

MONDAY  
MAY 21

TensorFlow class requires separate registration

7:30 – 9:00 am  
REGISTRATION  
MAIN LOBBY

7:30 am – 5:00 pm  
SUMMIT REGISTRATION  
MAIN LOBBY

8:00 – 9:00 am  
COFFEE & PASTRIES  
GREAT AMERICA  
MEETING ROOMS 1–3

9:00 am – 5:00 pm  
Deep Learning for  
Computer Vision  
with TensorFlow

GREAT AMERICA  
MEETING ROOMS 1–3

The first day of the Embedded Vision Summit is dedicated to our new training class: Deep Learning for Computer Vision with TensorFlow. The single-day program will provide you with a hands-on overview of deep learning applications of TensorFlow.

This one-day training covers:

- Introduction to TensorFlow
- Neural Networks in TensorFlow
- Object Recognition in TensorFlow
- Training Data and Issues
- Open Source CNN Models

If you missed this opportunity at the Summit to take charge of your professional development, visit [tensorflow.embedded-vision.com](http://tensorflow.embedded-vision.com) to find out about upcoming classes.

FEE: \$895

Organized by 

TUESDAY MAY 22

7:30 am – 7:00 pm  
REGISTRATION

8:00 – 9:00 am  
COFFEE & PASTRIES

9:00 – 10:30 am  
WELCOME & KEYNOTE MISSION CITY BALLROOM B2 – 5

Think Like an Amateur, Do As an Expert:  
Lessons from a Career in Computer Vision

Dr. Takeo Kanade  
U.A. and Helen Whitaker Professor, Carnegie Mellon University

10:40 am – 12:30 pm  
Technical Insights I  
MISSION CITY BALLROOM B2 – 5

Technical Insights II  
ROOM 203 – 4

Fundamentals  
MISSION CITY BALLROOM M1 – 3

Business Insights  
THEATER

12:30 – 1:30 pm  
LUNCH

1:30 – 6:00 pm

Technical Insights I  
MISSION CITY BALLROOM B2 – 5

Technical Insights II  
ROOM 203 – 4

Fundamentals  
MISSION CITY BALLROOM M1 – 3

Business Insights  
THEATER

Enabling Technologies I  
HALL A2

Enabling Technologies II  
HALL A3

12:00 – 8:00 pm  
Vision  
Technology  
Showcase  
HALL A1 & HALL B

6:00 – 8:00 pm  
Vision  
Technology  
Showcase  
Reception  
HALL A1 & HALL B  
Join us for food,  
drink & demos!

5:00 – 6:00 pm  
Vision Entrepreneur's Panel  
THEATER

6:15 – 6:45 pm  
Vision Product of the Year Awards  
HALL A3

WEDNESDAY MAY 23

7:30 am – 6:00 pm  
REGISTRATION

8:00 – 9:00 am  
COFFEE & PASTRIES

9:00 – 10:30 am  
WELCOME & KEYNOTE MISSION CITY BALLROOM B2 – 5

From Mobility to Medicine: Vision Enables  
the Next Generation of Innovation

Dean Kamen  
Founder, DEKA Research and Development

10:40 am – 12:30 pm  
Technical Insights I  
MISSION CITY BALLROOM B2 – 5

Technical Insights II  
ROOM 203 – 4

Fundamentals  
MISSION CITY BALLROOM M1 – 3

Business Insights  
THEATER

Enabling Technologies I  
HALL A2

Enabling Technologies II  
HALL A3

12:30 – 1:30 pm  
LUNCH

10:30 am – 6:00 pm  
Vision  
Technology  
Showcase  
HALL A1 & HALL B

1:30 – 6:00 pm  
Technical Insights I  
MISSION CITY BALLROOM B2 – 5

Technical Insights II  
ROOM 203 – 4

Fundamentals  
MISSION CITY BALLROOM M1 – 3

Business Insights  
THEATER

Enabling Technologies I  
HALL A2

Enabling Technologies II  
HALL A3

5:00 – 6:00 pm  
Vision Tank Competition  
THEATER

THURSDAY  
MAY 24

Workshops require separate registration.  
Badge pickup is available at times and locations listed below.

8:00 – 9:00 am  
COFFEE & PASTRIES


9:00 am – 5:30 pm  
Khronos Standards  
for Neural Networks and  
Embedded Vision  
ROOM 203— LUNCH IN ROOM 204

REGISTRATION 8:00 am – 2:00 pm  
OUTSIDE ROOM 203  
This workshop covers Khronos standards related to neural networks and computer vision. The primary focus is on neural network inference workflows based on the new NNEF (Neural Network Interchange Format) standard.

FEE: \$50 Organized by 

8:45 am – 5:30 pm  
Artificial Intelligence: From  
Concept to Implementation  
BALLROOM H — LUNCH IN BALLROOM H FOYER

REGISTRATION 8:00 am – 2:00 pm  
OUTSIDE BALLROOM H  
In this workshop, you will learn how the latest advances in deep learning, artificial intelligence and embedded vision are being implemented in designs from automotive ADAS and IoT to industrial design.







FEE: \$25 Organized by 

9:00 am – 5:00 pm  
Optimized Inference  
at the Edge with Intel®  
ROOM 209—LUNCH IN ROOM 210

REGISTRATION 8:00 am – 2:00 pm  
OUTSIDE ROOM 209  
This hands-on workshop will take you through a computer vision workflow using the latest Intel® technologies and comprehensive toolkits including support for deep learning algorithms that help accelerate smart video applications.

FEE: \$25 Organized by 

Technical Insights I MISSION CITY BALLROOM B2 – B5	Technical Insights II ROOM 203/204	Fundamentals MISSION CITY BALLROOM M1 – M3		Business Insights THEATER	Enabling Technologies I HALL A2	Enabling Technologies II HALL A3
<p><b>10:40 – 11:10 am</b> <b>Portability and Performance in Embedded Deep Neural Networks: Can We Have Both?</b> Cormac Brick, Movidius, an Intel company</p> <p><b>11:20 am – 12:20 pm</b>  <b>Words, Pictures and Common Sense: Visual Question Answering</b> Devi Parikh, Georgia Tech and Facebook AI Research</p>	<p><b>10:40 – 11:10 am</b> <b>How Simulation Accelerates Development of Self-driving Technology</b> László Kishonti, Almotive</p> <p><b>11:20 – 11:50 am</b> <b>Computer Vision HW Acceleration for Driver Assistance</b> Markus Tremmel, Robert Bosch</p> <p><b>12:00 – 12:30 pm</b> <b>Understanding Real-World Imaging Challenges for ADAS and Autonomous Vision Systems—IEEE P2020</b> Felix Heide, Algolux</p>	<p><b>10:40 – 11:10 am</b>  <b>Solving Vision Tasks Using Deep Learning: An Introduction</b> Pete Warden, Google</p> <p><b>11:20 am – 12:20 pm</b> <b>From Feature Engineering to Network Engineering</b> Auro Tripathy, AMD</p>		<p><b>10:40 – 11:10 am</b> <b>What's Hot? The M&amp;A and Funding Landscape for Machine Vision Companies</b> Rudy Burger, Woodside Capital</p> <p><b>11:20 – 11:50 am</b>  <b>Ubiquitous \$90B AR to Dominate Focused \$15B VR by 2022, Becoming a Major Computer Vision Market</b> Tim Merel, Digi-Capital</p> <p><b>12:00 – 12:30 pm</b>  <b>EXECUTIVE PERSPECTIVE</b> <b>Balancing Safety, Convenience and Privacy in the Era of Ubiquitous Cameras</b> Charlotte Dryden, Intel</p>		
<b>12:30 – 1:30 pm</b> LUNCH				<b>12:30 – 1:30 pm</b> LUNCH		
<p><b>1:30 – 2:30 pm</b> <b>Even Faster CNNs: Exploring the New Class of Winograd Algorithms</b> Gian Marco Iodice, Arm</p> <p><b>2:50 – 3:20 pm</b> <b>Developing Computer Vision Algorithms for Networked Cameras</b> Dukhwan Kim, Intel</p> <p><b>3:30 – 4:00 pm</b> <b>Building a Typical Visual SLAM Pipeline</b> YoungWoo Seo, Hyperloop-One</p> <p><b>4:10 – 4:40 pm</b> <b>Programming Techniques for Implementing Inference Software Efficiently</b> Andrew Richards, Codeplay Software</p> <p><b>4:50 – 5:20 pm</b> <b>The OpenVX Computer Vision and Neural Network Inference Library Standard for Portable, Efficient Code</b> Radhakrishna Giduthuri, AMD</p> <p><b>5:30 – 6:00 pm</b> <b>APIs for Accelerating Vision and Inferencing: Options and Trade-offs</b> Neil Trevett, Khronos Group and NVIDIA</p>	<p><b>1:30 – 2:00 pm</b>  <b>The Roomba 980: Computer Vision Meets Consumer Robotics</b> Mario Munich, iRobot</p> <p><b>2:10 – 2:40 pm</b> <b>Deep Understanding of Shopper Behaviors and Interactions Using Computer Vision</b> Emanuele Frontoni and Rocco Pietrini, Università Politecnica delle Marche</p> <p><b>2:50 – 3:50 pm</b> <b>Getting More from Your Datasets: Data Augmentation, Annotation and Generative Techniques</b> Peter Corcoran, FotoNation (an Xperi company) and National University of Ireland Galway</p> <p><b>4:10 – 4:40 pm</b> <b>Recognizing Novel Objects in Novel Surroundings with Single-shot Detectors</b> Alexander C. Berg, UNC Chapel Hill</p> <p><b>4:50 – 5:20 pm</b> <b>Deploying CNN-based Vision Solutions on a \$3 Microcontroller</b> Greg Lytle, Au-Zone Technologies</p> <p><b>5:30 – 6:00 pm</b> <b>How to Get the Labeled Data for Free</b> Matt King, IUNU</p>	<p><b>1:30 – 2:30 pm</b> <b>Visual-inertial Tracking for AR and VR</b> Timo Ahonen, Meta</p> <p><b>2:50 – 3:50 pm</b> <b>Understanding and Implementing Face Landmark Detection and Tracking</b> Jayachandra Dakala, PathPartner Technology</p> <p><b>4:10 – 4:40 pm</b>  <b>Building a Practical Face Recognition System Using Cloud APIs</b> Chris Adzima, Washington County Sheriff's Office</p> <p><b>4:50 – 5:20 pm</b> <b>Bad Data, Bad Network, or: How to Create the Right Dataset for Your Application</b> Mike Schmit, AMD</p> <p><b>5:30 – 6:00 pm</b> <b>Introduction to Creating a Vision Solution in the Cloud</b> Nishita Sant, GumGum</p>		<p><b>1:30 – 2:00 pm</b> <b>Reduce Risk in Computer Vision Design: Focus on the User</b> Paul Duckworth, Twistthink</p> <p><b>2:10 – 2:40 pm</b> <b>Data-driven Business Models Enabled by 3D Vision Technology</b> Christopher Scheubel, FRAMOS</p> <p><b>2:50 – 3:20 pm</b>  <b>EXECUTIVE PERSPECTIVE</b> <b>Leveraging Edge and Cloud for Visual Intelligence Solutions</b> Salil Raje, Xilinx</p> <p><b>3:30 – 4:00 pm</b> <b>From 2D to 3D: How Depth Sensing Will Shape the Future of Vision</b> Guillaume Girardin, Yole Développement</p> <p><b>4:10 – 4:40 pm</b>  <b>EXECUTIVE PERSPECTIVE</b> <b>Embedded AI for Smart Cities and Retail in China</b> Kai Yu, Horizon Robotics</p> <p><b>5:00 – 6:00 pm</b>  <b>Vision Entrepreneurs' Panel</b> Moderator: Nik Gagvani President, CheckVideo  Radha Basu CEO, iMerit  Gary Bradski CTO, Arraiy &amp; CEO, OpenCV.org  László Kishonti CEO, Almotive</p>	<p><b>1:30 – 2:00 pm</b> <b>Machine Learning Inference in Under 5 mW with a Binarized Neural Network on an FPGA</b> Abdullah Raouf, Lattice</p> <p><b>2:10 – 2:40 pm</b>  <b>EXECUTIVE PERSPECTIVE</b> <b>Energy-efficient Processors Enable the Era of Intelligent Devices</b> Ren Wu, NovuMind</p> <p><b>2:50 – 3:20 pm</b> <b>High-end Multi-camera Technology, Applications and Examples</b> Max Larin, Ximea</p> <p><b>3:30 – 4:00 pm</b> <b>Mythic's Analog Deep Learning Accelerator Chip: High Performance Inference</b> Frederick Soo, Mythic</p> <p><b>4:10 – 4:40 pm</b> <b>Programmable CNN Acceleration in Under 1 Watt</b> Gordon Hands, Lattice</p> <p><b>4:50 – 5:20 pm</b> <b>A Physics-based Approach to Removing Shadows and Shading in Real Time</b> Bruce Maxwell, Tandent Vision Science</p> <p><b>5:30 – 6:00 pm</b> <b>At the Edge of AI at the Edge: Ultra Efficient AI on Low-power Compute Platforms</b> Mohammad Rastegari, XNOR.ai</p>	<p><b>1:30 – 2:00 pm</b> <b>A New Generation of Camera Modules: A Novel Approach and Its Benefits for Embedded Systems</b> Paul Maria Zalewski, Allied Vision</p> <p><b>2:10 – 2:40 pm</b> <b>Enabling Cross-platform Deep Learning Applications with the Intel CV SDK</b> Yury Gorbachev, Intel</p> <p><b>2:50 – 3:20 pm</b> <b>Achieving High-performance Vision Processing for Embedded Applications with Qualcomm SoC Platforms</b> Sahil Bansal, Qualcomm</p> <p><b>3:30 – 4:00 pm</b> <b>Infusing Visual Understanding in Cloud and Edge Solutions Using State-of-the-Art Microsoft Algorithms</b> Anirudh Koul and Jin Yamamoto, Microsoft</p> <p><b>4:10 – 4:40 pm</b> <b>Rapid Development of Efficient Vision Applications Using the Halide Language and CEVA Processors</b> Yair Siegel, CEVA and Gary Gitelson, mPerpetuo, Inc.</p>
Don't miss the <b>Vision Technology Showcase Reception 6:00 pm – 8:00 pm</b> for food, drink and demos!				Don't miss the <b>Vision Technology Showcase Reception 6:00 pm – 8:00 pm</b> for food, drink and demos!		

Technical Insights I MISSION CITY BALLROOM B2 – B5	Technical Insights II ROOM 203/204	Fundamentals MISSION CITY BALLROOM M1 – M3	Business Insights THEATER	Enabling Technologies I HALL A2	Enabling Technologies II HALL A3
<div>10:40 – 11:10 am Deep Quantization for Energy Efficient Inference at the Edge Hoon Choi, Lattice</div> <div>11:20 am – 12:20 pm  What is Neuromorphic Event-based Computer Vision? Sensors, Theory and Applications Ryad B. Benosman, University of Pittsburgh Medical Center, Carnegie Mellon University and Sorbonne Universit�s</div>	<div>10:40 – 11:10 am The Perspective Transform in Embedded Vision Aditya Joshi and Shrinivas Gadkari, Cadence</div> <div>11:20 – 11:50 am Harnessing the Edge and the Cloud Together for Visual AI S�bastien Taylor, Au-Zone Technologies</div> <div>12:00 – 12:30 pm New Deep Learning Techniques for Embedded Systems Tom Michiels, Synopsys</div>	<div>10:40 – 11:10 am Understanding Automotive Radar: Present and Future Arunesh Roy, NXP</div> <div>11:20 am – 12:20 pm Depth Cameras: A State-of-the-Art Overview Carlo Dal Mutto, Aquifi</div>	<div>10:40 – 11:10 am Leveraging Cloud Computer Vision for a Real-time Consumer Product Pavan Kumar, Cocoon Cam</div> <div>11:20 – 11:50 am Using Vision to Transform Retail Sumit Gupta, IBM</div> <div>12:00 – 12:30 pm Computer Vision for Industrial Inspection: The Evolution from PCs to Embedded Solutions Thomas D�ubler, NET New Electronic Technology GmbH</div>	<div>10:40 – 11:10 am Designing Smarter, Safer Cars with Embedded Vision Fergus Casey, Synopsys</div> <div>11:20 – 11:50 am Neural Network Compiler: Enabling Rapid Deployment of DNNs on Low-cost, Low-power Processors Megha Daga, Cadence</div> <div>12:00 – 12:30 pm New Memory-centric Architecture Needed for AI Sylvain Dubois, Crossbar</div>	<div>10:40 – 11:10 am Deep Learning on Arm Cortex-M Microcontrollers Vikas Chandra, Arm</div> <div>11:20 – 11:50 am Rethinking Deep Learning: Neural Compute Stick Ashish Pai, Intel</div> <div>12:00 – 12:30 pm Project Trillium: A New Suite of Machine Learning IP from Arm Steve Steele, Arm</div>
12:30 – 1:30 pm LUNCH					
<div>1:30 – 2:00 pm Real-time Calibration for Stereo Cameras Using Machine Learning Sheldon Fernandes, Lucid VR</div> <div>2:10 – 2:40 pm Building Efficient CNN Models for Mobile and Embedded Applications Peter Vajda, Facebook</div> <div>2:50 – 3:20 pm Utilizing Neural Networks to Validate Display Content in Mission Critical Systems Shang-Hung Lin, VeriSilicon</div> <div>3:30 – 4:00 pm Role of the Cloud in Autonomous Vehicle Vision Processing: A View from the Edge Ali Osman Ors, NXP</div> <div>4:10 – 4:40 pm  Generative Sensing: Reliable Recognition from Unreliable Sensor Data Lina Karam, Arizona State University</div> <div>4:50 – 5:20 pm Creating a Computationally Efficient Embedded CNN Face Recognizer G.B. Praveen, PathPartner Technology</div>	<div>1:30 – 2:00 pm  Implementing Image Pyramids Efficiently in Software Michael Stewart, Polymorphic Technologies</div> <div>2:10 – 2:40 pm Architecting a Smart Home Monitoring System with Millions of Cameras Hongcheng Wang, Comcast</div> <div>2:50 – 3:20 pm Improving and Implementing Traditional Computer Vision Algorithms Using DNN Techniques Paul Brasnett, Imagination Technologies</div> <div>3:30 – 4:00 pm Hybrid Semi-parallel Deep Neural Networks (SPDNN)—Example Methodologies and Use Cases Peter Corcoran, FotoNation (an Xperi company) and National University of Ireland Galway</div>	<div>1:30 – 2:30 pm  Approaches for Energy Efficient Implementation of Deep Neural Networks Vivienne Sze, MIT</div> <div>2:50 – 3:50 pm Introduction to Optics for Embedded Vision Jessica Gehlhar, Edmund Optics</div> <div>4:10 – 4:40 pm Introduction to Lidar for Machine Perception Mohammad Musa, DeepenAI</div> <div>4:50 – 5:20 pm Designing Vision Front Ends for Embedded Systems Friedrich Dierks, Basler</div> <div>5:30 – 6:00 pm Optimize Performance: Start Your Algorithm Development with the Imaging Subsystem Ryan Johnson, Twistthink</div>	<div>1:30 – 2:00 pm  EXECUTIVE PERSPECTIVE Building up a Start-up in Embedded Vision: Lessons from Machine Vision Arndt Bake, Basler</div> <div>2:10 – 2:40 pm The Four Key Trends Driving the Proliferation of Visual Perception Jeff Bier, Embedded Vision Alliance and BDTI</div> <div>2:50 – 3:20 pm  Overcoming Bias in Computer Vision—A Business Imperative Will Byrne, Entrepreneur</div> <div>3:30 – 4:00 pm  Intelligent Consumer Robots Powering the Smart Home Mario Munich, iRobot</div> <div>5:00 – 6:00 pm Vision Tank Start-up Competition AiFi Presented by Jo�o Diogo Falc�o Aquifi Presented by Carlo Dal Mutto Boulder AI Presented by Dan Connors Sturfee Presented by Anil Cheriyyadat VirtuSense Technologies Presented by Deepak Gaddipati</div>	<div>1:30 – 2:00 pm Enabling Software Developers to Harness FPGA Compute Accelerators Bernhard Friebe, Intel</div> <div>2:10 – 2:40 pm Deep Learning in MATLAB: From Concept to Optimized Embedded Code Girish Venkataramani and Avi Nehemiah, MathWorks</div> <div>2:50 – 3:20 pm Achieving 15 TOPS/s Equivalent Performance in Less Than 10 W Using Neural Network Pruning on Xilinx Zynq Nick Ni, Xilinx</div> <div>3:30 – 4:00 pm NovuTensor: Hardware Acceleration of Deep Convolutional Neural Networks for AI Mike Li, NovuMind</div> <div>4:10 – 4:40 pm The Journey and Sunrise Processors: Leading-edge Performance for Embedded AI Kai Yu, Horizon Robotics</div>	<div>1:30 – 2:00 pm Embedding Programmable DNNs in Low-power SoCs Petrone� Bigioi, FotoNation (an Xperi company)</div> <div>2:10 – 2:40 pm Exploiting Reduced Precision for Machine Learning on FPGAs Kees Vissers, Xilinx</div> <div>2:50 – 3:20 pm Optimizing Your System Software and BSP for Embedded Vision and AI Daniel Sun, ThunderSoft</div> <div>3:30 – 4:00 pm Pilot AI Vision Framework: From Doorbells to Defense Jonathan Su, Pilot AI</div>