We There Yet?
Moderator: Sally Ward-Foxton, Senior Reporter, EE Times
Panelists: Adel Ahmadyan, Staff Engineer, Meta Reality Labs
Jilei Hou, Vice President of Engineering and Head of AI Research, Qualcomm Technologies
Yong Jae Lee, Associate Professor, Department of Computer Sciences, University of Wisconsin-Madison
Pete Warden, CEO, Useful Sensors

11:00 pm - 5:00 pm Technology Exhibits Exhibit Hall

11:10 am - 11:25 am Break	
11:25 am - 12:30 pm	**Morning Sessions**
Technical Insights 1	Technical Insights 2
Fundamentals	Business Insights
Enabling Technologies 1	
Enabling Technologies 2	

12:30 pm - 1:30 pm Lunch Exhibit Hall	
1:30 pm - 3:10 pm	**Afternoon Sessions**
Technical Insights 1	Technical Insights 2
Fundamentals	Business Insights
Enabling Technologies 1	
Enabling Technologies 2	

3:10 pm - 4:15 pm Break	
4:15 pm - 5:55 pm	**Afternoon Sessions**
Technical Insights 1	Technical Insights 2
Fundamentals	Business Insights

4:50 pm - 5:55 pm
Vision Tank Start-Up Competition
Theater (Upstairs)

April 22, 2024 Based on R05.
Contents preliminary and subject to change.

Following Jeff Bier's opening remarks, "A View from the Summit" (9:00 am - 9:40 am on Wednesday), join us for our Keynote and General Session!

9:40 am - 10:40 am
Wednesday
Learning to Understand Our Multimodal World with Minimal Supervision
Yong Jae Lee
Associate Professor, Department of Computer Sciences, University of Wisconsin-Madison

The field of computer vision is undergoing another profound change. Recently, "generalist" models have emerged that can solve a variety of visual perception tasks. Also known as foundation models, they are trained on huge internet-scale unlabeled or weakly labeled data and can adapt to new tasks without any additional supervision or with just a small number of manually labeled samples. Moreover, some are multimodal: they understand both language and images and can support other perceptual modes as well.

In our 2024 Keynote, Professor Yong Jae Lee from the University of Wisconsin-Madison will present recent groundbreaking research on creating intelligent systems that can learn to understand our multimodal world with minimal human supervision. He will focus on systems that can understand images and text, and also touch upon those that utilize video, audio and LiDAR. Since training foundation models from scratch can be prohibitively expensive, Yong Jae will discuss how to efficiently repurpose existing foundation models for use in application-specific tasks. He will also discuss how these models can be used for image generation and, in turn, for detecting AI-generated images. He'll conclude by highlighting key remaining challenges and promising research directions.

Join us to learn how emerging techniques will address today's neural network training bottlenecks, facilitate new types of multimodal machine perception and enable countless new applications.

About Speaker

Yong Jae Lee is an Associate Professor in the Department of Computer Sciences at the University of Wisconsin-Madison. His research interests are in computer vision and machine learning, with a focus on robust visual recognition systems that learn to understand the visual world with minimal human supervision. Before joining UW-Madison in 2021, he spent one year as an AI Visiting Faculty at Cruise and six years as an Assistant and then Associate Professor at UC Davis. He received his PhD from the University of Texas at Austin in 2012 and was a postdoc at Carnegie Mellon University (2012-2013) and UC Berkeley (2013-2014).

Professor Lee is co-author of the widely cited paper "Visual Instruction Tuning," which proposes LLaVA (large language and vision assistant), an end-to-end trained large multimodal model that connects a vision encoder and an LLM for general-purpose visual and language understanding. He is also co-author of "Segment Everything Everywhere All at Once," which proposes a novel decoding mechanism enabling diverse prompting for all types of segmentation tasks.

Professor Lee is a recipient of the ARO Young Investigator Program Award (2017), UC Davis Hellman Fellowship (2017), NSF CAREER Award (2018), AWS Machine Learning Research Award (2018 and 2019), Adobe Data Science Research Award (2019 and 2022), UC Davis College of Engineering Outstanding Junior Faculty Award (2019), Sony Focused Research Award (2020 and 2023) and UW-Madison SACM Student Choice Professor of the Year Award (2022). He and his collaborators received the Most Innovative Award at the COCO Object Detection Challenge, ICCV 2019 and the Best Paper Award at BMVC 2020.

10:40 am - 11:10 am
Wednesday
Scaling Vision-Based Edge AI Solutions: From Prototype to Global Deployment
Maurits Kaptein
Chief Data Scientist, Network Optix and Professor, University of Eindhoven

The Embedded Vision Summit brings together innovators in silicon, devices, software and applications and empowers them to bring computer vision and perceptual AI into reliable and scalable products. However, integrating recent hardware, software and algorithm innovations into prime-time-ready products is quite challenging. Scaling from a proof of concept—for example, a novel neural network architecture performing a valuable task efficiently on a new piece of silicon—to an AI vision system installed in hundreds of sites requires surmounting myriad hurdles.

First, building on Network Optix's fourteen years of experience, Professor Kaptein will detail how to overcome the networking, fleet management, visualization and monetization challenges that come with scaling a global vision solution. Second, Maurits will discuss the complexities of making vision AI solutions device-agnostic and remotely manageable, proposing an open standard for AI model deployment to edge devices. The proposed standard aims to simplify market entry for silicon manufacturers and enhance scalability for solution developers. Maurits will outline the standard's core components and invite collaborative contributions to drive market expansion.

General Session and Panel

Following Jeff Bier's opening remarks, "A View from the Summit" (9:00 am - 9:50 am on Thursday), join us for our General Session and Panel Discussion!

9:50 am - 10:20 am

Thursday

What's Next in On-Device Generative AI

Jilei Hou
Vice President of Engineering and Head of AI Research, Qualcomm Technologies

The generative AI era has begun! Large multimodal models are bringing the power of language understanding to machine perception, and transformer models are expanding to allow machines to understand using multiple types of sensors. This new wave of approaches is poised to revolutionize user experiences, disrupt industries and enable powerful new capabilities. For generative AI to reach its full potential, however, we must deploy it on edge devices, providing improved latency, pervasive interaction and enhanced privacy.

In this talk, we will share Qualcomm's vision of the compelling opportunities enabled by efficient generative AI at the edge. We will also identify the key challenges that the industry must overcome to realize the massive potential of these technologies. And we will highlight research and product development work that Qualcomm is doing to lead the way via an end-to-end system approach—including techniques for efficient on-device execution of LLMs, LVMs and LMMs, methods for orchestration of large models at the edge and approaches for adaptation and personalization.

10:20 am - 11:10 am

Thursday

Multimodal LLMs at the Edge: Are We There Yet?

Moderator: Sally Ward-Foxton
Senior Reporter, EE Times

Panelists:

Adel Ahmadyan
Staff Engineer, Meta Reality Labs

Jilei Hou
Vice President of Engineering and Head of AI Research, Qualcomm Technologies

Yong Jae Lee
Associate Professor, Department of Computer Sciences, University of Wisconsin-Madison

Pete Warden
CEO, Useful Sensors

Large language models (LLMs) are fueling a revolution in AI. And, while chatbots are the most visible manifestation of LLMs, the use of multimodal LLMs for visual perception—for example, vision language models like LLaVA that are capable of understanding both text and images—may ultimately have greater impact given that so many AI use cases require an understanding of both language concepts and visual data, versus language alone.

To what extent—and how quickly—will multimodal LLMs change how we do computer vision and other types of machine perception? Are they needed for real-world applications, or are they a solution looking for a problem?

If they are needed, are they needed at the edge? What will be the main challenges in running them there? Is it the nature of the computation, the amount of computation, memory bandwidth, ease of development or some other factor? Is today's edge hardware up to the task? If not, what will it take to get there?

To answer these and many other questions around the rapidly evolving role of multimodal LLMs in machine perception applications at the edge, we've assembled an amazing set of panelists who have firsthand experience with these models and the challenges associated with implementing them at the edge. Join us for a lively and insightful discussion!

◈ = Invited presentation

Session codes (e.g., T1W01) allow for quick searching in our event app!

| Technical Insights 1 Mission City Ballroom—B1-B5	Technical Insights 2 Mission City Ballroom—M1-M3	Fundamentals Great America Ballroom	Business Insights Theater (Upstairs)	Enabling Technologies 1 Exhibit Hall—ET-1	Enabling Technologies 2 Exhibit Hall—ET-2	Enabling Technologies 3 Exhibit Hall—ET-3
11:25 am - 11:55 am T1W04 Augmenting Visual AI through Radar and Camera Fusion *Sébastien Taylor, Vice President of Research and Development, Au-Zone Technologies*	**11:25 am - 11:55 am T2W04** Data-Efficient and Generalizable: The Domain-Specific Small Vision Model Revolution *Heather Couture, Founder and Computer Vision Consultant, Pixel Scientia Labs*	**11:25 am - 12:30 pm FW04** ◈ Introduction to Computer Vision with Convolutional Neural Networks *Mohammad Haghighat, Senior Manager, Core AI, eBay*	**11:25 am - 11:55 am BW04** Making Alexa More Ambiently Intelligent with Computer Vision *Michael Giannangeli, Senior Manager, Product Management, Alexa Devices, Amazon*			
12:00 pm - 12:30 pm T1W05 ◈ Building Meaningful Products Using Complex Sensor Systems *Dirk van der Merwe, Autonomous Robotics Lead, DEKA Research & Development*	**12:00 pm - 12:30 pm T2W05** Continual, On-the-Fly Learning through Sequential, Lightweight Optimization *Guy Lavi, Managing Partner, Vision Elements*		**12:00 pm - 12:30 pm BW05** Using Artificial Intelligence to Enhance the Well-Being of the Elderly *Harro Stokman, CEO, Kepler Vision Technologies*			
12:30 pm - 1:30 pm Lunch in the Exhibit Hall			*12:30 pm - 1:30 pm Lunch in the Exhibit Hall*			
1:30 pm - 2:00 pm T1W06 ◈ Practical Strategies for Successful Implementation and Deployment of AI-Based Solutions *Ritesh Agarwal, Computer Vision Lead, Globus Medical*	**1:30 pm - 2:00 pm T2W06** Using Synthetic Data to Train Computer Vision Models *Brian Geisel, CEO, Geisel Software*	**1:30 pm - 2:35 pm FW06** ◈ Fundamentals of Training AI Models for Computer Vision Applications *Amit Mate, Founder and CEO, GMAC Intelligence*	**1:30 pm - 2:00 pm BW06** Implementing AI/Computer Vision for Corporate Security Surveillance *Prasad Saranjame, Head of Physical Security and Resiliency, VMware*	**1:30 pm - 2:00 pm E1W06** Squeezing the Last Milliwatt and Cubic Millimeter from Smart Cameras Using the Latest FPGAs and DRAMs *Richard Crisp, Vice President and Chief Scientist, Etron Technology America* *Mark Hoopes, Director of Industrial & Automotive Segments, Lattice Semiconductor*	**1:30 pm - 2:00 pm E2W06** OpenCV for High-Performance, Low-Power Vision Applications on Snapdragon *Xin Zhong, Computer Vision Product Manager, Qualcomm Technologies*	**1:30 pm - 2:00 pm E3W06** Transforming Enterprise Intelligence: The Power of Computer Vision and Gen AI at the Edge with OpenVINO *Leila Sabeti, Americas AI Technical Sales Lead, Intel*
2:05 pm - 2:35 pm T1W07 Leveraging Neural Architecture Search for Efficient Computer Vision on the Edge *Hiram Rayo Torres Rodriguez, Senior AI Research Engineer, NXP Semiconductors*	**2:05 pm - 2:35 pm T2W07** ◈ Innovative Applications of Computer Vision for Power Utility Infrastructure Inspection *Vikhyat Chaudhry, Co-Founder, Chief Technology Officer and Chief Operating Officer, Buzz Solutions*		**2:05 pm - 2:35 pm BW07** ◈ Ten Commandments for Building a Vision AI Product *Vaibhav Ghadiok, Chief Technology Officer, Hayden AI*	**2:05 pm - 2:35 pm E1W07** Image Signal Processing Optimization for Object Detection (OB) *Young-Jun Yoo, Executive VP, Nextchip*	**2:05 pm - 2:35 pm E2W07** Implementing Transformer Neural Networks for Visual Perception on Embedded Devices *Shang-Hung Lin, Vice President of Neural Processing Products, VeriSilicon*	**2:05 pm - 2:35 pm E3W07** How Axelera AI Uses Digital Compute-in-Memory to Deliver Fast and Energy-Efficient Computer Vision *Bram Verhoef, Head of Machine Learning, Axelera AI*
2:40 pm - 3:10 pm T1W08 Testing Cloud-to-Edge Deep Learning Pipelines: Ensuring Robustness and Efficiency *Rustem Feyzkhanov, Staff Machine Learning Engineer, Instrumental*	**2:40 pm - 3:10 pm T2W08** Federated ML Architecture for Computer Vision in the IoT Edge *Akram Sheriff, Senior Manager, Software Engineering, Cisco*	**2:40 pm - 3:10 pm FW08** Transformer Networks: How They Work and Why They Matter *Rakshit Agrawal, Co-Founder and CEO, Ryddle AI*	**2:40 pm - 3:10 pm BW08** ◈ Why Amazon Failed and the Future of Computer Vision in Retail: An Interview with Will Glaser of Grabango *Will Glaser, Founder and CEO, Grabango* *Interviewer: Junko Yoshida, Editor-in-Chief, Ojo-Yoshida Report*	**2:40 pm - 3:10 pm E1W08** Building and Scaling AI Applications with the Nx AI Manager *Robin van Emden, Senior Director, Data Science, Network Optix*	**2:40 pm - 3:10 pm E2W08** Addressing Tomorrow's Sensor Fusion and Processing Needs with Cadence's Newest Processors *Amol Borkar, Product Marketing Director, Cadence*	**2:40 pm - 3:10 pm E3W08** A Cutting-Edge Memory Optimization Method for Embedded AI Accelerators *Arnaud Collard, Technical Leader, Embedded AI, 7 Sensing Software*
3:10 pm - 4:15 pm Break - Be sure to visit the Exhibit Hall!			*3:10 pm - 4:15 pm Break - Be sure to visit the Exhibit Hall!*			
4:15 pm - 5:20 pm T1W09 ◈ Unveiling the Power of Multimodal Large Language Models: Revolutionizing Perceptual AI *István Fehérvári, Chief Scientist, Ingram Technologies*	**4:15 pm - 4:45 pm T2W09** Advancing Embedded Vision Systems: Harnessing Hardware Acceleration and Open Standards *Neil Trevett, President, Khronos Group*	**4:15 pm - 4:45 pm FW09** ◈ DNN Quantization: Theory to Practice *Dwith Chenna, Senior Embedded DSP Engineer, Computer Vision, Magic Leap*	**4:15 pm - 4:45 pm BW09** ◈ Interview: Embedded Vision Opportunities and Challenges in Retail Checkout *Anatoly Kotlarsky, Distinguished Member, Technical Staff, R & D, Zebra Technologies* *Interviewer: Phil Lapsley, Co-Founder and Vice President, BDTI*	**4:15 pm - 4:45 pm E1W09** How to Run Audio and Vision AI Algorithms at Ultra-Low Power *Deepak Mital, Senior Director, Architectures, Synaptics*	**4:15 pm - 4:45 pm E2W09** The Importance of Memory for Breaking the Edge AI Performance Bottleneck *Wil Florentino, Sr. Segment Marketing Manager - Industrial/IIoT, Micron Technology*	**4:15 pm - 5:20 pm E3W09** Deploying Large Language Models on a Raspberry Pi *Pete Warden, CEO, Useful Sensors*
	4:50 pm - 5:20 pm T2W10 Adventures in Moving a Computer Vision Solution from Cloud to Edge *Nate D'Amico, CTO and Head of Product, MetaConsumer*	**4:50 pm - 5:55 pm FW10** ◈ Deep Neural Network Training: Diagnosing Problems and Implementing Solutions *Fahed Hassanat, COO and Head of Engineering, Sensor Cortek*	**4:50 pm - 5:20 pm BW10** ◈ Navigating Challenges and Seizing Opportunities: Scaling Computer Vision in Edge AI for Manufacturing *Rutger Vrijen, Partner, McKinsey & Company*	**4:50 pm - 5:20 pm E1W10** Maximize Your AI Compatibility with Flexible Pre- and Post-Processing *Jayson Bethurem, VP, Marketing and Business Development, Flex Logix*	**4:50 pm - 5:20 pm E2W10** Numerical and Accuracy Challenges of Implementing Quantized State-of-the-Art Transformer Networks in Embedded Devices *Aman Sikka, Chief Architect, Quadric*	
5:25 pm - 5:55 pm T1W11 ◈ How Large Language Models Are Impacting Computer Vision *Jacob Marks, Senior ML Engineer and Researcher, Voxel51*			**5:25 pm - 5:55 pm BW11** ◈ Omnilert Gun Detect: Harnessing Computer Vision to Tackle Gun Violence *Chad Green, Director of Artificial Intelligence, Omnilert*	**6:00 pm - 6:30 pm E1W12** Special Interest Group: Generative AI at the Edge *Phil Lapsley, Vice President, Edge AI and Vision Alliance* *Kerry Shih, Founder, GenAI Nerds*		April 22, 2024 Based on R05. Contents preliminary and subject to change.

Don't miss the Technology Exhibits Reception! 6:00 pm - 7:30 pm for food, drink and demos in the Exhibit Hall!

Don't miss the Technology Exhibits Reception 6:00 pm - 7:30 pm for food, drink and demos in the Exhibit Hall!

= Invited presentation

Session codes (e.g., T1R01) allow for quick searching in our event app!

Technical Insights 1 Mission City Ballroom—B1-B5	Technical Insights 2 Mission City Ballroom—M1-M3	Fundamentals Great America Ballroom	Business Insights Theater (Upstairs)	Enabling Technologies 1 Exhibit Hall—ET-1	EnablingTechnologies 2 Exhibit Hall—ET-2
11:25 am - 11:55 am T1R04 **Enabling Smart Retail with Visual AI** *Himanshu Vajaria, Engineering Manager, 365 Retail Markets*	**11:25 am - 11:55 am** T2R04 **Harm and Bias Evaluation and Solution for Adobe Firefly** *Rebecca Li, Machine Learning Engineering Manager, Adobe*	**11:25 am - 12:30 pm** FR04 **Multiple Object Tracking Systems** *Javier Berneche, Senior Machine Learning Engineer, Tryolabs*	**11:25 am - 11:55 am** BR04 **Future Radar Technologies and Applications** *James Jeffs, Senior Technology Analyst, IDTechEx*	**11:25 am - 11:55 am** E1R04 **Optimized Vision Language Models for Intelligent Transportation System Applications** *Tae-Ho Kim, Co-Founder and CTO, Nota AI*	**11:25 am - 11:55 am** E2R04 **How Arm's Machine Learning Solution Enables Vision Transformers at the Edge** *Stephen Su, Senior Segment Marketing Manager, Arm*
12:00 pm - 12:30 pm T1R05 **Using Vision Systems, Generative Models and Reinforcement Learning for Sports Analytics** *Mehrsan Javan, Chief Technology Officer, Sportlogiq*	**12:00 pm - 12:30 pm** T2R05 **Identifying and Mitigating Bias in AI** *Nikita Tiwari, AI Enabling Engineer, OEM PC Experiences, Client Computing Group, Intel Corporation*		**12:00 pm - 12:30 pm** BR05 **Interview: Latest Trends in AI Semiconductors** *Jay Goldberg, CEO and Founder, D2D Advisory* *Interviewer: Phil Lapsley, Co-Founder and Vice President, BDTI*	**12:00 pm - 12:30 pm** E1R05 **Efficiency Unleashed: The Next-Gen NXP i.MX 95 Applications Processor for Embedded Vision** *James Prior, Senior Product Manager, NXP Semiconductors*	**12:00 pm - 12:30 pm** E2R05 **Nx EVOS: A New Enterprise Operating System for Video and Visual AI** *Nathan Wheeler, Co-Founder and CEO, Network Optix*
		12:30 pm - 1:30 pm Lunch in the Exhibit Hall		**12:30 pm - 1:30 pm Lunch in the Exhibit Hall**	
1:30 pm - 2:00 pm T1R06 **Cost-Efficient, High-Quality AI for Consumer-Grade Smart Home Cameras** *Lin Chen, Chief Scientist, Wyze*	**1:30 pm - 2:00 pm** T2R06 **Interview: Exploring MIPI Camera Interface Standards for Embedded Vision Applications** *Haran Thanigasalam, Camera and Imaging Consultant, MIPI Alliance* *Interviewer: Shung Chieh, Senior Vice President, Eikon Systems, Eikon Therapeutics*	**1:30 pm - 2:00 pm** FR06 **Seeing Through Machines: A Guide to Image Sensors for Edge AI Applications** *Armita Abadian, Advisor, SEEdar Consulting*	**1:30 pm - 2:00 pm** BR06 **Entering the Era of Multimodal Perception** *Simon Morris, Serial Tech Entrepreneur and Start-Up Advisor, Connected Vision Advisors*	**1:30 pm - 2:00 pm** E1R06 **Intel's Approach to Operationalizing AI in the Manufacturing Sector** *Tara Thimmanaik, AI Systems and Solutions Architect, Intel Corporation*	**1:30 pm - 2:00 pm** E2R06 **Deploying Large Models on the Edge: Success Stories and Challenges** *Vinesh Sukumar, Senior Director, Product Management, Qualcomm Technologies*
2:05 pm - 2:35 pm T1R07 **Smarter AI for Detecting Microscopic Material Defects** *Shradha Agarwal, Research Scientist, Oak Ridge National Laboratory and the University of Tennessee*	**2:05 pm - 3:10 pm** T2R07 **Bridging Vision and Language: Designing, Training and Deploying Multimodal Large Language Models** *Adel Ahmadyan, Staff Engineer, Meta Reality Labs*	**2:05 pm - 2:35 pm** FR07 **Introduction to Cameras for Embedded Applications** *Brian Rodricks, CTO, SensorSpace*	**2:05 pm - 2:35 pm** BR07 **Market and Technology Trends in Automotive ADAS** *Florian Domengie, Senior Technology and Market Analyst, Yole Group*	**2:05 pm - 2:35 pm** E1R07 **Thriving on the Edge with Efficient Neuromorphic Processors** *Nandan Nayampally, Chief Marketing Officer, BrainChip*	**2:05 pm - 2:35 pm** E2R07 **Build a Tiny Vision Application in Minutes with the Edge App SDK** *Dan Mihai Dumitriu, Chief Technology Officer, Midokura, a Sony Group company*
2:40 pm - 3:10 pm T1R08 **Better Farming through Embedded AI** *Chris Padwick, Director, Computer Vision Machine Learning, Blue River Technology*	**2:40 pm - 3:10 pm** T2R08 **Removing Weather-Related Image Degradation at the Edge** *Ramit Pahwa, Machine Learning Scientist, Rivian*	**2:40 pm - 3:10 pm** FR08 **Introduction to Depth Sensing** *Harish Venkataraman, Depth Cameras Architecture and Tech Lead, Meta*	**2:40 pm - 3:10 pm** BR08 **Recent Trends in Industrial Machine Vision: Challenging Times** *Axel Clouet, Technology and Market Analyst, Imaging, Yole Group*	**2:40 pm - 3:10 pm** E1R08 **Silicon Slip-Ups: The Ten Most Common Errors Processor Suppliers Make (Number Four Will Amaze You!)** *Phil Lapsley, Co-Founder and Vice President, BDTI*	**2:40 pm - 3:10 pm** E2R08 **Challenges and Solutions of Moving Vision LLMs to the Edge** *Costas Calamvokis, Distinguished Engineer, Expedera*
		3:10 pm - 4:15 pm Break - Be sure to visit the Exhibit Hall!		**3:10 pm - 4:15 pm Break - Be sure to visit the Exhibit Hall!**	
4:15 pm - 4:45 pm T1R09 **Edge AI Optimization on Rails—Literally** *Matthew Pietrzykowski, Principal Data Scientist, Wabtec*	**4:15 pm - 4:45 pm** T2R09 **Seeing the Invisible: Unveiling Hidden Details through Advanced Image Acquisition Techniques** *Raghava Kashyapa, CEO, Qualitas Technologies*	**4:15 pm - 4:45 pm** FR09 **Introduction to Modern Radar for Machine Perception** *Robert Laganière, Professor, University of Ottawa and CEO, Sensor Cortek*			
4:50 pm - 5:20 pm T1R10 **Real-Time Retail Product Classification on Android Devices inside the Caper AI Cart** *David Scott, Senior Machine Learning Engineer, Instacart*	**4:50 pm - 5:20 pm** T2R10 **Using MIPI CSI to Interface with Multiple Cameras** *Karthick Kumaran, Staff Software Engineer, Meta*	**4:50 pm - 5:20 pm** FR10 **Introduction to Visual Simultaneous Localization and Mapping (VSLAM)** *Amol Borkar, Product Marketing Director, Cadence* *Shrinivas Gadkari, Design Engineering Group Director, Cadence*	**4:50 pm - 5:55 pm** BR10 **Vision Tank Start-Up Competition** *Active Insights . . . . William West, Founder and CEO* *Edgehog Technologies . . . Nasim Sahraei, Chief Product Officer,* *EyePop.ai . . . . Brad Chisum, CEO* *OpenMV . . . . .Kwabena Agyeman, President and CEO* *Waveye . . . . Gor Hakobyan, CTO*		
5:25 pm - 5:55 pm T1R11 **Improved Navigation Assistance for the Blind via Real-Time Edge AI** *Aishwarya Jadhav, Software Engineer, Autopilot AI Team, Tesla*	**5:25 pm - 5:55 pm** T2R11 **Improved Data Sampling Techniques for Training Neural Networks** *Karthik Rao Aroor, AI Engineer, Independent*	**5:25 pm - 5:55 pm** FR11 **Introduction to Semantic Segmentation** *Sébastien Taylor, Vice President of Research and Development, Au-Zone Technologies*	**Vision Tank Judges:** *John Feland, Forrest Iandola, Vin Ratford, Shweta Shrivastava*		

April 22, 2024 Based on R05. Contents preliminary and subject to change.