

# Agenda

## Summit

Monday, May 11

7:30 am - 7:00 pm  
Registration

7:30 am - 9:00 am  
Coffee and Pastries

9:00 am - 11:10 am

**Introduction, Awards, Keynote and Panel**  
Mission City Ballroom—B1-B5

### INTRODUCTION

**A View from the Summit (Part 1)**

*Jeff Bier, Founder, Edge AI and Vision Alliance and President, BDTI*

**PRODUCT OF THE YEAR AWARDS**  
**AI INNOVATION AWARDS**

### KEYNOTE

**AGI Meets the Real World—Toward Reasoning, Planning and Acting Beyond Book Intelligence**

*Eric Xing, President, Mohamed bin Zayed University of Artificial Intelligence and Professor, Carnegie Mellon University*

### PANEL

**Edge AI and Vision in Robotics: From Benchmarks to Fleet-Scale Reality**

*Moderator: Dave Tokic, Vice President of Corporate Development, Torc Robotics*

*Panelists: Vlad Branzoi, Perception Sensors Team Lead and Senior Staff Engineer, Agility Robotics*

*Bob Kunz, Chief Architect, Ambarella*

*Rajan Mistry, Senior Staff Engineer—Developer Advocacy, Qualcomm Technologies, Inc.*

*Durgesh Tiwari, VP of Hardware Systems, R & D, Simbe Robotics*

11:10 am - 11:25 am Break

11:25 am - 12:30 pm  
**Technical Insights 1**  
Fundamentals

**Morning Sessions**  
**Technical Insights 2**  
Business Insights

12:30 pm - 1:30 pm Lunch Exhibit Hall

1:30 pm - 3:10 pm  
**Technical Insights 1**  
Fundamentals  
Enabling Technologies 1  
Enabling Technologies 2  
Enabling Technologies 3

**Afternoon Sessions**  
**Technical Insights 2**  
Business Insights

3:10 pm - 4:15 pm Break Exhibit Hall

4:15 pm - 5:55 pm  
**Technical Insights 1**  
Fundamentals  
Enabling Technologies 1  
Enabling Technologies 2  
Enabling Technologies 3

**Afternoon Sessions**  
**Technical Insights 2**  
Business Insights

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

Tuesday, May 12

7:30 am - 5:00 pm  
Registration

7:30 am - 9:00 am  
Coffee and Pastries

9:00 am - 11:10 am

**Introduction, Keynote and Panel**  
Mission City Ballroom—B1-B5

### INTRODUCTION

**A View from the Summit (Part 2)**

*Jeff Bier, Founder, Edge AI and Vision Alliance and President, BDTI*

### KEYNOTE

**Scaling Down Is the New Scaling Up**

*Vikas Chandra, Senior Director, Meta Reality Labs*

### PANEL

**Vision-Language Models in the Real World: What Ships, What Breaks, What's Next**

*Moderator: Phil Lapsley, Co-Founder and Vice President, BDTI and Vice President of Business Development, Edge AI and Vision Alliance*

*Panelists: Rakshit Agrawal, Principal Applied Scientist, Microsoft*

*Vaibhav Ghadiok, Co-Founder and CTO, Hayden AI*

*David Selinger, CEO, Deep Sentinel*

*Tushar Sheth, Co-Founder and CEO, Superfocus.ai*

11:10 am - 11:25 am Break

11:25 am - 12:30 pm  
**Technical Insights 1**  
Fundamentals  
Enabling Technologies 1  
Enabling Technologies 2

**Morning Sessions**  
**Technical Insights 2**  
Business Insights

12:30 pm - 1:30 pm Lunch Exhibit Hall

1:30 pm - 3:10 pm  
**Technical Insights 1**  
Fundamentals  
Enabling Technologies 1  
Enabling Technologies 2

**Afternoon Sessions**  
**Technical Insights 2**  
Business Insights

3:10 pm - 4:15 pm Break Exhibit Hall

4:15 pm - 5:55 pm  
**Technical Insights 1**  
Fundamentals

**Afternoon Sessions**  
**Technical Insights 2**  
Business Insights

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

## Trainings and Edge AI Deep Dive

Wednesday, May 13

Separate registration and fee required for each session.

### Edge AI Deep Dive™

Qualcomm

**Compiling and Optimizing AI Models for Deployment on Qualcomm Hexagon NPUs with Qualcomm AI Hub**

In this hands-on workshop you'll learn how to optimize, compile and test AI models for deployment on Qualcomm Hexagon NPUs using Qualcomm AI Hub—preserving accuracy while staying within memory constraints. You will work through an end-to-end workflow spanning model selection, quantization, compilation, runtime choice, profiling and application integration. You'll leave with a working app, performance data and reusable APIs for scalable on-device AI deployment.

9:00 am - 12:00 pm

SEMI Innovation Center • 673 S Milpitas Blvd, Milpitas, CA

\$25

### Vision-Language Model (VLM) Trainings



VLMs offer a promising path toward image and video understanding, and the possibility of building agentic physical AI systems based on them. But they're new, complex, fast changing and resource-hungry. These training courses focus on the latest techniques in VLMs, their integration with traditional computer vision methods and their use in video understanding and agentic applications.

9:00 am - 12:00 pm

**Introductory Course: Vision-Language Models for Computer Vision Applications: A Hands-On Introduction**

1:30 pm - 5:00 pm

**Advanced Course: Vision-Language Models for Video Understanding and Agentic AI**

Cadence Design Systems, Bldg. 5 • 2655 Seely Ave, San Jose, CA

\$495 per session



**There's more to see in the Exhibit Hall!**

**Visit the Speaker Square in the Exhibit Hall**

Chat one-on-one with some of our amazing speakers in the Speaker Square. Speaker schedule located in the Summit mobile app, on the Event Guide Addendum and in the Speaker Square.

Following Jeff Bier's opening remarks, "A View from the Summit (Part 1)" (9:00 am - 9:35 am on Monday), join us for our Keynote and Panel!



9:35 am - 10:25 am

## Monday

### AGI Meets the Real World—Toward Reasoning, Planning and Acting Beyond Book Intelligence

**Eric Xing**

*President, Mohamed bin Zayed University of Artificial Intelligence and Professor, Carnegie Mellon University*

LLMs have shown astounding capabilities on a wide range of tasks, such as genius-level standard test performance, Olympiad-level math reasoning and human-like conversation ability. However, LLMs are only "book intelligent"; they suffer from limitations for embodied, physical and social reasoning—and for strategic planning—in the real world.

In this talk, we'll explain how to build a true **world model**, rather than a video generator, to simulate actionable possibilities of the real world for reasoning and planning via thought experiment rather than pixel realism. We'll also explain how to build a true **agent model**, rather than an LLM wrapper or software pipeline, to enable learning and acting with flexibility, adaptability and autonomy, and with abilities of self-regulation, reflection, collaboration and socialization. We propose a Generative Latent Prediction architecture for world modeling that builds on stateful latent space, long-horizon and closed-loop action-conditioned latent reasoning and learning/inferencing grounded over realizable world states. In addition, we propose a Goal-Identity-Configurator architecture for agent modeling that can regulate reasoning modes between unconsciously reactive and consciously deliberative, generate real-world actions based on its goals and identity and self-learn off-line from the world model.

Finally, we will present PAN, a physical, agentive and nested framework over the proposed architectures that brings together perception, state, action and causality within one system to support

open-domain, interactable world simulation and agentive intelligence. PAN achieves strong performance in action-conditioned world simulation, long-horizon forecasting and simulative reasoning compared to other approaches.

#### About Eric Xing

Professor Eric Xing is the President of the Mohamed bin Zayed University of Artificial Intelligence and a Professor of Computer Science at Carnegie Mellon University. His main research interests are in the development of machine learning and statistical methodology, as well as large-scale distributed computational systems and architectures, for solving problems involving automated learning, reasoning and decision-making in artificial, biological and social systems. In recent years, he has been focused on building large language models, world models, agent models and foundation models for biology.

Professor Xing completed his undergraduate study at Tsinghua University and holds a PhD in Molecular Biology and Biochemistry from the State University of New Jersey and a PhD in Computer Science from the University of California, Berkeley.

Professor Xing has served on the editorial boards of several leading journals, including *JASA*, *AOAS* and *JMLR*; is a recipient of several awards, including NSF CAREER, Sloan and Carnegie Science, and Best Paper in conferences such as ACL, ISMB, NeurIPS and OSDI; and is a fellow of several societies, including AAAI, ACM, ASA, IEEE and IMS.

10:25 am - 11:10 am

## Monday

### Edge AI and Vision in Robotics: From Benchmarks to Fleet-Scale Reality

**Moderator:**

**Dave Tokic**

*Vice President of Corporate Development, Torc Robotics*

**Panelists:**

**Vlad Branzoi**

*Perception Sensors Team Lead and Senior Staff Engineer, Agility Robotics*

**Bob Kunz**

*Chief Architect, Ambarella*

**Rajan Mistry**

*Senior Staff Engineer—Developer Advocacy, Qualcomm Technologies, Inc.*

**Durgesh Tiwari**

*VP of Hardware Systems, R & D, Simbe Robotics*

Edge AI is helping robots see, decide and act in the real world—but the hardest work starts after the demo. In this plenary panel, we'll unpack where edge AI is creating measurable value today, then dive into the system choices that determine whether a robot ships: vision-only versus multimodal sensing and fusion, the real compute bottlenecks (latency, bandwidth, power, thermals, memory and software) and the trade-offs between modular pipelines and end-to-end learned stacks. We'll also discuss the data problem—collection, labeling, simulation and continuous improvement—plus the practical role of foundation and vision-language models in embodied systems. Finally, we'll cover safety and trust around people, why pilots fail to scale, what changes from 10 robots to 1,000, and what breakthroughs are most likely to matter over the next three to five years.

Following Jeff Bier's opening remarks, "A View from the Summit (Part 2)" (9:00 am - 9:35 am on Tuesday), join us for our Keynote and Panel!



9:35 am - 10:25 am

## Tuesday

### Scaling Down Is the New Scaling Up

**Vikas Chandra**  
Senior Director, Meta Reality Labs

The AI industry spent the last decade scaling up. The next decade will be about scaling down: AI that runs on your device, reasons across what you see and hear and understands by utilizing context that never leaves your pocket.

Drawing on our recent work, I'll show how to make this shift algorithmically. Extreme quantization down to sub-2-bit precision doesn't just shrink models; it forces them into entirely new representations. KV cache compression and speculative decoding make interactive latency possible within mobile power budgets. Lightweight vision models bring real-time understanding to cameras. And temporal compression lets models process hours of video on constrained hardware.

These techniques unlock new experiences. Sub-billion parameter reasoning models will power private contextual assistants that understand your routines and daily trade-offs without anything leaving your device. Efficient vision encoders will enable real-time segmentation and object tracking for live video effects, and metric depth estimation on glasses for AR and spatial understanding. Temporal compression will make on-device video search over personal recordings practical.

The AI of choice won't be the biggest model. It will be the one that's always there.

#### About Vikas Chandra

Vikas Chandra is a Senior Director at Meta Reality Labs, where he leads an AI research team building efficient on-device AI for AI glasses and other mixed-reality products. These devices perceive the world as the wearer does, using context to anticipate needs and take action, laying the foundation for the next generation of human-device interaction.

Prior to joining Meta in 2018, Dr. Chandra was Director of Applied Machine Learning at Arm Research, where his team helped pioneer tinyML, a new field of machine learning that enables AI to run on small, resource-constrained devices.

Dr. Chandra holds an MS and PhD in Electrical and Computer Engineering from Carnegie Mellon University. He has also served as Visiting Scholar (2011-2014) and Visiting Faculty (2016-2017) in the Electrical Engineering Department at Stanford University. He has authored more than 200 research papers and is an inventor on over 40 US and international patents. In recognition of his technical leadership, Dr. Chandra received the ACM/SIGDA Technical Leadership Award in 2009 and was invited to the National Academy of Engineering's Frontiers of Engineering Symposium in 2017. In 2026, he received the Outstanding Alumni Award from Carnegie Mellon University's College of Engineering.

10:25 am - 11:10 am

## Tuesday

### Vision-Language Models in the Real World: What Ships, What Breaks, What's Next

#### Moderator:

**Phil Lapsley**  
Co-Founder and Vice President, BDTI and Vice President of Business Development, Edge AI and Vision Alliance

#### Panelists:

**Rakshit Agrawal**  
Principal Applied Scientist, Microsoft

**Vaibhav Ghadiok**  
Co-Founder and CTO, Hayden AI

**David Selinger**  
CEO, Deep Sentinel

**Tushar Sheth**  
Co-Founder and CEO, Superfocus.ai

▶ Vision-language models are moving fast, but it's not always clear where they add value, and many teams struggle to turn demos into dependable products. In this plenary panel, we'll cut through the hype and focus on where VLMs make sense and the barriers to deploying them in real systems. We'll discuss where VLMs are delivering clear value and where classic computer vision still wins on cost, latency and determinism. Panelists will compare practical hybrid architectures, examine failure modes such as weak grounding and hallucination and outline guardrails, evaluation methods and monitoring that work in production. We'll also debate the edge vs. cloud split, domain adaptation strategies and the privacy/security governance required when models can answer open-ended questions about people and places. Attendees will leave with insights into requirements and potential pitfalls to de-risk their VLM road map.



















# Monday Sessions Overview

Monday

Read the full session abstracts on pages 32-42

Session slides will be available about a week after the Summit. You will receive an email when they are available.

 = Invited presentation

Technical Insights 1 Mission City Ballroom—B1-B5	Technical Insights 2 Mission City Ballroom—M1-M3	Fundamentals Room 203-204 (Upstairs)	Business Insights Room 209-210 (Upstairs)	Enabling Technologies 1 Exhibit Hall—ET-1	Enabling Technologies 2 Exhibit Hall—ET-2	Enabling Technologies 3 Exhibit Hall—ET-3
<p><b>11:25 am - 11:55 am</b> <b>TIM04</b>   <b>From CNNs to LLMs at the Edge: Why Memory, Not TOPS, Is the Bottleneck</b>  <i>David Selinger, CEO, Deep Sentinel</i></p> <p><b>12:00 pm - 12:30 pm</b> <b>TIM05</b>  <b>Decomposing AI Pipelines for Robust Performance at the Edge</b>  <i>Pullarao Maddu, Embedded AI Architect, Valeo</i></p>	<p><b>11:25 am - 11:55 am</b> <b>T2M04</b>  <b>Efficient State-Space Models</b>  <i>Raja Giryes, Professor, Tel Aviv University and Research Scientist, Apple</i></p> <p><b>12:00 pm - 12:30 pm</b> <b>T2M05</b>  <b>Exploring Radar SLAM: Advancing Localization and Mapping for Automotive, Robotics and Beyond</b>  <i>Amit Kumar, Director of Product Management and Marketing, Cadence</i>  <i>Amit Sulakhe, Director in the Vision Group, Cadence Pune</i></p>	<p><b>11:25 am - 12:30 pm</b> <b>FM04</b>   <b>Introduction to Deep Learning and Visual AI: Fundamentals, Architectures and Real-World Trade-Offs</b>  <i>Gowdhaman Sadhasivam, Chief Technology Officer, Labelbees AI</i></p>	<p><b>11:25 am - 11:55 am</b> <b>BM04</b>   <b>Revolutionizing Physical Security with Agentic AI</b>  <i>Steve Lindsey, Co-Founder and Chief Strategist, LVT</i></p> <p><b>12:00 pm - 12:30 pm</b> <b>BM05</b>  <b>Reality Check: Vision and AI Developer Trends and Pain Points and Implications</b>  <i>Phil Lapsley, Co-Founder and Vice President, BDTI and Vice President of Business Development, Edge AI and Vision Alliance</i></p>			
<b>12:30 pm - 1:30 pm Lunch in the Exhibit Hall</b>			<b>12:30 pm - 1:30 pm Lunch in the Exhibit Hall</b>			
<p><b>1:30 pm - 2:00 pm</b> <b>TIM06</b>   <b>Edge AI Processor SDKs: A Reality Check</b>  <i>Francesco Mattioli, Partnership and Ecosystem Manager, Ultralytics</i></p> <p><b>2:05 pm - 2:35 pm</b> <b>TIM07</b>   <b>Deploying VLMs at the Edge: Model Selection, Memory Reality and System Design</b>  <i>Patrick Farry, Founder and Architect, Intelligence at the Edge</i></p> <p><b>2:40 pm - 3:10 pm</b> <b>TIM08</b>   <b>Can Cheap Chips Run Multimodal AI? Lessons from Testing 17 SoCs</b>  <i>Tushar Sheth, Co-Founder and CEO, Superfocus.ai</i></p>	<p><b>1:30 pm - 2:00 pm</b> <b>T2M06</b>  <b>Self-Compression for Edge Inference</b>  <i>James Imber, Director of Research, Imagination Technologies</i></p> <p><b>2:05 pm - 2:35 pm</b> <b>T2M07</b>   <b>Embedded Vision for Real-Time Gravity Infusion Monitoring and Control</b>  <i>Dirk van der Merwe, Autonomous Robotics Lead, DEKA Research &amp; Development</i></p> <p><b>2:40 pm - 3:10 pm</b> <b>T2M08</b>  <b>Smaller Beats Bigger: Developing a Custom VLM for Marketing Image Evaluation</b>  <i>Shradha Agrawal, ML Lead, Adobe</i></p>	<p><b>1:30 pm - 2:35 pm</b> <b>FM06</b>   <b>Introduction to DNN Training: Fundamentals, Process and Best Practices</b>  <i>Kevin Weekly, CEO, Think Circuits</i></p> <p><b>2:40 pm - 3:10 pm</b> <b>FM08</b>   <b>Object Detection Models: Balancing Speed, Accuracy and Efficiency</b>  <i>Sage Elliott, AI Engineer, Union.ai</i></p>	<p><b>1:30 pm - 2:00 pm</b> <b>BM06</b>   <b>Ear-Worn Edge AI Devices: Multifunctional Hearing Aids with Embedded Deep Neural Networks</b>  <i>Achin Bhowmik, Chief Technology Officer and Executive Vice President of Engineering, Starkey</i></p> <p><b>2:05 pm - 2:35 pm</b> <b>BM07</b>  <b>The Rise of the AI-Defined Enterprise: From Autonomous Vehicles to Robots and the Future of Intelligent Machines</b>  <i>Andy Xiao, Staff AI Engineer, Rivian and Volkswagen Group Technologies and Adjunct Professor, San Jose State University</i></p> <p><b>2:40 pm - 3:10 pm</b> <b>BM08</b>   <b>ElliQ: Redefining Healthy Aging through Proactive, Human-Centric AI</b>  <i>Assaf Gad, Chief Strategy Officer and General Manager, Intuition Robotics</i></p>	<p><b>1:30 pm - 2:00 pm</b> <b>E1M06</b>  <b>AMD Vitis AI Workflow: Compilation, Hardware Deployment and Profiling</b>  <i>Thomas Zerbs, Technical Marketing Engineer, Adaptive and Embedded Computing Group, AMD</i></p> <p><b>2:05 pm - 2:35 pm</b> <b>E1M07</b>  <b>Invertible Light Technology: A Paradigm Shift for Depth Sensing</b>  <i>Takeo Miyazawa, Founder and CEO, MagikEye</i></p> <p><b>2:40 pm - 3:10 pm</b> <b>E1M08</b>  <b>One Silicon, Two Worlds: NPU Optimization for Autoregressive and Diffusion Transformers</b>  <i>Shang-Hung Lin, Vice President of NPU Technology, VeriSilicon</i></p>	<p><b>1:30 pm - 2:00 pm</b> <b>E2M06</b>  <b>From Generative AI to Physical AI: Enabling VLM/VLA Multimodal Performance on Enhanced NPX6 NPU IP</b>  <i>Gordon Cooper, Principal Product Manager, Synopsys</i></p> <p><b>2:05 pm - 2:35 pm</b> <b>E2M07</b>  <b>HiFi iQ: Enabling Voice AI and Immersive Audio for Smart Home, Mobile and Automotive</b>  <i>Amol Borkar, Group Director of Product Management and Marketing—Tensilica DSPs, Silicon Solutions Group (SSG), Cadence</i></p> <p><b>2:40 pm - 3:10 pm</b> <b>E2M08</b>  <b>Roads to Robots: How Generative AI Is Redefining Memory and Storage for Embodied AI</b>  <i>Saideep Tiku, Principal System Architect, Micron</i></p>	<p><b>1:30 pm - 2:00 pm</b> <b>E3M06</b>  <b>Enabling "GenAI Everywhere": Flexible Model Compatibility and Insight-Driven Workflows</b>  <i>Tae-Ho Kim, Co-Founder and CTO, Nota AI</i></p> <p><b>2:05 pm - 2:35 pm</b> <b>E3M07</b>  <b>Why Your Next AI Accelerator Should Be an FPGA</b>  <i>Mark Oliver, VP of Marketing and Business Development, Efinix</i></p> <p><b>2:40 pm - 3:10 pm</b> <b>E3M08</b>  <b>Democratizing Physical AI: Arduino's Open Door to Qualcomm's Platform</b>  <i>Olivier Bloch, Director of Developer Relations, Qualcomm Technologies, Inc.</i></p>
<b>3:10 pm - 4:15 pm Break - Be sure to visit the Exhibit Hall!</b>			<b>3:10 pm - 4:15 pm Break - Be sure to visit the Exhibit Hall!</b>			
<p><b>4:15 pm - 4:45 pm</b> <b>TIM09</b>   <b>Document Intelligence with Vision-Language Augmentation</b>  <i>Sanjay Nichani, VP of AI and Computer Vision, ABBYY</i></p> <p><b>4:50 pm - 5:20 pm</b> <b>TIM10</b>   <b>Shrinking Vision-Language Models for Edge Deployment via Distillation and Pruning</b>  <i>Denis Gudovskiy, Distinguished AI Research Engineer, Panasonic AI Lab</i></p>	<p><b>4:15 pm - 5:20 pm</b> <b>FM09</b>   <b>Understanding Transformers: From LLMs to Context-Aware Multimodal Models</b>  <i>Tom Michiels, System Architect, Synopsys</i></p> <p><b>4:50 pm - 5:20 pm</b> <b>T2M10</b>   <b>From iPhone Scan to Structured Building Models: Lessons from 150M Square Feet</b>    <i>Anton Yakubenko, Head of AI and Computer Vision Strategy, Twindo</i></p>	<p><b>5:25 pm - 5:55 pm</b> <b>FM11</b>  <b>Training Models Starts with Data: Collection, Curation and the Refinement Loop</b>  <i>Umabharathi Govindarajan, Senior Machine Learning Engineer, Blue River Technology</i></p>	<p><b>4:15 pm - 4:45 pm</b> <b>BM09</b>   <b>What Kills Industrial Vision AI Pilots? (Hint: It's Not Model Performance)</b>  <i>Kasper De Smaele, Founder and CEO, Kasqade</i></p> <p><b>4:50 pm - 5:20 pm</b> <b>BM10</b>   <b>The Robotaxi Market Ramps Up</b>  <i>Greg Basich, Associate Director, Counterpoint Research</i></p> <p><b>5:25 pm - 5:55 pm</b> <b>BM11</b>  <b>Building a Virtual Product Placement System: A Case Study</b>  <i>Yash Chaturvedi, Head of Product, Generative Media, Amazon Prime Video</i></p>	<p><b>4:15 pm - 4:45 pm</b> <b>E1M09</b>  <b>Efficient Computer Vision at the Far Edge: Design and Training Under Constraints</b>  <i>Nicolas Widynski, Head of Lattice AI Lab, Lattice Semiconductor</i></p> <p><b>4:50 pm - 5:20 pm</b> <b>E1M10</b>  <b>From Compute-Bound to Memory-Bound: Edge AI Architectures for VLMs</b>  <i>Athish Rahul Rao, Staff Software Engineer, Expedera</i></p>	<p><b>4:15 pm - 4:45 pm</b> <b>E2M09</b>  <b>COOL: Accelerating Computer Vision Workloads with Cloud Optimized OpenCV on AWS</b>  <i>Frantz Lohier, Senior WW Specialist—Advanced Computing, AI and Robotics, Amazon Web Services</i>  <i>Satya Mallick, CEO, OpenCV.org</i></p> <p><b>4:50 pm - 5:20 pm</b> <b>E2M10</b>  <b>Speeding Time to Market with Production-Ready Edge AI Solutions: From Wake Word Detection to Face Recognition</b>  <i>Nick De Rosa, Edge AI Marketing Manager, Microchip Technology</i>  <i>Kannan Srinivasagam, CEO, WG Tech Solutions</i></p> <p><b>5:25 pm - 5:55 pm</b> <b>E2M11</b>  <b>Building a Local Voice Agent on a Raspberry Pi</b>  <i>Pete Warden, CEO, Moonshine AI</i></p>	<p><b>4:15 pm - 4:45 pm</b> <b>E3M09</b>  <b>From Prototype to Production: What Computer Vision Teams Wish They Knew at 100 Devices</b>  <i>Justin Schneck, Co-Founder and CTO, Peridio</i></p> <p><b>4:50 pm - 5:20 pm</b> <b>E3M10</b>  <b>Navigating Physical AI Deployment Across Multiple Platforms for Automated Optical Inspection</b>  <i>Barrie Mullins, Assistant Vice President, elfinchips (an Arrow company)</i></p>









































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!

Read the full session abstracts on pages 44-53

Session slides will be available about a week after the Summit. You will receive an email when they are available.

 = Invited presentation

Technical Insights 1 Mission City Ballroom—B1-B5	Technical Insights 2 Mission City Ballroom—M1-M3	Fundamentals Room 203-204 (Upstairs)	Business Insights Room 209-210 (Upstairs)	Enabling Technologies 1 Exhibit Hall—ET-1	Enabling Technologies 2 Exhibit Hall—ET-2
<p><b>11:25 am - 11:55 am</b> <span style="float:right">T1T04</span></p> <p> <b>Building Trustworthy Autonomous Driving Systems: Waymo's Holistic AI Safety Approach</b></p> <p><i>Chen Wu, Director, Head of Perception, Waymo</i></p>	<p><b>11:25 am - 11:55 am</b> <span style="float:right">T2T04</span></p> <p> <b>From YOLO to SAM: Segmentation Models on Real Edge Hardware</b></p> <p><i>Sébastien Taylor, VP of R &amp; D, Au-Zone Technologies</i></p>	<p><b>11:25 am - 12:30 pm</b> <span style="float:right">FT04</span></p> <p> <b>Vision-Language Models in Practice: Architecture and Performance</b></p> <p><i>Rajy Rawther, PMTS Software Architect, AMD</i></p>	<p><b>11:25 am - 11:55 am</b> <span style="float:right">BT04</span></p> <p> <b>Interview: Beyond TOPS—Evaluating and Selecting Processors for Edge AI and Vision</b></p> <p><i>Sally Ward-Foxton, Senior Reporter, EE Times</i></p> <p><i>Mark Jamtgaard, Director of Technology, RetailNext</i></p>	<p><b>11:25 am - 11:55 am</b> <span style="float:right">E1T04</span></p> <p> <b>Porting and Optimizing Advanced Vision-Language-Action Models for Embedded Autonomous Systems</b></p> <p><i>Mike Leonard, Software Architect, Quadric</i></p>	<p><b>11:25 am - 11:55 am</b> <span style="float:right">E2T04</span></p> <p> <b>Why Edge Vision Models Keep Breaking—and What Complete Training Data Changes</b></p> <p><i>David Scott, Founder and CEO, Syntec</i></p>
<p><b>12:00 pm - 12:30 pm</b> <span style="float:right">T1T05</span></p> <p> <b>From Decentralized ECUs to Zonal Compute: Emerging Vehicle Architectures for SDV and Edge AI</b></p> <p><i>Frank Moesle, System Architecture and Strategy Manager, Valeo</i></p>	<p><b>12:00 pm - 12:30 pm</b> <span style="float:right">T2T05</span></p> <p> <b>Challenges in Deploying Edge AI in the Defense Industry</b></p> <p><i>Roger Garcia, Senior Principal Engineer, Advanced Radar Products Department, Raytheon</i></p>		<p><b>12:00 pm - 12:30 pm</b> <span style="float:right">BT05</span></p> <p> <b>Automating Forklift Movement in Real-World Warehouses</b></p> <p><i>Brian Nachtigall, General Manager, ArcBest   Vaux</i></p>	<p><b>12:00 pm - 12:30 pm</b> <span style="float:right">E1T05</span></p> <p> <b>Always-On Edge Perception Via a Heterogeneous Near-Memory AI Architecture</b></p> <p><i>Petronel Bigioi, CEO, FotoNation</i></p>	<p><b>12:00 pm - 12:30 pm</b> <span style="float:right">E2T05</span></p> <p> <b>OpenCV 5 First Impressions: What Breaks, What Improves and What Matters</b></p> <p><i>Mark Antonelli, CTO, Boston.AI</i></p>
<b>12:30 pm - 1:30 pm Lunch in the Exhibit Hall</b>					
<p><b>1:30 pm - 2:00 pm</b> <span style="float:right">T1T06</span></p> <p> <b>Beyond "What's in the Image?": World Models for Understanding and Predicting</b></p> <p><i>Gowdhaman Sadhasivam, Chief Technology Officer, Labelbees AI</i></p>	<p><b>1:30 pm - 2:00 pm</b> <span style="float:right">T2T06</span></p> <p> <b>Designing Vision-Based, Battery-Powered Edge AI Sensors</b></p> <p><i>Larry Arne, Director of Hardware, VergeSense</i></p>	<p><b>1:30 pm - 2:35 pm</b> <span style="float:right">FT06</span></p> <p> <b>Specifying and Designing Cameras for Computer Vision Applications</b></p> <p><i>Richard Crisp, Vice President / CTO, Etron Technology America</i></p>	<p><b>1:30 pm - 2:00 pm</b> <span style="float:right">BT06</span></p> <p> <b>Vision AI as Catalyst for Accelerating Agricultural Discovery Research</b></p> <p><i>Sudhir Sornapudi, Principal Investigator, Corteva Agriscience</i></p>	<p><b>1:30 pm - 2:00 pm</b> <span style="float:right">E1T06</span></p> <p> <b>Accelerating Physical AI with ROCm: High-Performance ML on AMD Embedded iGPUs</b></p> <p><i>Alok Gupta, Senior Technical Marketing Manager, AMD</i></p>	<p><b>1:30 pm - 2:00 pm</b> <span style="float:right">E2T06</span></p> <p> <b>How Qualcomm Is Making Computer Vision Accessible Across Edge Verticals</b></p> <p><i>Derrick Chang, Senior Product Manager, Qualcomm Technologies, Inc.</i></p>
<p><b>2:05 pm - 2:35 pm</b> <span style="float:right">T1T07</span></p> <p> <b>Trains, Containers and Embedded Vision AI: How We Track 5M Containers at BNSF</b></p> <p><i>Michael Ibanez, Manager of Autonomous Systems, BNSF Railway</i></p>	<p><b>2:05 pm - 2:35 pm</b> <span style="float:right">T2T07</span></p> <p> <b>Small but Mighty: The Rise of Small Language Models and Their Role in the Future of AI</b></p> <p><i>Rakshit Agrawal, Principal Applied Scientist, Microsoft</i></p>		<p><b>2:05 pm - 2:35 pm</b> <span style="float:right">BT07</span></p> <p> <b>Closing the Visibility Gap: Digitizing the Store as the Foundation for Retail Transformation</b></p> <p><i>Kishor Taywade, CTO, Zippedi</i></p>	<p><b>2:05 pm - 2:35 pm</b> <span style="float:right">E1T07</span></p> <p> <b>How to Train an AI Model Using Roboflow and Deploy to an MCU with an NPU</b></p> <p><i>Kwabena Agyeman, President, OpenMV</i></p> <p><i>Joseph Nelson, Co-Founder and CEO, Roboflow</i></p>	<p><b>2:05 pm - 2:35 pm</b> <span style="float:right">E2T07</span></p> <p> <b>No RISC, No Reward: Unlocking Extreme Efficiency in Physical AI with RISC-V</b></p> <p><i>Mayank Mangla, AI Product Manager and Systems Architect, MIPS, a GlobalFoundries company</i></p>
<p><b>2:40 pm - 3:10 pm</b> <span style="float:right">T1T08</span></p> <p> <b>Pushing the Limits: Sensing and Compute Requirements for Humanoid Robots</b></p> <p><i>Vlad Branzoi, Perception Sensors Team Lead and Senior Staff Engineer, Agility Robotics</i></p>	<p><b>2:40 pm - 3:10 pm</b> <span style="float:right">T2T08</span></p> <p> <b>Small Language Models for Edge AI: Trade-Offs and Quantization in Practice</b></p> <p><i>Dwith Chenna, MTS Product Engineer, AI Inference, AMD</i></p>	<p><b>2:40 pm - 3:10 pm</b> <span style="float:right">FT08</span></p> <p> <b>Auto-Exposure, Auto-White-Balance, and Auto-Focus Algorithms for Embedded Camera Systems</b></p> <p><i>Yu-Chun Chen, Software Development Engineer, Amazon Lab126</i></p>	<p><b>2:40 pm - 3:10 pm</b> <span style="float:right">BT08</span></p> <p> <b>Turning Tech Differentiation into a Strategy for Dominance</b></p> <p><i>Simon Morris, Serial Tech Entrepreneur and Start-Up Advisor, Connected Vision Advisors</i></p>	<p><b>2:40 pm - 3:10 pm</b> <span style="float:right">E1T08</span></p> <p> <b>From Chips to Platforms: Scaling Edge AI with SoMs, Production Linux and Secure Life-Cycle Ops</b></p> <p><i>Amir Sherman, Head of Global Business Development, S-BD</i></p>	<p><b>2:40 pm - 3:10 pm</b> <span style="float:right">E2T08</span></p> <p> <b>Edge-First Coding Agents: Trustworthy Agentic Development for Real Devices</b></p> <p><i>Pietro Antonio Cicalese, Senior Technical Marketing Engineer, Ambarella</i></p>
<b>3:10 pm - 4:15 pm Break - Be sure to visit the Exhibit Hall!</b>					
<p><b>4:15 pm - 4:45 pm</b> <span style="float:right">T1T09</span></p> <p> <b>Industrial Vision Meets Physical AI: From Automated Vision to Human-Centric Agentic Systems</b></p> <p><i>Mahbubul Alam, Chief Research Scientist, Industrial AI Lab, Research &amp; Development, Hitachi America</i></p>	<p><b>4:15 pm - 4:45 pm</b> <span style="float:right">T2T09</span></p> <p> <b>Generative AI-Driven Image Dataset Curation for Industrial IoT at Scale</b></p> <p><i>Apurva Godghase, Lead Computer Vision Engineer, Brambles</i></p> <p><i>Nikhil Dange, Senior Computer Vision Engineer, Brambles</i></p>	<p><b>4:15 pm - 5:20 pm</b> <span style="float:right">FT09</span></p> <p> <b>Introduction to Depth Sensing: Technologies, Trade-Offs and Applications</b></p> <p><i>Chris Sarantos, Independent Consultant, Think Circuits</i></p>	<p><b>4:15 pm - 4:45 pm</b> <span style="float:right">BT09</span></p> <p> <b>Interview: Ben Pouladian of BEP Research on Emerging Opportunities in Edge AI and World Models</b></p> <p><i>Ben Pouladian, Founder and CEO, BEP Research and BEP Holdings</i></p> <p><i>Phil Lapsley, Co-Founder and Vice President, BDTI and Vice President of Business Development, Edge AI and Vision Alliance</i></p>		
<p><b>4:50 pm - 5:20 pm</b> <span style="float:right">T1T10</span></p> <p> <b>Deploying Embedded Vision in Healthcare: Lessons from Real-World Retina Screenings</b></p> <p><i>Sam Kavusi, Head of Retina Imaging, Google Life Sciences</i></p>	<p><b>4:50 pm - 5:20 pm</b> <span style="float:right">T2T10</span></p> <p> <b>The Immune System for Vision AI: Recovering from Data Poisoning in Critical Systems</b></p> <p><i>Dippu Kumar Singh, Leader of Emerging Technologies (App Services), Fujitsu North America</i></p>		<p><b>4:50 pm - 5:20 pm</b> <span style="float:right">BT10</span></p> <p> <b>From Clinical Problem to Regulated Product: Vision-Based Monitoring for Gravity Infusions</b></p> <p><i>Drew Blais, Engineering Program Manager, DEKA Research &amp; Development</i></p>		
<p><b>5:25 pm - 5:55 pm</b> <span style="float:right">T1T11</span></p> <p> <b>Multimodal AI Agents for Content Editing</b></p> <p><i>Yong Jae Lee, Professor, Department of Computer Sciences, University of Wisconsin-Madison and Research Scientist, Adobe Research</i></p>	<p><b>5:25 pm - 5:55 pm</b> <span style="float:right">T2T11</span></p> <p> <b>Vision AI in Retail: From Lab to 2,000+ Autonomous Checkout Points of Sale</b></p> <p> <i>José Benitez, Founder and Chief AI Officer, Intuitivo</i></p>	<p><b>5:25 pm - 5:55 pm</b> <span style="float:right">FT11</span></p> <p> <b>Introduction to Simultaneous Localization and Mapping: From Block Diagrams to Real Systems</b></p> <p><i>Amit Gupta, Associate Director—Solution Architecture and Head of Robotics CoE, elfinichips (an Arrow company)</i></p>	<p><b>5:25 pm - 5:55 pm</b> <span style="float:right">BT11</span></p> <p> <b>Interview: Robotaxis at the Crossroads—New AI Models, Sensor Fusion and the Road to Scalable AVs</b></p> <p><i>Greg Basich, Associate Director, Counterpoint Research</i></p> <p><i>Junko Yoshida, Independent Journalist</i></p>		