embedded® VISION SUMMIT 2018

Enabling Computer Vision, at the Edge and in the Cloud

EVENT GUIDE

WELCOME

Embedded Vision Alliance founder Jeff Bier with Grégoire Gentil of Always Innovating's In/Out, audience choice winner at the 2017 Vision Tank Start-up Competition.

On behalf of the Embedded Vision Alliance, welcome to the 2018 Embedded Vision Summit!

Computer vision is becoming one of the most important technologies of our era, enabling products that are safer, more capable, easier to use and more autonomous, in applications that impact virtually every industry. The Summit is designed to inspire and empower product creators to incorporate visual intelligence into all types of systems.

I'd like to give you an idea of what you can expect over the next few days. We've introduced new elements, including a hands-on training class on Monday: Deep Learning for Computer Vision with TensorFlow. Tuesday and Wednesday comprise the core program and present our keynote addresses and over 90 practical technical and business sessions, including the popular Vision Entrepreneurs' panel. We'll also host our Vision Technology Showcase, which features more than 50 exhibitors and over 100 demos. Thursday is dedicated to Vision Technology Workshops, presented by experienced engineers from Embedded Vision Alliance Member companies.

The 2018 Summit will also include the first annual Vision Product of the Year Awards, which recognize the innovation of the companies that are enabling the next generation of practical applications for computer vision. And you won't want to miss our Vision Tank competition, which features the industry's most promising vision technology startups pitching their companies to well-known industry judges, in the hopes of winning the grand prize!

Before I close, I'd like to thank you for attending our conference and bringing your expertise to our community. If the Summit enables you to learn and teach, and to connect with vision technology professionals from other companies and countries, then it will have achieved our goals.

Thank you for joining us!

Jeff Bier, Founder
Embedded Vision Alliance

ABOUT THE embedded VISION ALLIANCE

The Embedded Vision Alliance® is a global partnership that brings together technology providers with end product and system developers who are enabling innovative, practical applications of computer vision.

Our mission is to inspire and empower product creators to incorporate visual intelligence into new products and applications, and enable Member companies to accelerate success in computer vision by:

- Bringing together suppliers, end-product designers and partners to speed the adoption of computer vision in products
- Delivering timely insights into market research, technology trends, standards and application requirements
- Enabling companies to become more visible as thought leaders

Membership is open to any company that supplies or uses technology for computer vision systems and applications.

For more information, visit **www.embedded-vision.com**.

CONTENTS

SUMMIT AGENDA

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** to find out about upcoming classes.

FEE: $895

Organized by embedded VISION ALLIANCE

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

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

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

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

SUMMIT AGENDA

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

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

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

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 KHRONOS GROUP

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 SYNOPSYS®

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 intel

Honoring innovation and achievement in computer vision.

The Vision Product of the Year Awards celebrate the innovations that are enabling the next generation of computer vision applications from the industry's leading companies.

With awards for best:

Camera
Processor
Developer Tools
AI Technology
Automotive Solutions
Cloud Technologies
Software or Algorithm
End Product

2018 winners will be announced on Tuesday, May 22 at the Welcome & Keynote.

KEYNOTE PRESENTATIONS

9:00 – 10:30 am
TUESDAY

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

Dr. Takeo Kanade will share his experiences and lessons learned in developing a vast range of pioneering computer vision systems and autonomous robots, including face recognition, autonomously-driven cars, computer-assisted surgical robots, robot helicopters, biological live cell tracking and a system for sports broadcasts.

Most researchers, when asked their fondest desire, respond that they want to do good research. If asked what constitutes "good research," they often find it difficult to give a clear answer. For Dr. Kanade, good research derives from solving real-world problems, delivering useful results to society. "Think like an amateur, do as an expert" is his research motto: When conceptualizing a problem and its possible solution, think simply and openly—as a novice in that field—without preconceived notions. When implementing a solution, on the other hand, do so thoroughly, meticulously and with expert skill.

In his research projects, Dr. Kanade has met and worked with people from diverse backgrounds, and has encountered many challenges. While exploring the technical side of some of his most important projects, he will also describe experiences that highlight the enjoyable aspects of a researcher's life—those that have occurred accidentally or inevitably as his "think like an amateur, do as an expert" approach has guided his interactions with problems and people.

ABOUT DR. TAKEO KANADE

Dr. Takeo Kanade is the U.A. and Helen Whitaker Professor for Carnegie Mellon University, and one of the world's foremost researchers in computer vision and robotics. His work spans theoretical foundations through practical use cases, including applications in facial recognition, motion tracking, virtual reality and robotics. He is the recipient of the 2016 Kyoto Prize, has published more than 300 technical articles, and holds over 30 patents.

9:00– 10:30 am
WEDNESDAY

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

Dean Kamen
Founder, DEKA Research and Development

In this presentation, legendary inventor Dean Kamen will explain why he believes the time is right for computer vision to be used everywhere. In Kamen's work, feedback control systems are central. For example, the Segway relies on feedback control to keep the rider stable on two wheels. The iBot agile wheelchair (from which the Segway was derived) relies on feedback control to perform complex maneuvers, such as ascending stairs. In Kamen's view, computer vision has now advanced to the point where it can serve as a ubiquitous, versatile sensor enabling feedback control in countless applications. Eventually, says Kamen, embedded vision sensors will be as common as simple microcontrollers or mechanical sensors are today.

Kamen will share what he's learned from the introduction of computer vision into the FIRST Robotics competition, a worldwide program that has inspired millions of young people's interest and participation in science and technology. Kamen will also present his company's work on its next-generation agile wheelchair. And, in perhaps his most ambitious initiative ever, Kamen will explain how his Advanced Regenerative Manufacturing Institute plans to enable the large-scale manufacturing of engineered tissues and tissue-related technologies, with the eventual goal of mass-producing replacement organs for humans. Kamen expects computer vision to play a key role, enabling monitoring and feedback control of tissue-growing processes without requiring physical contact with tissue.

ABOUT DEAN KAMEN

Dean Kamen is an innovator, but not just of things. As the founder of DEKA Research and Development, he has helped revolutionize attitudes, quality of life and awareness by developing life-changing products including a portable dialysis machine, a vascular stent and the iBot motorized wheelchair that climbs stairs. Kamen is also the founder of FIRST, the non-profit behind the successful robotics competitions which last year saw half a million students participating.

TUESDAY OVERVIEW

Technical Insights I
MISSION CITY BALLROOM B2 – B5

10:40 – 11:10 am
Portability and Performance in Embedded Deep Neural Networks: Can We Have Both?
Cormac Brick, Movidius, an Intel company

11:20 am – 12:20 pm
Words, Pictures and Common Sense: Visual Question Answering
Devi Parikh, Georgia Tech and Facebook AI Research

1:30 – 2:30 pm
Even Faster CNNs: Exploring the New Class of Winograd Algorithms
Gian Marco Iodice, Arm

2:50 – 3:20 pm
Developing Computer Vision Algorithms for Networked Cameras
Dukhwan Kim, Intel

3:30 – 4:00 pm
Building a Typical Visual SLAM Pipeline
YoungWoo Seo, Hyperloop-One

4:10 – 4:40 pm
Programming Techniques for Implementing Inference Software Efficiently
Andrew Richards, Codeplay Software

4:50 – 5:20 pm
The OpenVX Computer Vision and Neural Network Inference Library Standard for Portable, Efficient Code
Radhakrishna Giduthuri, AMD

5:30 – 6:00 pm
APIs for Accelerating Vision and Inferencing: Options and Trade-offs
Neil Trevett, Khronos Group and NVIDIA

Technical Insights II
ROOM 203/204

10:40 – 11:10 am
How Simulation Accelerates Development of Self-driving Technology
László Kishonti, AImotive

11:20 – 11:50 am
Computer Vision HW Acceleration for Driver Assistance
Markus Tremmel, Robert Bosch

12:00 – 12:30 pm
Understanding Real-World Imaging Challenges for ADAS and Autonomous Vision Systems—IEEE P2020
Felix Heide, Algolux

1:30 – 2:00 pm
The Roomba 980: Computer Vision Meets Consumer Robotics
Mario Munich, iRobot

2:10 – 2:40 pm
Deep Understanding of Shopper Behaviors and Interactions Using Computer Vision
Emanuele Frontoni and Rocco Pietrini, Università Politecnica delle Marche

2:50 – 3:50 pm
Getting More from Your Datasets: Data Augmentation, Annotation and Generative Techniques
Peter Corcoran, FotoNation (an Xperi company) and National University of Ireland Galway

4:10 – 4:40 pm
Recognizing Novel Objects in Novel Surroundings with Single-shot Detectors
Alexander C. Berg, UNC Chapel Hill

4:50 – 5:20 pm
Deploying CNN-based Vision Solutions on a $3 Microcontroller
Greg Lytle, Au-Zone Technologies

5:30 – 6:00 pm
How to Get the Labeled Data for Free
Matt King, IUNU

Fundamentals
MISSION CITY BALLROOM M1 – M3

10:40 – 11:10 am
Solving Vision Tasks Using Deep Learning: An Introduction
Pete Warden, Google

11:20 am – 12:20 pm
From Feature Engineering to Network Engineering
Auro Tripathy, AMD

1:30 – 2:30 pm
Visual-inertial Tracking for AR and VR
Timo Ahonen, Meta

2:50 – 3:50 pm
Understanding and Implementing Face Landmark Detection and Tracking
Jayachandra Dakala, PathPartner Technology

4:10 – 4:40 pm
Building a Practical Face Recognition System Using Cloud APIs
Chris Adzima, Washington County Sheriff's Office

4:50 – 5:20 pm
Bad Data, Bad Network, or: How to Create the Right Dataset for Your Application
Mike Schmit, AMD

5:30 – 6:00 pm
Introduction to Creating a Vision Solution in the Cloud
Nishita Sant, GumGum

12:30 – 1:30 pm LUNCH

Don't miss the **Vision Technology Showcase Reception 6:00 pm – 8:00 pm** for food, drink and demos!

= Invited presentation

TUESDAY OVERVIEW

Business Insights
THEATER

10:40 – 11:10 am
What's Hot? The M&A and Funding Landscape for Machine Vision Companies
Rudy Burger, Woodside Capital

11:20 – 11:50 am
Ubiquitous $90B AR to Dominate Focused $15B VR by 2022, Becoming a Major Computer Vision Market
Tim Merel, Digi-Capital

12:00 – 12:30 pm
EXECUTIVE PERSPECTIVE
Balancing Safety, Convenience and Privacy in the Era of Ubiquitous Cameras
Charlotte Dryden, Intel

Enabling Technologies I
HALL A2

Enabling Technologies II
HALL A3

12:30 – 1:30 pm LUNCH

1:30 – 2:00 pm
Reduce Risk in Computer Vision Design: Focus on the User
Paul Duckworth, Twisthink

2:10 – 2:40 pm
Data-driven Business Models Enabled by 3D Vision Technology
Christopher Scheubel, FRAMOS

2:50 – 3:20 pm
EXECUTIVE PERSPECTIVE
Leveraging Edge and Cloud for Visual Intelligence Solutions
Salil Raje, Xilinx

3:30 – 4:00 pm
From 2D to 3D: How Depth Sensing Will Shape the Future of Vision
Guillaume Girardin, Yole Développement

4:10 – 4:40 pm
EXECUTIVE PERSPECTIVE
Embedded AI for Smart Cities and Retail in China
Kai Yu, Horizon Robotics

5:00 – 6:00 pm
Vision Entrepreneurs' Panel
Moderator:
Nik Gagvani
President, CheckVideo

Radha Basu
CEO, iMerit

Gary Bradski
CTO, Arraiy & CEO, OpenCV.org

László Kishonti
CEO, AImotive

1:30 – 2:00 pm
Machine Learning Inference in Under 5 mW with a Binarized Neural Network on an FPGA
Abdullah Raouf, Lattice

2:10 – 2:40 pm
EXECUTIVE PERSPECTIVE
Energy-efficient Processors Enable the Era of Intelligent Devices
Ren Wu, NovuMind

2:50 – 3:20 pm
High-end Multi-camera Technology, Applications and Examples
Max Larin, Ximea

3:30 – 4:00 pm
Mythic's Analog Deep Learning Accelerator Chip: High Performance Inference
Frederick Soo, Mythic

4:10 – 4:40 pm
Programmable CNN Acceleration in Under 1 Watt
Gordon Hands, Lattice

4:50 – 5:20 pm
A Physics-based Approach to Removing Shadows and Shading in Real Time
Bruce Maxwell, Tandent Vision Science

5:30 – 6:00 pm
At the Edge of AI at the Edge: Ultra Efficient AI on Low-power Compute Platforms
Mohammad Rastegari, XNOR.ai

1:30 – 2:00 pm
A New Generation of Camera Modules: A Novel Approach and Its Benefits for Embedded Systems
Paul Maria Zalewski, Allied Vision

2:10 – 2:40 pm
Enabling Cross-platform Deep Learning Applications with the Intel CV SDK
Yury Gorbachev, Intel

2:50 – 3:20 pm
Achieving High-performance Vision Processing for Embedded Applications with Qualcomm SoC Platforms
Sahil Bansal, Qualcomm

3:30 – 4:00 pm
Infusing Visual Understanding in Cloud and Edge Solutions Using State-of-the-Art Microsoft Algorithms
Anirudh Koul and Jin Yamamoto, Microsoft

4:10 – 4:40 pm
Rapid Development of Efficient Vision Applications Using the Halide Language and CEVA Processors
Yair Siegel, CEVA and Gary Gitelson, mPerpetuo, Inc.

Don't miss the **Vision Technology Showcase Reception 6:00 pm – 8:00 pm** for food, drink and demos!

= Invited presentation

WEDNESDAY OVERVIEW

Technical Insights I
MISSION CITY BALLROOM B2 – B5

10:40 – 11:10 am
Deep Quantization for Energy Efficient Inference at the Edge
Hoon Choi, Lattice

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 Universitas

1:30 – 2:00 pm
Real-time Calibration for Stereo Cameras Using Machine Learning
Sheldon Fernandes, Lucid VR

2:10 – 2:40 pm
Building Efficient CNN Models for Mobile and Embedded Applications
Peter Vajda, Facebook

2:50 – 3:20 pm
Utilizing Neural Networks to Validate Display Content in Mission Critical Systems
Shang-Hung Lin, VeriSilicon

3:30 – 4:00 pm
Role of the Cloud in Autonomous Vehicle Vision Processing: A View from the Edge
Ali Osman Ors, NXP

4:10 – 4:40 pm
Generative Sensing: Reliable Recognition from Unreliable Sensor Data
Lina Karam, Arizona State University

4:50 – 5:20 pm
Creating a Computationally Efficient Embedded CNN Face Recognizer
G.B. Praveen, PathPartner Technology

Technical Insights II
ROOM 203/204

10:40 – 11:10 am
The Perspective Transform in Embedded Vision
Aditya Joshi and Shrinivas Gadkari, Cadence

11:20 – 11:50 am
Harnessing the Edge and the Cloud Together for Visual AI
Sébastien Taylor, Au-Zone Technologies

12:00 – 12:30 pm
New Deep Learning Techniques for Embedded Systems
Tom Michiels, Synopsys

1:30 – 2:00 pm
Implementing Image Pyramids Efficiently in Software
Michael Stewart, Polymorphic Technologies

2:10 – 2:40 pm
Architecting a Smart Home Monitoring System with Millions of Cameras
Hongcheng Wang, Comcast

2:50 – 3:20 pm
Improving and Implementing Traditional Computer Vision Algorithms Using DNN Techniques
Paul Brasnett, Imagination Technologies

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

Fundamentals
MISSION CITY BALLROOM M1 – M3

10:40 – 11:10 am
Understanding Automotive Radar: Present and Future
Arunesh Roy, NXP

11:20 am – 12:20 pm
Depth Cameras: A State-of-the-Art Overview
Carlo Dal Mutto, Aquifi

1:30 – 2:30 pm
Approaches for Energy Efficient Implementation of Deep Neural Networks
Vivienne Sze, MIT

2:50 – 3:50 pm
Introduction to Optics for Embedded Vision
Jessica Gehlhar, Edmund Optics

4:10 – 4:40 pm
Introduction to Lidar for Machine Perception
Mohammad Musa, DeepenAI

4:50 – 5:20 pm
Designing Vision Front Ends for Embedded Systems
Friedrich Dierks, Basler

5:30 – 6:00 pm
Optimize Performance: Start Your Algorithm Development with the Imaging Subsystem
Ryan Johnson, Twisthink

12:30 – 1:30 pm LUNCH

= Invited presentation

WEDNESDAY OVERVIEW

Business Insights
THEATER

10:40 – 11:10 am
Leveraging Cloud Computer Vision for a Real-time Consumer Product
Pavan Kumar, Cocoon Cam

11:20 – 11:50 am
Using Vision to Transform Retail
Sumit Gupta, IBM

12:00 – 12:30 pm
Computer Vision for Industrial Inspection: The Evolution from PCs to Embedded Solutions
Thomas Däubler, NET New Electronic Technology GmbH

1:30 – 2:00 pm
EXECUTIVE PERSPECTIVE
Building up a Start-up in Embedded Vision: Lessons from Machine Vision
Arndt Bake, Basler

2:10 – 2:40 pm
The Four Key Trends Driving the Proliferation of Visual Perception
Jeff Bier, Embedded Vision Alliance and BDTI

2:50 – 3:20 pm
Overcoming Bias in Computer Vision—A Business Imperative
Will Byrne, Entrepreneur

3:30 – 4:00 pm
Intelligent Consumer Robots Powering the Smart Home
Mario Munich, iRobot

5:00 – 6:00 pm
Vision Tank Start-up Competition
AiFi Presented by João Diogo Falcão

Aquifi Presented by Carlo Dal Mutto

Boulder AI Presented by Dan Connors

Sturfee Presented by Anil Cheriyadat

VirtuSense Technologies
Presented by Deepak Gaddipati

Enabling Technologies I
HALL A2

10:40 – 11:10 am
Designing Smarter, Safer Cars with Embedded Vision
Fergus Casey, Synopsys

11:20 – 11:50 am
Neural Network Compiler: Enabling Rapid Deployment of DNNs on Low-cost, Low-power Processors
Megha Daga, Cadence

12:00 – 12:30 pm
New Memory-centric Architecture Needed for AI
Sylvain Dubois, Crossbar

1:30 – 2:00 pm
Enabling Software Developers to Harness FPGA Compute Accelerators
Bernhard Friebe, Intel

2:10 – 2:40 pm
Deep Learning in MATLAB: From Concept to Optimized Embedded Code
Girish Venkataramani and Avi Nehemiah, MathWorks

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

3:30 – 4:00 pm
NovuTensor: Hardware Acceleration of Deep Convolutional Neural Networks for AI
Mike Li, NovuMind

4:10 – 4:40 pm
The Journey and Sunrise Processors: Leading-edge Performance for Embedded AI
Kai Yu, Horizon Robotics

Enabling Technologies II
HALL A3

10:40 – 11:10 am
Deep Learning on Arm Cortex-M Microcontrollers
Vikas Chandra, Arm

11:20 – 11:50 am
Rethinking Deep Learning: Neural Compute Stick
Ashish Pai, Intel

12:00 – 12:30 pm
Project Trillium: A New Suite of Machine Learning IP from Arm
Steve Steele, Arm

1:30 – 2:00 pm
Embedding Programmable DNNs in Low-power SoCs
Petronel Bigioi, FotoNation (an Xperi company)

2:10 – 2:40 pm
Exploiting Reduced Precision for Machine Learning on FPGAs
Kees Vissers, Xilinx

2:50 – 3:20 pm
Optimizing Your System Software and BSP for Embedded Vision and AI
Daniel Sun, ThunderSoft

3:30 – 4:00 pm
Pilot AI Vision Framework: From Doorbells to Defense
Jonathan Su, Pilot AI

12:30 – 1:30 pm LUNCH

= Invited presentation

PRESENTATIONS MAP

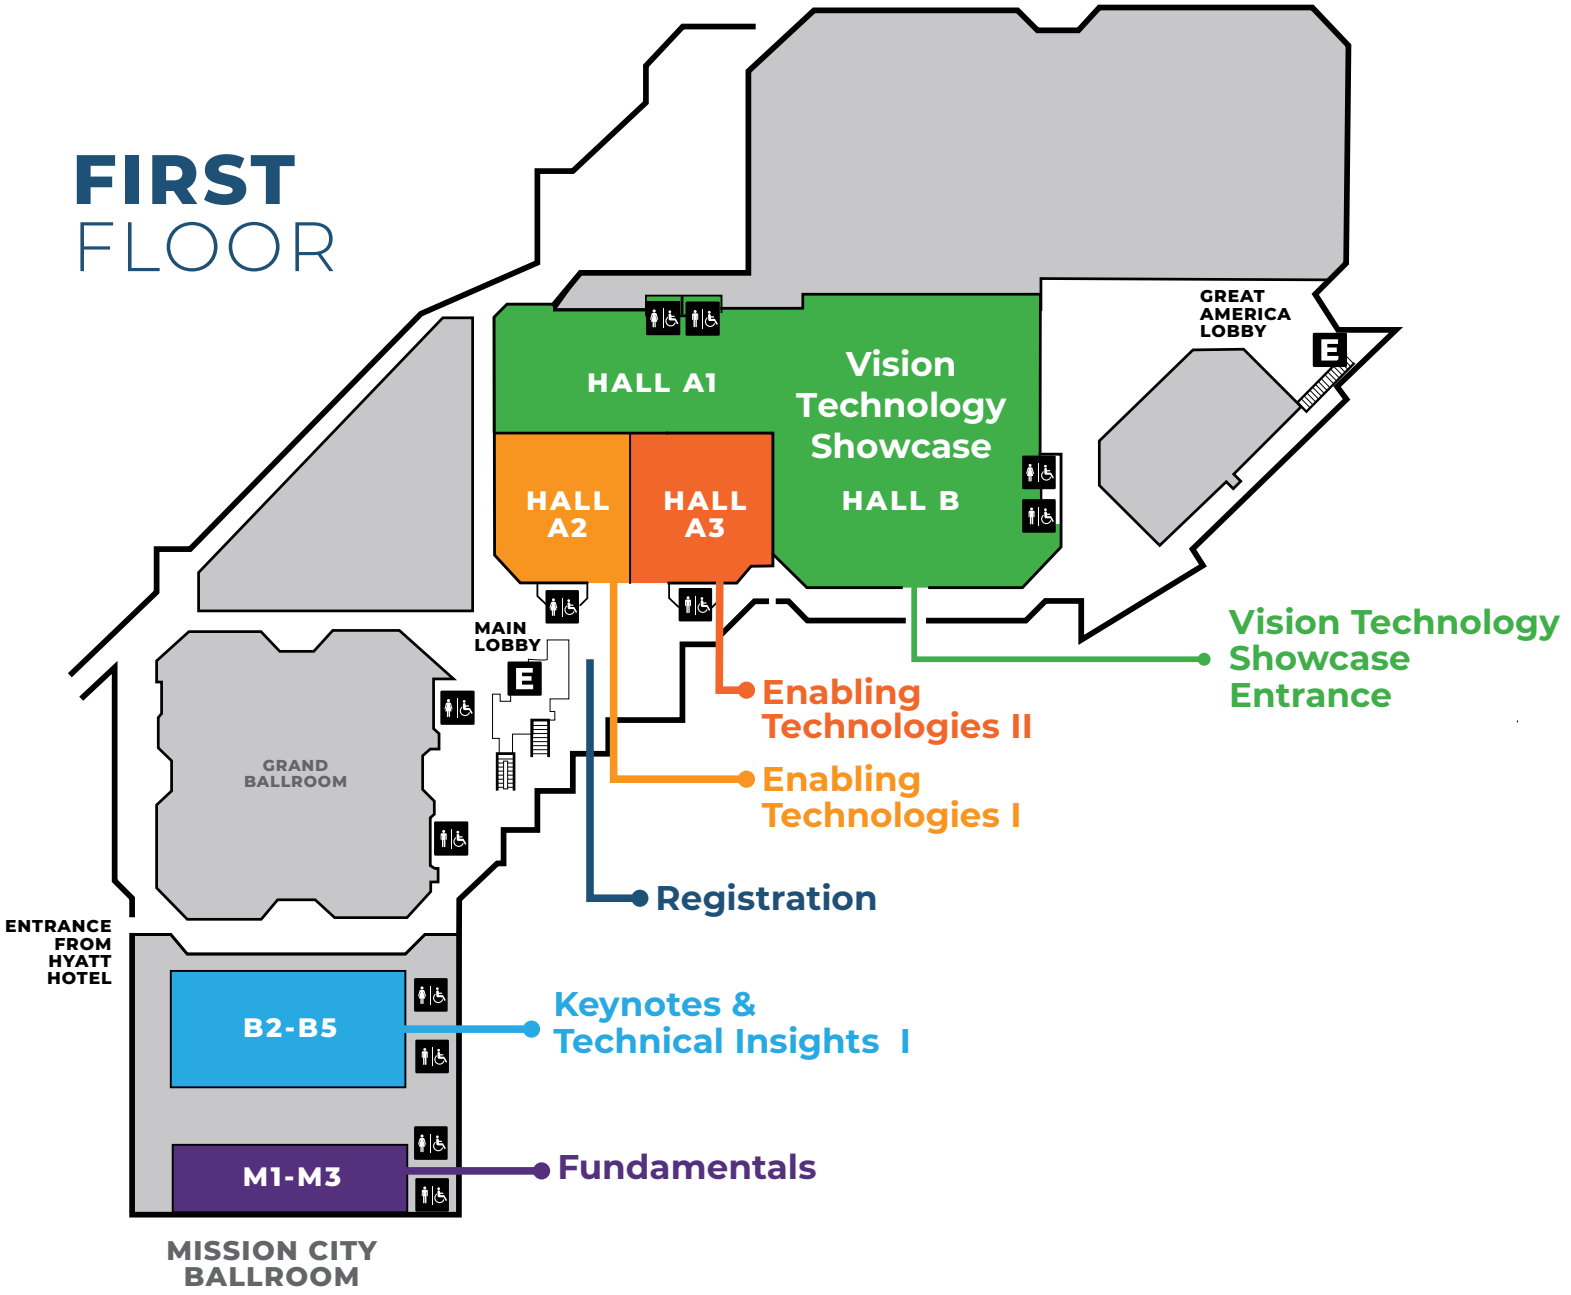

PRESENTATIONS MAP

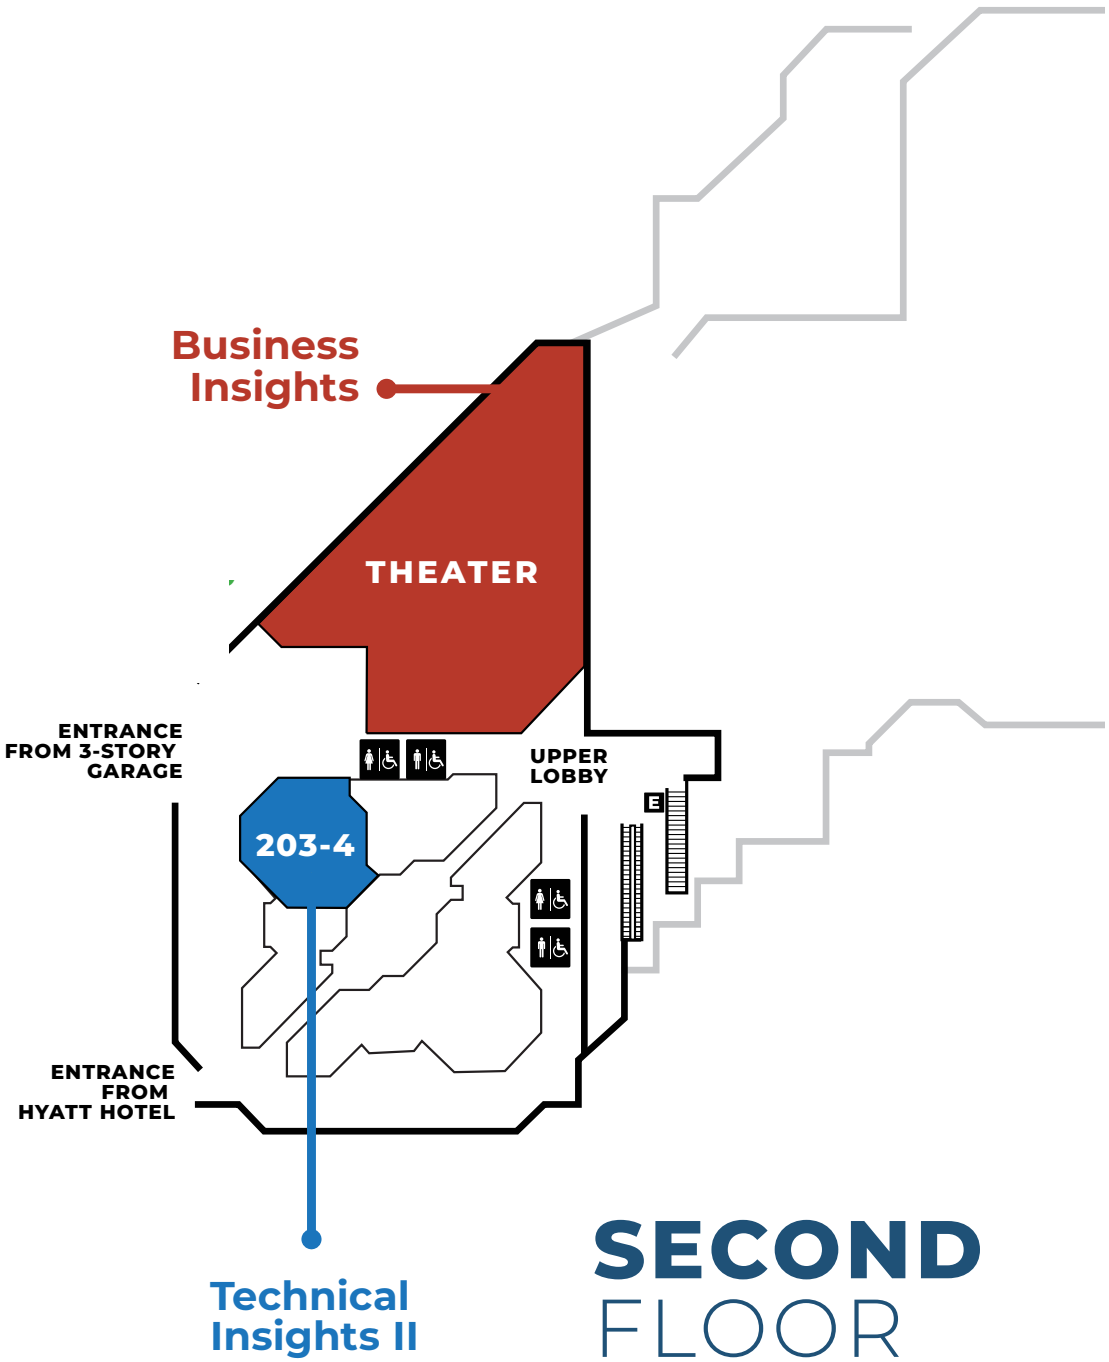

SECOND FLOOR

WORKSHOPS

Khronos Standards for Neural Networks and Embedded Vision

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. Using examples, we illustrate the mapping of neural networks and computer vision algorithms to the OpenVX graph API. Also covered is the deployment model that pre-compiles a graph to create optimized binaries for deployment use cases such as neural network inference. The workshop includes a demo session that shows the participants how to solve real computer vision and neural network problems using Khronos standards.

BY ATTENDING THIS WORKSHOP YOU WILL:

Gain an understanding of the architecture of Khronos standards for computer vision and neural networks

Develop fluency in actually using NNEF and OpenVX for real-time computer vision and neural networks

Become familiar with SYCL and its ecosystem

ORGANIZED BY

KHRONOS GROUP

Artificial Intelligence from Concept to Implementation

Join our full-day workshop to learn about how the latest advances in artificial intelligence, deep learning and embedded vision are being implemented in designs from automotive ADAS and IoT to industrial design. You'll hear from technology leaders about the current and future state of AI, and dive deep into the hardware and software needed to make it happen. Gain practical "how-to" knowledge from Synopsys experts and our customers on topics including porting your code to OpenCL C, using OpenVX and how to compare CNN graphs. Finally, you'll get insights into hot new technologies including the latest Synopsys embedded vision processors.

BY ATTENDING THIS WORKSHOP YOU WILL:

Hear from technology leaders about the latest AI advances

Learn about state-of-the-art hardware and software for embedded designs

Gain practical how-to knowledge from Synopsys experts and our customers

ORGANIZED BY

SYNOPSYS®

Optimized Inference at the Edge with Intel®

Join us for a one-day, hands-on workshop where Intel 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. You will learn how to optimize and improve performance with and without external accelerators and utilize tools to help you identify the best hardware configuration for your needs. This workshop will also outline the various frameworks and topologies supported by Intel accelerator tools in addition to a brief discussion on how to implement custom layers.

BY ATTENDING THIS WORKSHOP YOU WILL:

Learn more about Intel computer vision technology

Gain insight into how to accelerate computer vision applications

Explore performance optimization techniques with hardware acceleration

ORGANIZED BY

intel®

All workshops take place on Thursday, May 24. On Monday, Tuesday or Wednesday, visit the Registration desk in the main lobby to check availability of workshop seats. On Thursday, visit the workshop room to check availability.

WORKSHOPS

SECOND FLOOR

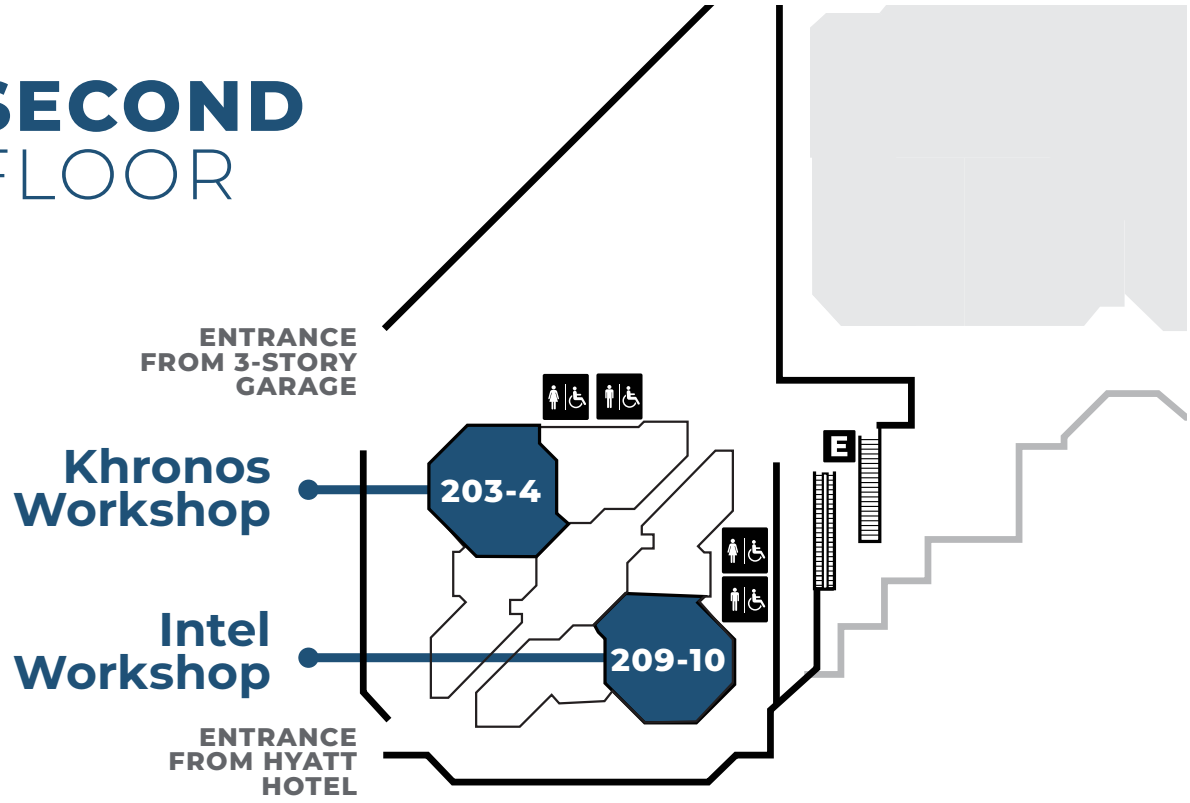

FIRST FLOOR

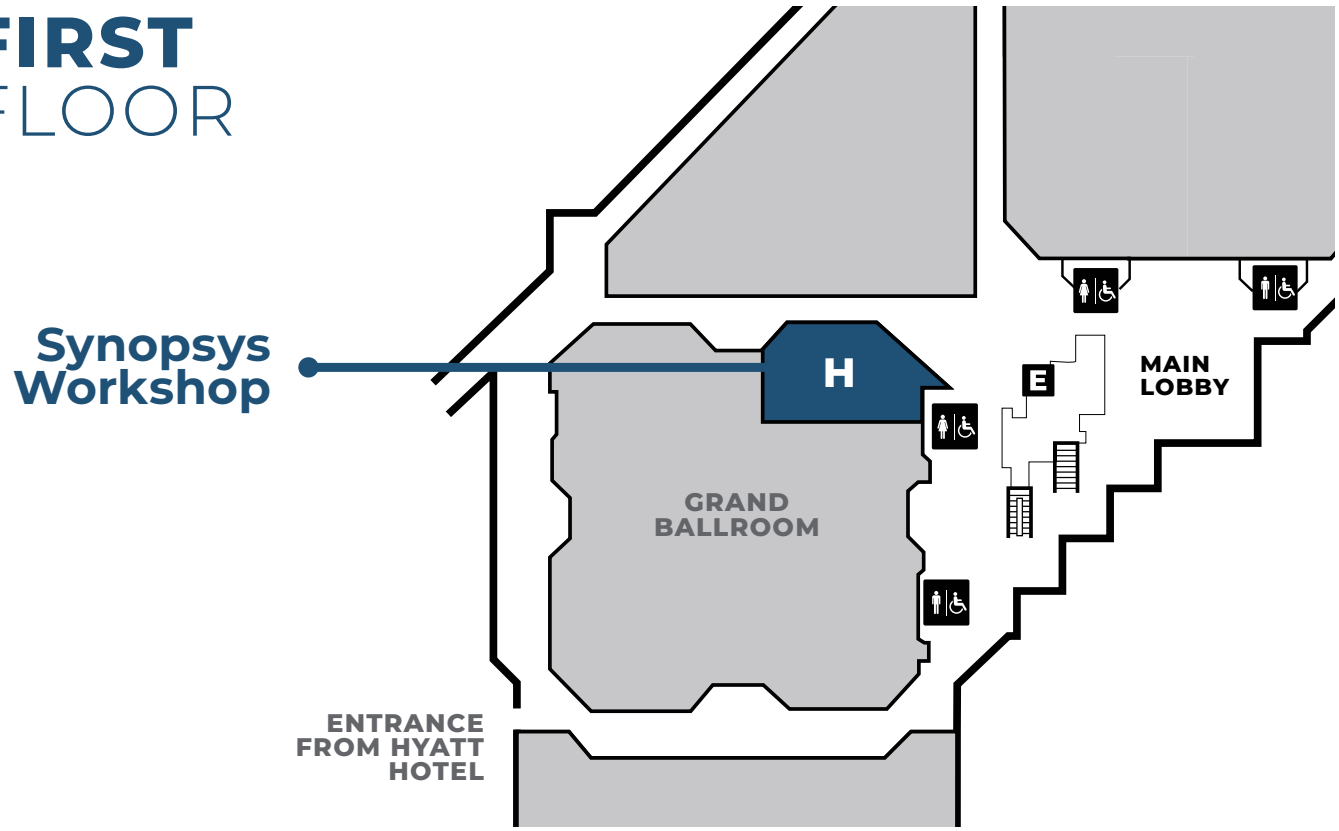

VISION TECHNOLOGY SHOWCASE

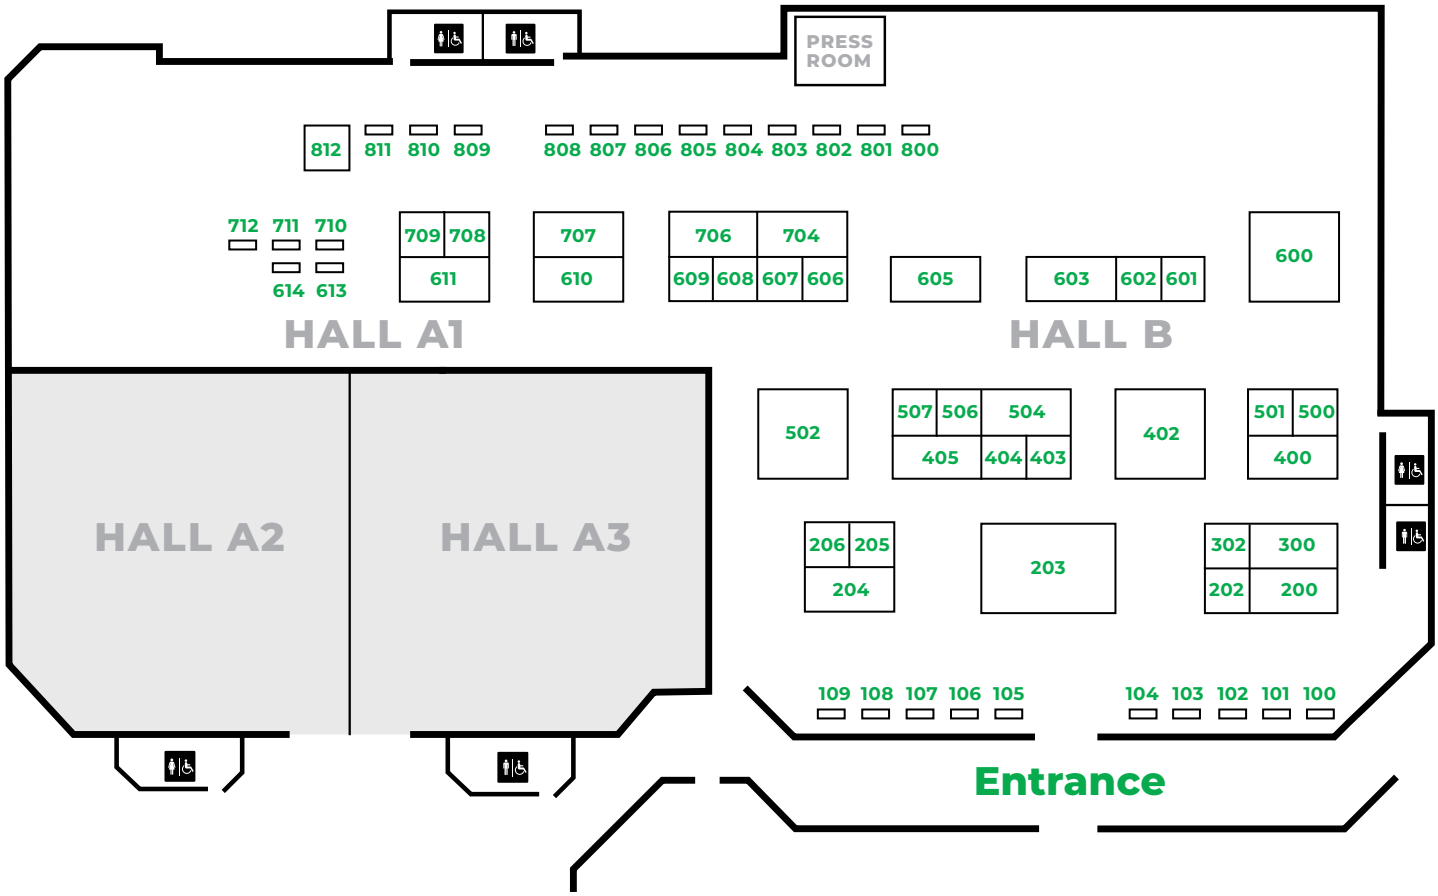

achronix™

Achronix, a privately held, fabless semiconductor corporation, is showcasing its Speedcore eFPGA IP, which can be integrated into an ASIC or SoC to provide a customized programmable fabric. Users specify their logic, memory and DSP resource needs, then Achronix configures the Speedcore IP to meet their individual requirements.

CONTACT

Alok Sanghavi,
Sr. Marketing Manager
aloksanghavi@achronix.com
www.achronix.com

Vision Tank Finalist

AIFI

AiFi is building a scalable version of "Amazon Go" to empower stores of the future to be check-out free. AiFi's innovative AI-powered sensor networks also provide retailers with valuable insights about shopping behavior and product preference, as well as improved inventory management.

CONTACT

Steve Gu, CEO, Aifi, Inc.
steve@aifi.io

SILVER SPONSOR

AIMOTIVE

AImotive is a global provider of vision-first self-driving technology. We utilize artificial intelligence, simulation and supporting hardware architectures for a safe autonomous experience. aiDrive, our scalable vision-first self-driving solution will be showcased through videos of real-world testing. Our technology will also drive live in our advanced simulator for autonomous vehicle development, aiSim. Alongside aiWare, these solutions form the foundation of AImotive's vision for the scalable future of autonomy.

CONTACT

Szabolcs Jánky,
Business Development
szabolcs.janky@aimotive.com
aimotive.com

ALDEC
THE DESIGN VERIFICATION COMPANY

Aldec will present two demos:

- ADAS reference designs based on the TySOM-3-ZU7EV (Zynq Ultrascale+ MPSoC) + FMC-ADAS, including 360-degree surround view, bird's eye view, driver drowsiness detection, and a smart rear camera
- Zynq hardware/software co-simulation solution for Zynq architecture based on QEMU and Riviera-PRO

CONTACT

Farhad Fallah
Farhadf@aldec.com
www.aldec.com

VISION TECHNOLOGY SHOWCASE

algolux

Algolux demos:

- **CRISP-ML**: Workflow tool using machine learning to automatically optimize your imaging/vision system through objective metrics, shrinking image quality tuning from months to hours.
- **CANA**: Full DNN stack for more robust perception in difficult conditions (e.g. low light, adverse weather), 30%+ better accuracy than state-of-the-art alternatives.

CONTACT

Dave Tokic, VP Mktg/Partnerships
dave.tokic@algolux.com
www.algolux.com

PREMIER & LANYARD SPONSOR

Allied Vision

The 1 product line is a new versatile technical platform that is designed and optimized for embedded vision. Instead of a traditional FPGA-based design used in machine vision cameras, the 1 product line are powered by our new ALVIUM® Technology. The 1 product line 130 and 140 Series support a very large range of sensors from 0.5 to 18 megapixel, several industry approved interfaces such as MIPI CSI-2 and USB3 Vision, different feature sets, as well as various housing options like board level, open end, complete housing and various lens mounts.

CONTACT

Contact sales at info@alliedvision.com

AMD

For more than 45 years AMD has driven innovation in high-performance computing, graphics and visualization technologies — the building blocks for gaming, immersive platforms and the datacenter. Hundreds of millions of consumers, Fortune 500 businesses and cutting-edge scientific research facilities worldwide rely on AMD technology daily to improve how they live, work and play.

CONTACT

Guy Ludden
Radeon Open Compute (ROCm) Group
www.amd.com

Vision Tank Finalist

Aquifi provides visual inspection services for logistics and manufacturing, based on the combination of 3D reconstruction and deep learning. The company's solution, a trainable virtual inspection system, increases the throughput of human workers and reduces errors due to fatigue and repetition.

CONTACT

Carlo Dal Mutto, CTO
cdm@aquifi.com
www.aquifi.com

GOLD SPONSOR

arm

Arm recently announced Project Trillium, a new suite of Arm® IP that brings machine learning (ML) to edge devices. Arm will demonstrate a variety of its machine learning technologies for embedded through client devices, as well as showing a few use case products for attendees to explore.

CONTACT

Tim Hartley
tim.hartley@arm.com
+44 7788 750 900

Au-Zone Technologies Inc.

Au-Zone is a leading provider of development tools, engineering design services, and enabling IP for intelligent embedded vision products. By utilizing our tools, our customers can quickly develop and securely deploy machine learning solutions. Through our engineering consulting engagements, we help our clients lower development costs, mitigate program risk and shorten time to revenue.

CONTACT

Brad Scott, President
brad@au-zone.com
www.embeddedml.com

BASLER

Basler is the leading global provider of high quality industrial cameras and camera modules for a wide range of applications. With 30 years of vision expertise and a dedicated embedded portfolio, Basler supports manufacturers worldwide in incorporating cutting edge vision technology into their products and applications.

CONTACT

Daniel Toth, Partner Manager
daniel.toth@baslerweb.com
www.baslerweb.com

VISION TECHNOLOGY SHOWCASE

GOLD SPONSOR

BDTI helps companies create products that incorporate computer vision and deep learning. BDTI specializes in designing custom algorithms that meet unique customer requirements, creating efficient software that executes demanding algorithms within tight cost and power budgets, and enabling informed decisions on the best techniques and technologies for customer products. See demos of DNN object detection, 3D sensing, and object measurement at Booth 603.

CONTACT

Jeremy Giddings 1.925.954.1411
giddings@bdti.com
www.BDTI.com

Vision Tank Finalist

Boulder AI has created an intelligent GPU-enabled deep-learning neural network camera, DNNcam, that is waterproof and dust-proof. The camera executes AI/machine learning and computer vision algorithms at the edge, distilling visual information into actionable event data. The end-to-end Boulder AI platform enables collecting edge data events into cloud environments.

CONTACT

Dan Conners, Co-Founder & CTO
dan@boulderAI.com
www.boulderAI.com

Brodmann17's advanced deep-learning algorithms produce state-of-the-art vision accuracy with only a fraction of the usual computation load. Brodmann17 will demonstrate how edge devices such as autonomous vehicles and ADAS can handle deep-learning vision on standard low-power processors. Brodmann17 is making IoT and automotive edge-devices cloud-free and autonomous.

CONTACT

Adi Pinhas, Co-founder & CEO
Adi@brodmann17.com
www.brodmann17.com

GOLD SPONSOR

Cadence will demonstrate its high-performance and low-power vision and AI DSPs, designed to handle complex imaging, computer vision, and AI processing functions in mobile handset, automotive, AR/VR, surveillance, drone, and wearable products.

CONTACT

Pulin Desai,
Product Marketing Director
pulin@cadence.com
www.cadence.com

CEVA is the leading licensor of signal processing platforms and artificial intelligence processors for a smarter, connected world and a range of end markets. CEVA's ultra-low-power IP for vision, audio, communications and connectivity includes DSP-based platforms for advanced imaging, computer vision and deep learning for any camera-enabled device.

CONTACT

Yair Siegel
Yair.Siegel@ceva-dsp.com
www.ceva-dsp.com

We will demonstrate fully hardwired deep learning inference IP performing object detection on 4Kp30 video from real-time camera inputs. Lens distortion correction will also be demonstrated, where images from wide-angled lenses are corrected for enhanced input images. The demonstrations will be performed on FPGA-based boards.

CONTACT

Philip Han, Head of Marketing
marketing@chipsnmedia.com
www.chipsnmedia.com

Crossbar ReRAM embedded non-volatile memory technology enables massive amounts of computational bandwidth at the lowest energy consumption when used to store AI trained models on the same silicon die as computing cores running neural networks and algorithms. The demos showcase ReRAM for object classification, face recognition and license plate recognition.

CONTACT

Sylvain DuBois, VP Bus Dev & Mktg
sylvain.dubois@crossbar-inc.com
www.crossbar-inc.com

DEKA Research & Development Corporation develops internally generated inventions and provides research and development for major corporate clients. DEKA's innovative devices have expanded the frontiers of health care worldwide. Some of DEKA's notable inventions include the first wearable insulin pump for diabetics, the HomeChoice™ portable peritoneal dialysis machine, the LUKE prosthetic arm, the iBot stair climbing wheel chair and the Segway Human Transporter.

CONTACT

Dirk Van Der Merwe
dmerwe@dekaresearch.com

VISION TECHNOLOGY SHOWCASE

If your company develops vision-based end products, visit the Embedded Vision Alliance booth to learn how we can accelerate your development and reduce risk.

If you're a provider of vision components, software or services, learn how the Alliance can connect you with customers, partners and provide early insights into key market and technology trends.

CONTACT

Kim Vaupen or Ruthann Fisher
info@embedded-vision.com
www.embedded-vision.com

FIRST® is a movement. The oldest and largest nonprofit organization of its kind, *FIRST* inspires innovation by teaching science, engineering, technology, math (STEM), and leadership skills through hands-on robotics challenges developed to ignite curiosity and passion in K-12 students.

CONTACT

www.firstinspires.org

FLIR Systems, Inc. is a global leader in the design and manufacture of innovative, high-performance digital cameras for industrial, medical and life science, traffic, biometric, GIS, and people counting applications.

CONTACT

Preston Barrett,
Territory Account Manager
preston.barrett@flir.com
www.flir.com/mv

FRAMOS® will be showcasing SONY's CMOS Starvis Rolling Shutter and Pregius Global Shutter sensors, including their new IMX250 polarized sensor, Intel®'s RealSense™ Technology, and a sampling of modules and cameras ready for any embedded vision application.

CONTACT

Chris Donegan,
Sales Mgr. North America
C.donegan@framos.com
www.framos.com

GreenWaves Technologies is a fabless semiconductor start-up designing disruptive ultra-low power embedded solutions for image, sound and vibration AI processing in sensing devices. GreenWaves will be showing GAP8, the industry's first ultra-low power processor enabling battery operated interpretation of images, sounds and vibrations in Internet of Things (IoT) applications.

CONTACT

Martin Croome, VP Business Dev.
sales@
greenwaves-technologies.com
www.greenwaves-technologies.com

Gidel provides intelligent FPGA solutions for Acceleration and Imaging. Its Infinivision technology enables companies to develop 3D mapping, VR and AR products, by capturing high-quality image content from a large array of cameras to create panoramic, 360 content for automotive applications and next-generation immersive fan experiences across media, sports and entertainment.

CONTACT

Nurit Ben Moshe
bm_nurit@gidel.com
www.gidel.com

GOLD SPONSOR

地 平 线
Horizon Robotics

Horizon Robotics is a leading technology powerhouse, dedicated to providing integrated and open edge artificial intelligence (AI) solutions with high performance, low power and low cost. After two years' R&D, we unveiled China's first world-leading, Brain Processing Unit (BPU) proprietary Gauss architecture-based edge AI computer vision processors, Journey and Sunrise. They power smart cars and smart cameras, providing industrial customers with a complete solution including chips, algorithms and cloud support.

CONTACT

Yufeng Zhang, VP Global Business
yufeng.zhang@hobot.cc
www.horizon.ai

Using object recognition, segmentation, and variable precision demonstrations Imagination will show the benefits of neural network acceleration for edge devices using either Imagination's PowerVR 2NX, a complete, highly efficient standalone hardware IP neural network accelerator solution for SoC; or a combination of the 2NX NNA and PowerVR GPUs together.

CONTACT

David Harold,
VP Marketing Communications
David.Harold@imgtec.com
www.imgtec.com

VISION TECHNOLOGY SHOWCASE

PREMIER PLUS SPONSOR

Intel, a leader in computing innovation, is driving the evolution of edge-to-cloud vision solutions, helping unlock new possibilities for the data that businesses generate with a comprehensive stack of products designed for AI. The company's robust hardware and software portfolio gives OEMs/ODMs, system integrators, ISVs, and solution providers the tools required to accelerate the design, development and deployment of high performance computer vision solutions. With heterogeneous camera-to-cloud inference and acceleration silicon—Intel® Movidius-VPUs; Intel® FPGAs, CPUs and CPUs with integrated graphics—along with high performance analytics development and deployment tools, Intel is enabling rich solutions for AI everywhere.

CONTACT

Brenda Christoffer, Marketing Specialist
Brenda.A.Christoffer@intel.com
www.intel.com

iMerit is a technology services company, delivering data to some of the most innovative companies in machine learning, eCommerce, and computer vision. iMerit's "humans in the loop" AI services are recognized globally for enabling advanced computing capabilities. iMerit does so while effecting positive social and economic change by empowering marginalized youth and young women.

CONTACT

Robert Frary, Director
robert@imerit.net
www.imerit.net

ImmerVision enables intelligent vision in the world's devices. The company designs patented, augmented resolution, wide-angle lenses and AI-ready image processing so consumer devices, and professional, automotive, robotics, and medical applications can See More, Smarter. ImmerVision licenses its technology to innovative component, OEM, and ODM manufacturers.

CONTACT

Angus Mackay, Director, Marketing & Communications
angus.mackay@immervision.com
www.immervisionenables.com

PREMIER PLUS SPONSOR

Lattice will showcase our small form factor, ultra-low power, production priced FPGAs addressing embedded vision needs in Industrial, Automotive and Consumer markets. Learn how we are creating innovative low power solutions for face tracking for surveillance cameras, collision avoidance for industrial robots, and speed sign detection for automotive aftermarket cameras. On display will be the latest demos on image sensor connectivity, computer vision, and machine learning inferencing, based on ECP5- and CrossLink-FPGAs.

CONTACT

Deepak Boppana
Deepak.bopaana@latticesemi.com
408-826-6336

SILVER SPONSOR

Luxoft is a leading independent software service provider for Automotive OEMs, Tier 1s and semiconductor companies. Transition to autonomous cars stimulates massive investments into advanced technologies that enable vehicles to drive themselves, and will change expectations of in-vehicle user experiences. We develop high-end automotive software solutions across UX/UI, HMI, ADAS, connectivity, IoT, telematics and navigation that enable gradual introduction of autonomous drive. Computer Vision and AI are crucial parts of our technology and industry transformation.

CONTACT

automotive.luxoft.com
automotive@luxoft.com

Deep Learning in MATLAB

MATLAB makes it easy to design deep learning based vision applications and deploy optimized, generated code to multiple embedded processors like Jetson TX2, Drive PX2, Intel®-based CPUs, or ARM-based platforms. Visit the MathWorks booth to experience the live demos and learn more!

CONTACT

Sandeep Hiremath, Product Mkting
sandeep.hiremath@mathworks.com
www.mathworks.com

VISION TECHNOLOGY SHOWCASE

Morpho, Inc., a global leader in embedded image processing software, will showcase SoftNeuro™, the world's fastest deep learning inference engines. SoftNeuro™ obtains profile data from target platforms that execute inference, and performs optimizations based on the data to achieve higher speeds, making it easy to deploy trained multi-framework networks.

CONTACT

Toshi Torihara, Vice President
h-torihara@morphoinc.com
www.morphoinc.com/en/

MVTec is demonstrating an inspection application that utilizes Deep Learning.
The images are being acquired live with a GigE camera and processed on the Jetson TX2 from NVIDIA. In addition we are showing an identification application on the Raspberry PI utilizing the integrated camera module.

CONTACT

Heiko Eisele
sales@mvtec.us
1.617.401.2112

Mentor provides custom compilers, HPC libraries, Catapult HLS (high level synthesis), and consultation services for deploying computer vision and machine leaning applications on embedded platforms requiring performance accelerators such as FPGAs, GPUs, DSPs and SIMD engines.

CONTACT

Pete Decher, Director Business Dev.
Pete_Decher@mentor.com
www.mentor.com/embedded

Microsoft's Cognitive Services offer a broad spectrum of vision capabilities which power imaginative and inspired uses. They include ready-to-use capabilities such as image tagging, content moderation, OCR and face detection, as well as fully customizable classifiers and object detectors that can be exported for edge delivery through the Custom Vision service.

CONTACT

Cornelia Carapcea,
Principal Program Manager
corncar@microsoft.com
www.microsoft.com/en-us/

NALBI specializes in deep learning for embedded systems. Our technologies include highly optimized deep learning models and a computing engine especially optimized for the specific target device. We present real-time, accurate human segmentation and detection on embedded devices.
The solution can be used in mobile apps, surveillance, smart home, and more.

CONTACT

Kina Jin, CEO
kinajin@nalbi.ai
www.nalbi.ai

NET will be presenting a real-time object tracking system using a linescan camera. Its FPGA is configurable with custom functionality, and the x86 architecture computer can run Windows/Linux and image processing libraries of a customer's choice. NET´s Open Camera Concept enables solution providers to create their own embedded vision solutions.

CONTACT

Grzegorz Kolodynski,
Marketing & PR
g.kolodynski@net-gmbh.com
www.net-gmbh.com

Nextchip is a vision solution company that offers an extraordinary ISP (image signal processor) with functions such as HDR, DeFog, 3DNR, LFM etc., as well as vision-based ADAS solutions and SVM (surround view monitoring) software.
The company's experience and capabilities are enabling Nextchip to take its next steps as a global image expert group, tackling myriad projects and challenges in the automotive field.

CONTACT

Mathias Sunghoon Chung,
Deputy General Manager
sunghoonch@nextchip.com
www.nextchip.com

Embedded vision is now possible in all light conditions, even in outdoor scenes with high illumination, thanks to our HDR sensor NSC1602 and MAGIC mono board.
Live demos will show high dynamic range, combined with advanced image processing for a reliable face recognition application.

CONTACT

Nicolas Baroan, BDM
Tel: +33 1 64 47 88 58
info@
new-imaging-technologies.com

VISION TECHNOLOGY SHOWCASE

GOLD SPONSOR

NOVUMIND
Making Things Think

NovuMind is dedicated to improving your life through Artificial Intelligence by making things think. Through cutting edge, in-house-developed artificial intelligence technology, NovuMind combines big data, high-performance, and heterogeneous computing to change the Internet of Things (IoT) into the Intelligent Internet of Things (I^2oT). Our NovuTensor chip does tensor computation at the speed of silicon, to provide unsurpassed performance-to-power ratios. NovuTensor is ideal for AI applications such as fast video object detection or video resolution enhancement.

CONTACT

NovuMind Inc, Santa Clara CA
info@novumind.com
www.novumind.com

SILVER SPONSOR

NXP

NXP Semiconductors N.V. enables secure connections and infrastructure for a smarter world, advancing solutions that make lives easier, better and safer.
As the world leader in secure connectivity solutions for embedded applications, NXP is driving innovation in embedded vision solutions for the secure connected vehicle, end-to-end security & privacy and smart connected solutions markets, built on more than 60 years of combined experience and expertise.

CONTACT

Ali Osman Ors
Director, AI Strategy & Partnerships
ali.ors@nxp.com
www.nxp.com

PATHPARTNER

PathPartner, a global product engineering specialist, is demonstrating its expertise in developing solutions for advanced embedded vision use-cases including :

- Advanced driver assistance systems for traffic sign detection, vehicle/ pedestrian detection
- Driver monitoring systems

Also, don't miss our speaker session on "Creating a Computationally-efficient Embedded CNN Face Recognizer."

CONTACT

Ramkishor Korada, Sales & Mkt
ramkishor.korada
@pathpartnertech.com
www.pathpartnertech.com

GOLD SPONSOR

QUALCOMM®

Qualcomm invents breakthrough technologies that transform how the world connects and communicates. When we connected the phone to the Internet, the mobile revolution was born. Today, our inventions are the foundation for life-changing products, experiences, and industries. As we lead the world to 5G, we envision this next big change in cellular technology spurring a new era of intelligent connected devices and enabling new opportunities in connected cars, networking, and the IoT —including smart cities, smart homes and wearables.

CONTACT

www.qualcomm.com

sama SOURCE

Companies like Google, Microsoft, Qualcomm and NVIDIA choose Samasource's training data services to power their human-in-the-loop artificial intelligence and machine learning projects. And because we offer fully managed services, we can guarantee your service levels and create reliability for your business. Get more out of your training data with Samasource.

CONTACT

Karolina Zajac
Senior Director of Global Sales
Karolina.zajac@samasource.org

scale

Scale accelerates the development of AI applications by helping generate high-quality ground truth data for computer vision teams including Cruise, Voyage and Embark. Scale specializes in a variety of perception use cases and industries including LiDAR point cloud, video annotation, and semantic segmentation with high accuracy and scalable volumes.

CONTACT

sales@scaleapi.com

STRADVISION

StradVision provides SVNet, an accurate production-ready perception software on automotive embedded hardware, for ADAS and autonomous driving. SVNet helps vehicles identify and navigate around objects for a better, safer driving experience. It enables vehicles to sense precisely where they are in space and in relation to their surroundings.

CONTACT

Hak-Kyoung Kim,
Algorithm Engineer
hak-kyoung.kim@stradvision.com
www.stradvision.ai

Vision Tank Finalist

STURFEE

Sturfee is building city-scale Visual Positioning Service (VPS) based on deep learning, computer vision and satellite imaging principles, enabling camera connected devices and machines to precisely locate themselves in the real world, identify where they are looking, and recognize what is around them—all based on visual input data. Cameras need VPS more than GPS.

CONTACT

Sheng Huang
Head of Business
Operations & Partnerships
www.sturfee.com

VISION TECHNOLOGY SHOWCASE

ThunderSoft

We will showcase embedded AI algorithms including face detection, scene recognition, object tracking, and food recognition. These AI algorithms can be optimized to run on various computing architectures, such as GPUs, DSPs, CPUs or dedicated AI chipsets.

CONTACT

Olivia Bai, Marketing Director
baijie@thundersoft.com
www.thundersoft.com

PREMIER SPONSOR

SYNOPSYS®

Come to Synopsys's booth to learn and discuss the latest embedded vision techniques and hardware to implement deep learning in edge applications including surveillance, AR/MR, mobile and automotive. Synopsys and our customers and partners will demonstrate new technology using the DesignWare EV6x Embedded Vision Processors including object and face recognition, Android neural networks and sparse optical flow. The programmable, scalable EV6x processors include scalar, vector DSP and CNN processing units for highly accurate and fast vision processing. They combine the flexibility of software solutions with the high performance and low power consumption of dedicated hardware.

CONTACT

Gordon Cooper, Product Marketing Manager, SG
gordonc@synopsys.com
www.synopsys.com

twisthink

Twisthink is a value-added partner helping companies bring IoT and vision-based products to life, using strategic insights and applying intuitive design with technology. Come to the Twisthink space (#109) to discuss what's next for your business and see product examples showcasing custom algorithms, cameras, UI and UX design and IoT.

CONTACT

Kaitlyn Marsman, Marketing Lead
kaitlyn@twisthink.com
www.twisthink.com

VeriSilicon

VeriSilicon's VIP8000 processors reach performance and memory efficiency of dedicated fixed-function logic with the customizable future-proofing of full programmability in OpenCL, OpenVX and a wide range of NN frameworks. VeriSilicon's exhibit will feature information about its OpenVX extensions in the i.MX8 applications processor.

CONTACT

www.verisilicon.com
vision IP: viv_info@verisilicon.com
Turnkey design services:
us_sales@verisilicon.com

VIA we connect

VIA Technologies, Inc. is a global leader in the development of highly integrated embedded platform and system solutions for M2M, IoT, and Smart City applications, ranging from video walls and digital signage to healthcare and industrial automation. VIA's customer base includes the world's leading high-tech, telecommunications, and consumer electronics industry brand names.

CONTACT

Jason Lee Gillikin, Bus Dev Manager
JasonLeeGillikin@ViaTech.com
www.ViaTech.com

videantis

videantis is a one-stop deep learning, computer vision and video processor IP provider. Together with our partners, we deliver low-power, high-performance, intelligent visual sensing and computing platforms to the automotive, mobile, consumer, and embedded markets.

CONTACT

Tony Picard, VP Sales
tony.picard@videantis.com
www.videantis.com

Vision Tank Finalist

VirtuSense Technologies

VirtuSense Technologies' product identifies people who are at risk of falls and injuries. The core technology is based on machine vision, using a 3D time-of-flight sensor to track a person's static and dynamic balance, identify sensory and muscular deficits and provide objective data to assess and treat issues.

CONTACT

Deepak Gaddipati, CTO
deepakg@virtusensetech.com
www.virtusensetech.com

VISION TECHNOLOGY SHOWCASE

Vision Components will showcase the new make & model feature of its multi-platform ALPR (automatic license plate recognition) that runs on Android phones, tablets, PCs and in the cloud. Our latest VC Nano 3D-Z embedded 3D triangulation system based on a dual-core ARM CPU and Linux O/S, will also be demonstrated with an angle measurement application.

CONTACT

Mariann M. Kiraly, Bus Dev Director
mariann.kiraly@
vision-components.com
www.vision-components.com

Wave Computing, Inc. is the Silicon Valley company that is revolutionizing artificial intelligence and deep learning with its dataflow-based systems. The company's vision is to "follow the data" and bring deep learning to customers' data wherever it may be—from the datacenter to the edge of the cloud.

CONTACT

info@wavecomp.com
wavecomp.ai

wrnch is teaching cameras to read human body language. The wrnchAI engine is a real-time AI software platform to digitize human motion and behaviour from standard video. wrnch inc is a computer vision/deep learning company based in Montréal, Canada.

CONTACT

Dr. Paul Kruszewski, CEO
paul@wrnch.ai
www.wrnch.ai

XIMEA designs and produces leading-edge, high-performance cameras with the lowest power consumption and the smallest footprint, as well as highly optimized cameras and imaging solutions, based on a PCI Express interface. The portfolio is targeted towards industrial, hyperspectral, scientific, high-speed, high-resolution, multi-sensor as well as OEM vision (sub-) systems.

CONTACT

Michael Cmok, Tech. Sales Dir.
mc@ximea.com
www.ximea.com

XNOR.ai offers state-of-the-art deep learning at the edge and will demonstrate object detection running in real-time on a single core of an iPhone 6s CPU, an Ambarella S5LM and a Samsung Galaxy 8. XNOR.ai is up to 10x faster, 200x more power efficient, and requires 8-15x less memory than floating-point CNNs.

CONTACT

Dan Waters, VP Mktg & Bus Dev
dan@xnor.ai | www.xnor.ai
AI at your fingertips

TUESDAY OVERVIEW

10:40 – 11:10 am

Portability and Performance in Embedded Deep Neural Networks: Can We Have Both?

Cormac Brick
Director of Machine Intelligence,
Movidius, an Intel company

In recent years, there has been much work done on low-precision deep neural network inference. The results show that by training for quantization, large gains in energy efficiency can be achieved. On the other hand, we have seen embedded runtime packages, like TensorFlow Lite and Caffe2go, emerge that offer portability over a number of platforms.

In this talk we look at the challenges presented by the choice between efficiency and portability, and ask, "Why can't we have both?" We will examine how big the gap between portability and efficiency truly is today, showing performance for a range of popular vision applications using state-of-the-art methods for trained networks. We present best-in-class design techniques for developing portable networks to maximize performance on different compute architectures. And, we identify key challenges remaining and progress needed to close the portability-performance gap.

11:20 am – 12:20 pm

INVITED PRESENTATION

Words, Pictures and Common Sense: Visual Question Answering

Devi Parikh
Professor, Georgia Tech and
Facebook AI Research

Wouldn't it be nice if machines could understand content in images and communicate this understanding as effectively as humans? Such technology would be immensely powerful, be it for aiding a visually-impaired user navigate a world built by the sighted, assisting an analyst in extracting relevant information from a surveillance feed, educating a child playing a game on a touch screen, providing information to a spectator at an art gallery or interacting with a robot. As computer vision and natural language processing techniques are maturing, we are closer to achieving this dream than we have ever been.

Visual Question Answering is one step in this direction. Given an image and a natural language question about the image (e.g., "What kind of store is this?" "Is it safe to cross the street?"), the machine's task is to automatically produce an accurate natural language answer. In this talk, we present our dataset, results obtained using neural models and open research questions in free-form and open-ended Visual Question Answering.

1:30 – 2:30 pm

Even Faster CNNs: Exploring the New Class of Winograd Algorithms

Gian Marco Iodice
Senior Software Engineer,
Machine Learning Group, Arm

Over the past decade, deep learning networks have revolutionized the task of classification and recognition in a broad area of applications. Deeper and more accurate networks have been proposed every year and more recent developments have shown how these workloads can be implemented on modern low-power embedded platforms. In this presentation, we discuss a recently introduced class of algorithms to reduce the arithmetic complexity of convolution layers with small filter sizes.

After an introduction to the latest optimization techniques for the most common solutions, such as GEMM, we will dive deeply into the design of Winograd algorithms, analyzing the complexity and the performance achieved for convolutional neural networks.

TECHNICAL INSIGHTS I

TUESDAY

2:50 – 3:20 pm

Developing Computer Vision Algorithms for Networked Cameras

Dukhwan Kim
Software Architect, Intel

Video analytics is one of the key elements in network cameras. Computer vision capabilities such as pedestrian detection, face detection and recognition and object detection and tracking are necessary for effective video analysis. With recent advances in deep learning technology, many developers are now utilizing deep learning to implement these capabilities. However, developing a deep learning algorithm requires more than just training models using Caffe or TensorFlow. It should start from an understanding of use cases, which affect the nature of required training datasets, and should be tightly bound with the hardware platform to get the best performance.

In this presentation, we explain how we have developed and optimized production-quality video analytics algorithms for computer vision applications.

3:30 – 4:00 pm

Building a Typical Visual SLAM Pipeline

YoungWoo Seo
Senior Director,
Hyperloop-One

Maps are important for both human and robot navigation. Simultaneous localization and mapping (SLAM) is one of the core techniques for map-based navigation. As SLAM algorithms have matured and hardware has improved, SLAM is spreading into many new applications, from self-driving cars to floor-cleaning robots.

In this talk, we walk through a typical pipeline for SLAM, specifically visual SLAM. A typical visual SLAM pipeline, based on visual feature tracking, begins with extracting visual features and matching the extracted features to previously surveyed features. It then continues with an estimation of the current camera poses based on the feature matching results, executing a (local) bundle adjustment to jointly optimize camera poses and map points and lastly performing a loop-closure routine to complete maps. While explaining each of these steps, we also cover challenges, tips, open source libraries, performance metrics and publicly available benchmark-datasets.

4:10 – 4:40 pm

Programming Techniques for Implementing Inference Software Efficiently

Andrew Richards
CEO & Founder,
Codeplay Software

When writing software to deploy deep neural network inferencing, developers are faced with an overwhelming range of options, from a custom-coded implementation of a single model to using a deep learning framework like TensorFlow or Caffe. If you custom code your own implementation, how do you balance the competing needs of performance, portability and capability? If you use an off-the-shelf framework, how do you get good performance? Andrew and his company have been building and standardizing developer tools for GPUs and AI accelerators for over 15 years.

We explore the available approaches for implementing deep neural networks in software, from the low-level details of how to map software to the highly parallel processors needed for AI all the way up to major AI frameworks. This will start with the LLVM compiler chain used to compile for most GPUs, through the OpenCL, HSA and SYCL programming standards (including how they compare with CUDA), all the way up to TensorFlow and Caffe and how they affect the key metrics like performance.

4:50 – 5:20 pm

The OpenVX Computer Vision and Neural Network Inference Library Standard for Portable, Efficient Code

Radhakrishna Giduthuri
Software Architect,
Advanced Micro Devices

OpenVX is an industry-standard computer vision and neural network inference API designed for efficient implementation on a variety of embedded platforms. The API incorporates the concept of a dataflow graph, which enables implementers to apply a range of optimizations appropriate to their architectures, such as image tiling and kernel fusion. Application developers can use this API to create high-performance computer vision and AI applications quickly, without having to perform complex device-specific optimizations for data management and kernel execution, since these optimizations are handled by the OpenVX implementation provided by the processor vendor.

We describe the current status of OpenVX, with particular focus on neural network inference capabilities and the most recent enhancements. We then conclude with a summary of the currently available implementations and an overview of the roadmap for the OpenVX API and its implementations.

5:30 pm – 6:00 pm

APIs for Accelerating Vision and Inferencing: Options and Trade-offs

Neil Trevett
President, Khronos and
VP, NVIDIA

The landscape of SDKs, APIs and file formats for accelerating inferencing and vision applications continues to evolve rapidly. Low-level compute APIs, such as OpenCL, Vulkan and CUDA are being used to accelerate inferencing engines, such as OpenVX, CoreML, NNAPI and TensorRT. Inferencing engines are being fed via neural network file formats, such as NNEF and ONNX. Some of these APIs, like OpenCV, are vision-specific, while others, like OpenCL, are general-purpose. Some engines, like CoreML and TensorRT, are supplier-specific, while others, such as OpenVX, are open standards that any supplier can adopt. Which ones should you use for your project?

In this presentation, we present the current landscape of APIs, file formats, and SDKs for inferencing and vision acceleration, explaining where each one fits in the development flow. We also highlight where these APIs overlap and where they complement each other, and preview some of the latest developments in these APIs.

WEDNESDAY OVERVIEW

10:40 – 11:10 am

Deep Quantization for Energy Efficient Inference at the Edge

Hoon Choi
Senior Director of
Design Engineering, Lattice

Intelligence at the edge is different from intelligence in the cloud relative to requirements for energy, cost, accuracy and latency. Due to limits on battery power and cooling systems in edge devices, energy consumption is strictly limited. Further, low cost and small size requirements make it hard to use packages with large numbers of pins, thus limiting the bandwidth to DRAM chips commonly used for storing neural network algorithm information. Despite these limitations, most applications require real-time operation. To tackle this issue, we have developed networks that heavily rely on deep quantization.

In this talk, we show how to use the deep quantization in real applications without degrading accuracy. Specifically, we explain the use of different quantizations for each layer of a deep neural network and how to use deep layered neural networks with deep quantization. We also explains the use of this deep quantization approach with recent lightweight networks.

11:20 am – 12:20 pm

INVITED PRESENTATION

What is Neuromorphic Event-based Computer Vision? Sensors, Theory and Applications

Ryad B. Benosman
Professor, University of Pittsburgh Medical Center, Carnegie Mellon University and Sorbonne Universitas

This talk introduces neuromorphic, event-based approaches for image sensing and processing. State-of-the-art image sensors suffer from severe limitations imposed by their very principle of operation. These sensors acquire the visual information as a series of "snapshots" recorded at discrete points in time, hence time-quantized at a predetermined frame rate, resulting in limited temporal resolution, low dynamic range and a high degree of redundancy in the acquired data. Nature suggests a different approach: Biological vision systems are driven and controlled by events happening within the scene in view, and not—like conventional image sensors—by artificially created timing and control signals that have no relation to the source of the visual information.

Translating the frameless paradigm of biological vision to artificial imaging systems implies that control over the acquisition of visual information is no longer imposed externally on an

array of pixels, but rather the decision making is transferred to each pixel, which handles its own information individually.

We will introduce the fundamentals underlying such bio-inspired, event-based image sensing and processing approaches, and explore strengths and weaknesses. We will show that bio-inspired vision systems have the potential to outperform conventional, frame-based vision acquisition and processing systems and to establish new benchmarks in terms of data compression, dynamic range, temporal resolution and power efficiency in applications such as 3D vision, object tracking, motor control, and visual feedback loops, in real-time.

1:30 – 2:00 pm

Real-time Calibration for Stereo Cameras Using Machine Learning

Sheldon Fernandes
Senior Software and Algorithms Engineer, Lucid VR

Calibration involves capturing raw data and processing it to get useful information about a camera's properties. Calibration is essential to ensure that a camera's output is as close as possible to what it "sees." Calibration for a stereo pair of cameras is even more critical because it also obtains data on the cameras' positions relative to each other. These extrinsic parameters ensure that 3D data can be properly rectified for viewing, and enable further advanced processing, such as obtaining disparity and depth maps and performing 3D reconstruction. In order for advanced processing to work correctly, calibration data should be error-fee. With age, heat and external conditions, extrinsic properties of a camera can change.

In this presentation, we discuss calibration techniques and a model for calibration, and propose advanced techniques using machine learning to estimate changes in extrinsic parameters in real time.

2:10 – 2:40 pm

Building Efficient CNN Models for Mobile and Embedded Applications

Peter Vajda
Research Scientist, Facebook

Recent advances in efficient deep learning models have led to many potential applications in mobile and embedded devices. In this talk, we discuss state-of-the-art model architectures, and introduce our work on real-time style transfer and pose estimation on mobile phones.

2:50 – 3:20 pm

Utilizing Neural Networks to Validate Display Content in Mission Critical Systems

Shang-Hung Lin
Vice President, Vision & Imaging Products, VeriSilicon

Mission-critical display systems in aerospace, automotive and industrial markets require validation of the content presented to the user in order to enable detection of potential failures and triggering of failsafe mechanisms. Traditional validation methods are based on pixel-perfect matching between expected and presented content. As user interface (UI) designs in these systems become more elaborate, the traditional validation methods become obsolete, and must be replaced with more robust methods that can recognize the mission critical information in a dynamic UI.

In this talk, we explore the limitations of the current content integrity checking systems and how they can be overcome by deployment of neural network pattern classification in the display pipeline. we also discuss the downscaling of these neural networks to run efficiently in a functionally safe microcontroller environment, and the requirements imposed on such solutions by the safety standards enforced in these domains.

3:30 – 4:00 pm

Role of the Cloud in Autonomous Vehicle Vision Processing: A View from the Edge

Ali Osman Ors
Director, Automotive Microcontrollers and Processors, NXP

Regardless of the processing topology—distributed, centralized or hybrid—sensor processing in automotive is an edge compute problem. However, with the improvement of connectivity technologies, vehicles are transforming into connection hubs, and this provides opportunities to augment the edge compute capabilities of today's vehicles. Conversely, the amount of data created at the sensors is increasing rapidly, creating a bandwidth shortage.

Drawing on design experience from deployed applications and emerging technologies, we explore the role and impact cloud-based services will have on applications that are predominantly reliant on

4:10 – 4:40 pm

INVITED PRESENTATION

Generative Sensing: Reliable Recognition from Unreliable Sensor Data

Lina Karam
Professor and Computer Engineering Director, Arizona State University

While deep neural networks (DNNs) perform on par with—or better than—humans on pristine high-resolution images, DNN performance is significantly worse than human performance on images with quality degradations, which are frequently encountered in real-world applications.

We introduce a new generative sensing framework that integrates low-end sensors with computational intelligence to attain recognition accuracy on par with that attained using high-end sensors. This generative sensing framework aims to transform low-quality sensor data into higher-quality data in terms of classification accuracy.

In contrast with existing methods for image generation, our framework is based on discriminative models and targets to maximize the recognition accuracy rather than a similarity measure. This is achieved through the introduction of selective feature regeneration in a deep neural network.

4:50 – 5:20 pm

Creating a Computationally Efficient Embedded CNN Face Recognizer

G.B. Praveen
Technical Lead,
PathPartner Technology

Face recognition systems have made great progress thanks to availability of data, deep learning algorithms and better image sensors. Face recognition systems should be tolerant to variations in illumination, pose and occlusions, and should be scalable to large numbers of users with minimal need for capturing images during registration. Machine learning approaches are limited by their scalability. Existing deep learning approaches make use of either "too-deep" networks with increased computational complexity or customized layers that require large model files.

In this talk, we explore low complexity to high complexity CNN architectures for face recognition and show that with the right combination of training data and cost functions, we can train a low-complexity CNN architecture (e.g., an Alexnet-like model) that could achieve reasonably good accuracy compared with more-complex networks. Then, we explore system-level algorithmic customization that will enable us to create a robust real-time embedded face recognition system using low-complexity CNN architectures.

TECHNICAL INSIGHTS II

TUESDAY OVERVIEW

10:40 – 11:10 am

How Simulation Accelerates Development of Self-driving Technology

László Kishonti
CEO, AImotive

Virtual testing is the only solution that scales to address the billions of miles of testing required to make autonomous vehicles robust. But integrating simulation technology into the development of a self-driving system requires consideration of several factors. One of these is the difference between real-time and fixed time step simulation. Both real-time simulation, which demands high-performance computing resources, and fixed-time step simulation, which can be executed more economically, are needed. Specialized frameworks are also needed to support the work of developers and maximize the utilization of hardware. AI researchers and developers should be able to request server-based tests from their workstations. These tests run versions of the simulator and self-driving solution and provide feedback about the effects of changes to the solution or to the simulator. Using established test scenarios and different software versions allows the simulator and self-driving solution to be used to validate each other. This prevents wasting resources on changes that degrade or destabilize either the simulator or the self-driving solution.

11:20 – 11:50 am

Computer Vision HW Acceleration for Driver Assistance

Markus Tremmel
Chief Expert ADAS,
Robert Bosch

With highly automated and fully automated driver assistance system just around the corner, next generation ADAS sensors and central ECUs will have much higher safety and functional requirements with which to cope. This directly translates into a huge uptake of the required calculation performance and hence into much higher power consumption and into system cost.

Due to the increased amount and complexity of visual sensors around the car, the embedded computer vision subsystem carries a major stake of this increase. To realize efficient, safe and affordable L3 and higher ADAS solutions, it is important to use the best possible compromise of HW acceleration between fixed logic and specialized processing architectures. In this presentation, we give an overview of the different CV methods in a next generation ADAS system and look at the trade-offs between fixed logic and programmable processing units, with an eye on the latest developments in deep learning.

12:00 – 12:30 pm

Understanding Real-world Imaging Challenges for ADAS and Autonomous Vision Systems— IEEE P2020

Felix Heide
CTO and Co-founder,
Algolux

ADAS and autonomous driving systems rely on sophisticated sensor, image processing and neural-network-based perception technologies. This has resulted in effective driver assistance capabilities and is enabling the path to full autonomy, which we see being shown at demonstration events and in controlled regions of operation. But current systems are significantly challenged by real-world operating conditions, such as darkness, poor weather and lens issues.

In this talk, we examine the difficult use cases that hamper effective ADAS for drivers and cause autonomous vision systems to fail. We explore a range of challenging scenarios and explain the key reasons for vision system failure in these situations. We also introduce industry initiatives that are tackling these challenges, such as IEEE P2020.

TECHNICAL INSIGHTS II TUESDAY

1:30 – 2:00 pm

INVITED PRESENTATION

The Roomba 980: Computer Vision Meets Consumer Robotics

Mario Munich
SVP Technology, iRobot

In 2015, iRobot launched the Roomba 980, introducing intelligent visual navigation to its successful line of vacuum cleaning robots. The availability of affordable electro-mechanical components, powerful embedded microprocessors and efficient algorithms created the opportunity for a successful deployment of a vision-enabled consumer robot.

In this talk, we discuss some of the challenges and the benefits of introducing a vision-based navigating robot and explore what made the timing right for bringing vision to Roomba. Based on the learnings gathered during development of the Roomba 980, we elaborate on the potential impact of computer vision in the future of robotics applications.

2:10 – 2:40 pm

Deep Understanding of Shopper Behaviors and Interactions Using Computer Vision

Emanuele Frontoni
Professor and
Rocco Pietrini
PhD Student, Università Politecnica delle Marche

In retail environments, there's great value in understanding how shoppers move in the space and interact with products. And, while the retail environment has some favorable characteristics for computer vision (such as reasonable lighting), the large number and diversity of products sold, along with the potential ambiguity of shopper movements, mean that accurately measuring shopper behavior is challenging.

In this talk, we explore some of these challenges and presents a set of deep-learning algorithms that we have used to address them. These algorithms have been deployed in the cloud and have been used to measure the activity of more than one million shoppers worldwide.

2:50 – 3:50 pm

Getting More from Your Datasets: Data Augmentation, Annotation and Generative Techniques

Peter Corcoran
Co-Founder, FotoNation
(an Xperi company) &
Lead PI, National University of Ireland Galway

Deep learning for embedded vision requires large datasets. The more varied training data is, the more accurate the trained network. But, acquiring and accurately annotating datasets costs time and money. In this talk, we show how to get more from existing datasets. First, state-of-the-art data augmentation techniques are reviewed, and a new approach, smart augmentation, is explained. CNN network-A vs. trained, learning optimal augmentation strategies for CNN network-B. Second, generative adversarial networks (GAN) learn the structure of an existing dataset and several example use cases show how GANs can generate "new" data corresponding to the original dataset. The example of creating a very large dataset of facial training data is presented. But, building a dataset is not the whole problem—data must be annotated in a way that is meaningful for the training process. An example of training a GAN from a dataset that incorporates "'annotations"' is given. This enables "pre-annotated data" to be generated, providing an exciting way to create large datasets at significantly reduced costs.

4:10 – 4:40 pm

Recognizing Novel Objects in Novel Surroundings with Single-shot Detectors

Alexander C. Berg
Associate Professor,
UNC Chapel Hill

Dr. Berg's 2016 work on single-shot object detection (SSD) reduced the computation cost for accurate detection of object categories to be in the same range as image classification, enabling deployment of general object detection at scale. Subsequent extensions add segmentation and improve accuracy, but still require many training examples in real-world contexts for each object category. In applications, it may be desirable to detect new objects or categories, for which many training examples are not readily available. In this talk, we consider two approaches to address this challenge. The first takes a small number of examples of objects not in context and composes them into scenes in order to construct training examples. The other approach learns to detect objects that are similar to a small number of target images provided during detection and does not require retraining the network for new targets.

4:50 – 5:20 pm

Deploying CNN-based Vision Solutions on a $3 Microcontroller

Greg Lytle
VP Engineering,
Au-Zone Technologies

In this presentation, we explains how Au-Zone designed, trained and deployed a CNN-based embedded vision solution on a low-cost, Cortex-M-based microcontroller (MCU). We describe the steps taken to design an appropriate neural network and then to implement it within the limited memory and computational constraints of an embedded MCU. We highlight the technical challenges encountered in the implementation and explains the methods used to meet design objectives. We also explore how this type of solution can be scaled across a range of low-cost processors with different price and performance metrics. Finally, we present and interpret benchmark data for representative devices.

5:30 – 6:00 pm

How to Get the Labeled Data for Free

Matt King
CTO, IUNU

Contextual machine understanding depends on labeled data and gathering this data can be a significant constraint on growth. Paying for data—and paying for labeled data—is putting the cart before the horse.

With the right strategy, you can get the customer to pay you while collecting and labeling the data for you, resulting in both a financial and technical virtuous cycle driving growth. We explain how IUNU offers commercial plant growers practical value in exchange for labeled data that drives creating more value.

WEDNESDAY OVERVIEW

10:40 – 11:10 am

The Perspective Transform in Embedded Vision

Aditya Joshi
Lead Design Engineer and
Shrinivas Gadkari
Design Engineering Director,
Cadence

This presentation, we focus on the perspective transform and its role in many state-of-the-art embedded vision applications, such as video stabilization, high dynamic range (HDR) imaging and super resolution imaging. The perspective transform accurately recomputes an image as seen from a different camera position (perspective).

Following an introduction to this transform, we explain its applicability in applications that involve joint processing of multiple frames taken from slightly different camera positions. We then discuss considerations for efficient implementation of this transform on an embedded processor. We highlight the computational challenge posed by the division operation involved in a canonical implementation of the perspective transform. Finally, we describe a division-free modification of the perspective transform that achieves a 3x reduction in processing cycles for typical use cases.

TECHNICAL INSIGHTS II WEDNESDAY

11:20 – 11:50 am

Harnessing the Edge and the Cloud Together for Visual AI

Sébastien Taylor
Vision Technology Architect, Au-Zone Technologies

Embedded developers are increasingly comfortable deploying trained neural networks as static elements in edge devices, and also with using cloud-based vision services to implement visual intelligence remotely.

In this presentation, we explore the benefits of combining edge and cloud computing to bring added capability and flexibility to edge devices. For example, an edge device can use a locally implemented neural network to address common cases at the edge and utilize larger models in the cloud for unfamiliar events, localization differences and corner cases. And, cloud resources can be used to provide updated neural network models to edge devices.

We also explore a cascaded machine learning architecture that takes advantage of both edge and cloud computing to create a system that can dynamically adapt to new conditions. Using image classification use cases, we examine the solution in detail, including system capabilities, applications and design trade-offs.

12:00 – 12:30 pm

New Deep Learning Techniques for Embedded Systems

Tom Michiels
System Architect, Embedded Vision, Synopsys

In the past few years, the application domain of deep learning has rapidly expanded. Constant innovation has improved the accuracy and speed of learning and inference. Many techniques are proposed to represent and learn more knowledge with smaller/more compact networks. Mapping these new techniques on low-power embedded platforms is challenging because they are often very demanding on compute, bandwidth and accuracy. In this presentation, we discuss new deep-learning techniques and their implications for the requirements of embedded platforms.

1:30 – 2:00 pm

INVITED PRESENTATION

Implementing Image Pyramids Efficiently in Software

Michael Stewart
Proprietor, Polymorphic Technologies

An image pyramid is a series of images, derived from a single original image, wherein each successive image is at a lower resolution than its predecessors. Image pyramids are widely used in computer vision, for example to enable detection of features at different scales.

After a brief introduction to image pyramids and their uses in vision applications, we explore techniques for efficiently implementing image pyramids on various processor architectures, including CPUs and GPUs. We'll illustrate the use of fixed- and floating-point arithmetic, vectorization, and parallelization to speed up image pyramid implementations. We'll also examine how memory caching approaches impact the performance of image pyramid code and discuss considerations for applications requiring real-time response or minimum latency.

2:10 – 2:40 pm

Architecting a Smart Home Monitoring System with Millions of Cameras

Hongcheng Wang
Senior Manager of Technical R&D, Comcast

Video monitoring is a critical capability for the smart home. With millions of cameras streaming to the cloud, efficient and scalable video analytics become essential.
To create a cost-effective smart home monitoring system that appeals to consumers, he and his team have architected and implemented a cost-effective and scalable video analytics system that detects events of interest and delivers them to consumers. To enable this, we developed a hybrid edge-cloud computing solution and innovated efficient deep-learning-based algorithms to detect events of interest. In this talk, we explore the trade-offs that informed our decisions in designing, implementing and deploying this system for millions of Xfinity Home customers.

WEDNESDAY

2:50 – 3:20 pm

Improving and Implementing Traditional Computer Vision Algorithms Using DNN Techniques

Paul Brasnett
Senior Research Manager
Vision & AI, PowerVR,
Imagination Technologies

There has been a very significant shift in the computer vision industry over the past few years from traditional vision algorithms to deep neural network (DNN) algorithms. Many companies with experience and investment in classical vision algorithms want to utilize DNNs without discarding their existing investment. For these companies, can classical vision algorithms provide insights and techniques to assist in the development of DNN-based approaches?

In this talk, we look at the similarities between classical and deep vision. We also look at how a classical vision algorithm can be expressed and adapted to become a trainable DNN. This strategy can provide a low-risk path for developers transitioning from traditional vision algorithms to DNN-based approaches.

3:30 – 4:00 pm

Hybrid Semi-parallel Deep Neural Networks (SPDNN)—Example Methodologies and Use Cases

Peter Corcoran
Co-Founder, FotoNation
(an Xperi company) and
Lead PI, National University
of Ireland Galway

Deep neural networks (DNN) are typically trained on specific datasets, optimized with particular discriminating capabilities. Often several different DNN topologies are developed solving closely related aspects of a computer vision problem. But to utilize these topologies together, leveraging their individual discriminating capabilities requires implementing each DNN separately, increasing the cost of practical solutions.

We present a methodology to merge multiple deep networks using graph contraction. This provides a single network topology, achieving significant reduction in size over the individual networks. This merged SPDNN network can be re-trained across the combined datasets used to train the original networks, improving its accuracy over the original networks. The result is a single network that is more generic, but with equivalent or enhanced performance.

Examples of several problems in contemporary computer vision are solved using SPDNNs, including eye-iris segmentation and depth mapping

TUESDAY OVERVIEW

10:40 – 11:10 am

INVITED PRESENTATION

Solving Vision Tasks Using Deep Learning: An Introduction

Pete Warden
Mobile TensorFlow Tech Lead, Google

This talk introduces deep learning for vision tasks. It is presented by Pete Warden, Google research engineer and the company's technology lead on the mobile and embedded TensorFlow team. The talk provides an overview of deep learning, explores its weaknesses and strengths and highlights best approaches to applying it to solving vision problems. The audience will learn to think about vision problems from a different perspective, understand what questions to ask and where to find the answers to these questions. The talk will conclude with insights on the challenges of deploying deep learning solutions on mobile devices.

11:20 am – 12:20 pm

From Feature Engineering to Network Engineering

Auro Tripathy
Principal Machine Learning Engineer, AMD

The availability of large labeled image datasets is tilting the balance in favor of "network engineering" instead of "feature engineering." Hand-designed features dominated recognition tasks in the past, but now features can be automatically learned by back-propagating errors through the layers of a hierarchical "network" of feature maps. We're now seeing a plethora of network topologies that satisfy such design objectives as reduced parameter count, lower compute complexity, and faster learning. Core network building blocks have emerged, such as split-transform-merge (Inception), skipping layers (Resnet, Densenet, etc.), weight-sharing across two independent networks for similarity learning (a Siamese network), and encoder/decoder network topologies for segmentation (Unet/Linknet).

In this talk, we highlight these topologies from a feature representation and learning angle and shows how they are succinctly implemented in the Keras high-level deep learning framework (with short Python code snippets). We also scrutinize model size and compute complexity of the topologies.

1:30 – 2:30 pm

Visual-inertial Tracking for AR and VR

Timo Ahonen
Director of Engineering, Computer Vision, Meta

This tutorial covers the main current approaches to solve the problem of tracking the motion of a display for AR and VR use cases. We cover methods for inside-out tracking that uses cameras and inertial sensors, such as accelerometers and gyroscopes, for both wearable head-mounted displays and for handheld devices, such as cell phones.

2:50 – 3:50 pm

Understanding and Implementing Face Landmark Detection and Tracking

Jayachandra Dakala
Technical Architect,
PathPartner Technology

Face landmark detection is of profound interest in computer vision because it enables tasks ranging from facial expression recognition to understanding human behavior. Face landmark detection and tracking can be quite challenging, though, due to a wide range of face appearance variations caused by different head positions, lighting conditions, occlusions and other factors.

In this tutorial, we introduce face landmarks and discusses some of the applications in which face landmark detection and tracking are used. We also highlight some of the key challenges that must be addressed in designing and implementing a robust face landmark detection and tracking algorithm. We survey algorithmic approaches, highlighting their complexities and trade-offs. And, we conclude with a discussion of implementation approaches for a real-time embedded face landmark tracking system.

4:10 – 4:40 pm

INVITED PRESENTATION

Building a Practical Face Recognition System Using Cloud APIs

Chris Adzima, Sr. Information Systems Analyst, Washington County Sheriff's Office

In this presentation, we will walk through the design and implementation of a face recognition system utilizing cloud computing and cloud computer vision APIs. We demonstrate how the Washington County Sheriff's Office leveraged hundreds of thousands of archived mugshots, along with off-the-shelf computer vision APIs, quickly to create a practical system that's helping law enforcement officers identify suspects.

4:50 – 5:20 pm

Bad Data, Bad Network, or: How to Create the Right Dataset for Your Application

Mike Schmit
Director of Software Engineering, AMD

When training deep neural networks, having the right training data is key. In this talk, we explore what makes a successful training dataset, common pitfalls and how to set realistic expectations. We illustrate these ideas using several object classification models that have won the annual ImageNet challenge. By analyzing accurate and inaccurate classification examples (some humorous and some amazingly accurate), you will gain intuition on the workings of neural networks. Mike's results are based on his personal dataset of over 10,000 hand-labeled images from around the world.

5:30 – 6:00 pm

Introduction to Creating a Vision Solution in the Cloud

Nishita Sant
Computer Vision Scientist,
GumGum

A growing number of applications utilize cloud computing for execution of computer vision algorithms. In this presentation, we introduce the basics of creating a cloud-based vision service, based on her experience developing and deploying a computer vision-based service for enterprises.

We explore the architecture of a cloud-based computer vision solution in three parts: an API, computer vision modules (housing both algorithm and server) and computer vision features (complex pipelines built with modules). Execution in the cloud requires the API to handle a continuous, but unpredictable, stream of data from multiple sources and task the appropriate computer vision modules with work. These modules consist of an AWS Simple Queue Service and an EC2 auto-scaling group and are built to handle both images and video.

We discuss in detail how these modules optimally utilize instance hardware for executing a computer vision algorithm. Further, we discuss our method of inter-process communication, which allows for the creation of complex computer vision pipelines that require several modules to be linked. We also address cost and run-time tradeoffs between GPU and CPU instances.

WEDNESDAY OVERVIEW

10:40 – 11:10 am

Understanding Automotive Radar: Present and Future

Arunesh Roy
Radar Algorithms Architect, NXP

Thanks to its proven, all-weather range detection capability, radar is increasingly used for driver assistance functions such as automatic emergency braking and adaptive cruise control. Radar is considered a crucial sensing technology for autonomous vehicles not only for its range finding ability, but also because it can be used to determine target velocity and target angle. In this tutorial, we introduce the basic principles of operation of a radar system, highlighting its main parameters and comparing radar with computer vision and other types of sensors typically found in ADAS and autonomous vehicles.

After examining the features and the limitations of current automotive radar systems, we discuss how automotive radar is evolving, particularly in light of safety performance assessment programs such as the European New Car Assessment Programme (eNCAP). We conclude with a discussion of how radar systems may compete with or complement vision-based sensors in future ADAS-equipped and autonomous vehicles.

11:20 am – 12:20 pm

Depth Cameras: A State-of-the-Art Overview

Carlo Dal Mutto
CTO, Aquifi

In the last few years, depth cameras have reached maturity and they are being incorporated into an increasing variety of commercial products. Typical applications span gaming, contactless authentication in smartphones, AR/VR and IoT. State-of-the-art depth cameras are based on three fundamental technologies: stereo vision, time-of-flight (ToF) and structured light (SL). These technologies are available in rather different implementations, which are characterized by specific properties in terms of measurement quality (e.g., precision, accuracy, range and resolution) and in terms of system requirements (e.g., power consumption, size and computational requirements).

In this talk, we provide an overview of the fundamentals behind stereo, ToF and SL technologies. We present measurement characterization and system requirements of different implementations and explore a selection of depth cameras available in the market.

1:30 – 2:30 pm

INVITED PRESENTATION

Approaches for Energy Efficient Implementation of Deep Neural Networks

Vivienne Sze
Associate Professor, MIT

Deep neural networks (DNNs) are proving very effective for a variety of challenging machine perception tasks. But these algorithms are very computationally demanding. To enable DNNs to be used in practical applications, it's critical to find efficient ways to implement them. In this talk, we explore how DNNs are being mapped onto today's processor architectures, and how these algorithms are evolving to enable improved efficiency. We explore energy consumption of commonly used CNNs versus their accuracy and provide insights into "energy-aware" pruning of these networks.

2:50 – 3:50 pm

Introduction to Optics for Embedded Vision

Jessica Gehlhar
Vision Solutions Engineer, Edmund Optics

This talk introduces optics for embedded vision system and algorithm developers. We begins by presenting fundamental imaging lens specifications and quality metrics. We explain key parameters and concepts, such as: field of view, f-number, working f-number, NA, focal length, working distance, depth of field, depth of focus, resolution, MTF, distortion, keystoning and telecentricity and their relationships. Optical design basics and trade-offs are introduced, such as: design types, aberrations, aspheres, pointing accuracy, sensor matching, color and protective coatings, filters, temperature and environmental considerations and their relation to sensor artifacts.

We also explore manufacturing considerations, including testing the optical components and imaging lenses in your product, and industrial optics used for a wide range of manufacturing tests. Depending on requirements, a wide variety of tests and calibrations may be performed. These tests and calibrations become important with designs that include technologies such as multi-camera, 3D, color and NIR.

4:10 – 4:40 pm

Introduction to Lidar for Machine Perception

Mohammad Musa
Founder & CEO, DeepenAI

Lidar sensors use pulsed laser light to construct 3D representations of objects and terrain. Recently, interest in lidar has grown, for example for generating high-definition maps required for autonomous vehicles and other robotic systems. In addition, many systems use lidar for 3D object detection, classification and tracking. In this presentation, we explain how lidar sensors operate, compare various lidar implementations available today and explore the pros and cons of lidar in comparison to other vision sensors.

4:50 – 5:20 pm

Designing Vision Front Ends for Embedded Systems

Friedrich Dierks
Director of Product Marketing and Development, Module Business, Basler

In this presentation, we guide the audience through the process of specifying and selecting a vision front end for an embedded system. This covers topics like selecting the right sensor, a suitable optical setup, sensor interface, real-time process coupling and, last but not least, how and where to do the image pre-processing, such as de-Bayering, white balancing, etc. While there are many experts around for processing images once they are in memory, the detailed knowledge of how to create these images in the first place is not so widely spread.
A lot of these more “analog” topics are critical for many projects’ success in terms of meeting the performance and cost targets. We give an overview of the design flow for the vision front end, address common pitfalls and describe solutions for typical applications.

5:30 – 6:00 pm

Optimize Performance: Start Your Algorithm Development with the Imaging Subsystem

Ryan Johnson
Lead Engineer, Twisthink

Image sensor and algorithm performance are rapidly increasing, and software and hardware development tools are making embedded vision systems easier to develop. Even with these advancements, optimizing vision-based detection systems can be difficult. To optimize performance, it’s important to understand the imaging subsystem and its impact on image quality and the detection algorithm. Whether performance improvement involves tuning an imaging subsystem parameter or increasing algorithm capability, it is the designer’s responsibility to navigate these relationships and trade-offs.

In this presentation, we describe a design approach that allows the designer to iteratively adjust imaging subsystem performance while increasing the fidelity of the detection algorithm. The audience will gain an understanding of several high-impact imaging subsystem noise sources, methods for evaluation and ways to determine requirements driven by the detection algorithm. The audience will also learn how datasets enable evaluation of non-obvious noise sources and system performance.

TUESDAY OVERVIEW

10:40 – 11:10 am

What's Hot? The M&A and Funding Landscape for Machine Vision Companies

Rudy Burger
Managing Partner,
Woodside Capital

The six primary markets driving computer vision are automotive, sports and entertainment, consumer and mobile, robotics and machine vision, medical, and security and surveillance. In this presentation, we examine the flow of venture money and the volume of M&A activity in each of these sectors and highlight notable transactions and trends. So far, computer vision has not created any unicorns (private companies worth over a billion dollars). Why not? Which existing computer vision companies may become unicorns over the next couple of years?

We then focus on a couple of hot sectors driving the vision market: automotive and consumer. Where are the growth opportunities for innovative private companies in these sectors? He'll highlight a few leading private computer vision companies within these two markets, and the disruptive vision technologies they are developing. We also consider which public companies are likely to be the winners (and losers) within these sectors.

11:20 – 11:50 am

INVITED PRESENTATION

Ubiquitous $90B AR to Dominate Focused $15B VR by 2022, Becoming a Major Computer Vision Market

Tim Merel
Managing Director,
Digi-Capital

Augmented reality (AR) including mobile AR and smartglasses, could approach an installed base of 3.5 billion units and revenue of $85-90 billion within 5 years, becoming one of the major computer vision consumer and enterprise markets. At the same time, virtual reality (VR–including mobile, standalone, console, and PC VR– could reach an installed base of 50-60 million and revenue of $10-15 billion. That's a big difference, and it all has to do with AR's ubiquity and VR's focus.

To understand why the sister markets are shaping up so differently, we dig into AR's and VR's installed bases, use cases, app store category revenues (IAP/premium), eCommerce category sales, ad spend by industry, enterprise revenues by industry, and geographic splits. The devil's in the details, and this is the most detailed dive ever done on these markets.

12:00 – 12:30 pm

EXECUTIVE PERSPECTIVE

Balancing Safety, Convenience and Privacy in the Era of Ubiquitous Cameras

Charlotte Dryden
Director, Intel

Computer vision-enabled cameras are proliferating rapidly and will soon be ubiquitous—in, on and around vehicles, homes, toys, stores, public transit, schools and more. This offers tremendous benefits in terms of safety, security, convenience and efficiency. But what about privacy? Are we doomed to give up our privacy as cameras proliferate? Not necessarily.

In this talk, we explore trade-offs related to privacy in a world filled with connected cameras. Many of the same technologies that are fueling the proliferation of visual intelligence can also be used to enhance privacy, if product developers choose to do so, and if consumers, enterprises and governments prioritize privacy. Accelerating innovation in sensors means that system designers have many choices of sensor types beyond the typical CMOS image sensor, enabling sensors that capture only the information required for the application. Progress in embedded processors and algorithms makes it increasingly feasible to consume images at the edge or in the fog, and then discard them, retaining only the required metadata.

BUSINESS INSIGHTS

1:30 – 2:00 pm

Reduce Risk in Computer Vision Design: Focus on the User

Paul Duckworth
Director of Engineering, Twisthink

Companies across a wide range of industries are considering ways to apply computer vision to innovate their products and services. With the vast potential of this exciting technology, a common mistake is building something that is powerful but not necessarily needed or desired. User input, when harnessed, can be a powerful antidote to this risk, driving focus into computer vision development.

Human-centered design is a process to manage the risk and complexity of developing computer vision systems by balancing three focus areas: product usability/desirability, technical feasibility, and business viability. This process identifies valuable insights to help navigate conflicting development options and ensure the final product is desirable to the end user.

In this presentation, we outline the essentials of human-centered design and how it can be used to drive focus and clarity in the development of complex vision-based products. We will illuminate the value of this approach through examples, including a wearable connected camera and a vision-based space tracker.

2:10 – 2:40 pm

Data-driven Business Models Enabled by 3D Vision Technology

Christopher Scheubel
Head of IP and Business Development, FRAMOS

In this presentation, we describe which applications are enabled by low-cost 3D vision technology, such as home robotics, smart cities/communities and drones for precision farming, and which business models can be based on the generated data. All industry verticals directly benefit from 3D vision, but even more value can be generated based on the enormous generated 3D data set. E.g., a vacuum cleaner robot will map its environment using SLAM algorithms, to maneuver intelligently around obstacles and avoid running over valuables such as jewelry.

Beyond this direct benefit, an extremely rich data set is created, containing information on apartment size and furniture. This information can be used for target marketing, as the robot has already mapped out how big the new couch could be and the furniture arrangement across many households could be used for future apartment design. We explore these kinds of data-driven business models for key industry verticals.

2:50 – 3:20 pm

EXECUTIVE PERSPECTIVE

Leveraging Edge and Cloud for Visual Intelligence Solutions

Salil Raje
Senior Vice President, Software and IP Products Group, Xilinx

For many computer vision systems, a critical decision is whether to implement vision processing at the edge or in the cloud. In a growing number of cases, designers are choosing to use both edge and cloud processing, which opens the possibility of leveraging the strengths of both approaches. But determining the best mix of edge and cloud processing for an application can be challenging because the trade-offs are often complex and subtle, involving numerous factors, such as latency, cost, bandwidth and power consumption.

In this talk, we explore the advantages of combining edge and cloud processing for visual intelligence, and outline ways that solution developers can optimize their applications for the right blend.

3:30 – 4:00 pm

From 2D to 3D: How Depth Sensing Will Shape the Future of Vision

Guillaume Girardin
Director of Photonics, Sensing and Display, Yole Développement

For several decades, 3D imaging and sensing technologies have matured, thanks to extensive, successful deployments in high-end applications, mainly in medical and industrial markets. More recently, as 3D sensor costs have dropped, 3D imaging and sensing devices have become a significant business, generating almost $2B in revenue in 2017, with strong growth prospects driven by consumer devices such as the iPhone X with its TrueDepth camera module. Currently, there are three main depth sensing approaches on the market: stereo vision, structured light, and time-of-flight. Each of these approaches has pros and cons, from a technical point of view and with respect to cost.

In this talk, we explain the key strengths and weaknesses of each of these technologies and analyze the cost of adding them to a system based on recent teardown analyses. We also give an overview of this depth-sensing sensor market and the main factors driving it and explore how this market could shape the future of vision.

BUSINESS INSIGHTS TUESDAY

4:10 – 4:40 pm

EXECUTIVE PERSPECTIVE

Embedded AI for Smart Cities and Retail in China

Kai Yu
Founder & CEO, Horizon Robotics

Over the past ten years, online shopping has changed the way we do business. Now, with the development of AI technology, we are seeing the beginning of the so-called "new retail revolution," in which nearly all China-based internet giants, such as Alibaba, Tencent and JD are active. These companies want to use big data, internet and AI technologies to transform brick-and-mortar retail. Embedded AI will play a critical role in this trend: it is an essential ingredient for extracting analyzable digital information from physical shops, and for connecting offline retail to online big data. Using embedded AI technology, cameras installed in shops can analyze customer behavior as well as interactions between customers, goods and the location real time, improving the shopping experience and operational efficiency at the same time.

In this talk, we analyze recent developments in the new retail revolution in China, and identifies key challenges that must be addressed for this trend to achieve its full potential.

5:00 – 6:00 pm

Vision Entrepreneurs' Panel

What can we learn from leaders of successful vision-based start-ups? The expanding applications of embedded vision are opening up exciting business opportunities, and countless entrepreneurs are developing diverse vision-based end-products and enabling technologies. But building a vision-based company brings unique risks and challenges.

This panel brings together an amazing group of visionary leaders who have conceived and scaled vision-based businesses. Sharing their failures as well as their successes, along with key lessons learned, these successful entrepreneurs will "pay it forward"—helping to enable the next generation of vision-based start-up leaders.

Panelists:
Nik Gagvani (Moderator)
President, CheckVideo

Radha Basu
CEO, iMerit

Gary Bradski
CTO & co-founder
Arraiy.com;
CEO & founder of
OpenCV.org

László Kishonti
CEO, AImotive

WEDNESDAY

WEDNESDAY OVERVIEW

10:40 – 11:10 am
Leveraging Cloud Computer Vision for a Real-time Consumer Product

11:20 – 11:50 am
Using Vision to Transform Retail

12:00 – 12:30 pm
Computer Vision for Industrial Inspection: The Evolution from PCs to Embedded Solutions

1:30 – 2:00 pm
Building up a Start-up in Embedded Vision: Lessons from Machine Vision

2:10 – 2:40 pm
The Four Key Trends Driving the Proliferation of Visual Perception

2:50 – 3:20 pm
Overcoming Bias in Computer Vision—A Business Imperative

3:30 – 4:00 pm
Intelligent Consumer Robots Powering the Smart Home

5:00 – 6:00 pm
Vision Tank Start-up Competition

10:40 – 11:10 am

Leveraging Cloud Computer Vision for a Real-Time Consumer Product

Pavan Kumar,
Co-Founder & CTO,
Cocoon Cam

The capabilities of cloud computing are expanding rapidly. At the same time, cloud computing costs are falling. This makes it increasingly attractive to implement computer vision in the cloud, even for cost-sensitive applications requiring real-time response.

In this presentation, we explore the benefits and limitations of computer vision in the cloud today—both for initial prototyping and for product deployment—based on Cocoon Cam's experience creating the first vision-enabled baby monitor.

11:20 – 11:50 am

Using Vision to Transform Retail

Sumit Gupta
Vice President,
AI, Machine Learning
and HPC, IBM

In this talk, we explore how recent advances in deep learning-based computer vision have fueled new opportunities in retail. Using case studies based on deployed systems, we explore how deep learning-based computer vision is enabling new ways to improve in-store consumer experiences, to gather more customer insights, to understand which types of in-store displays work better and to optimize workforce scheduling and management.

We will also touch upon how a combination of robotics and vision are improving inventory management and worker safety by automating aspects of warehouse operations.

12:00 – 12:30 pm

Computer Vision for Industrial Inspection: The Evolution from PCs to Embedded Solutions

Thomas Däubler
CTO, NET New Electronic
Technology GmbH

In this presentation, he will introduce current industrial inspection computer vision applications and solutions and explore how vision solutions are evolving for this market. In recent years, a growing number of system integrators and machine builders have recognized that embedded vision systems can often overcome key shortcomings of conventional PC-based vision systems.

We illustrate typical embedded vision inspection systems and compare them with PC-based solutions, pointing out the key advantages and limitations through several real-world examples.

1:30 – 2:00 pm

EXECUTIVE PERSPECTIVE

Building up a Start-up in Embedded Vision: Lessons from Machine Vision

Arndt Bake
Chief Marketing Officer, Basler

As embedded vision becomes more capable, it is proliferating into many new markets and attracting many new start-ups. How can you gauge when the time is right to launch an embedded vision start-up? Should you focus your business on one application, or try to address many? How should you decide when to expand your start-up into additional markets and regions?

One way to inform such decisions is to examine the experience of past start-ups in related fields. In this presentation, we briefly examine the development of the machine vision market (vision mainly for factory automation), including the main phases of the growth of this market over the past 30 years. We look at examples of successful start-up companies in machine vision, distills key learnings from these examples, and transfers the resulting success recipes to the current situation of embedded vision start-ups. Finally, we identify which market verticals are ripe for embedded vision start-ups and illustrates one representative opportunity.

2:10 – 2:40 pm

The Four Key Trends Driving the Proliferation of Visual Perception

Jeff Bier
President, BDTI
Founder, Embedded Vision
Alliance

With so much happening so fast in computer vision applications and technology it can be difficult to see the big picture. In this talk, we examine the four most important trends that are fueling proliferation of vision applications and influencing the future of the industry:

Deep learning–more than a fad, our survey of vision developers shows that deep learning is re-shaping the future of vision.

Democratization–thanks to cloud computing and higher levels of abstraction, developers now have easier access to vision than ever.

Fast, cheap, energy-efficient processors–massive investment in specialized processors is delivering 1000x improvements in performance and efficiency, enabling vision to be deployed even in very cost- and energy-constrained applications.

3D perception–the introduction of 3D optical sensors into high-volume applications like mobile phones and automobiles has catalyzed a dramatic acceleration in innovation, collapsing the size, cost and complexity of 3D perception.

We explain what's fueling each of these trends, and highlight implications for technology suppliers, solution developers and end users.

BUSINESS INSIGHTS WEDNESDAY

2:50 – 3:20 pm

Overcoming Bias in Computer Vision—A Business Imperative

Will Byrne
Entrepreneur

The rollout of computer vision applications—even by some of the world's most influential tech companies—have been marred by issues of bias against specific groups. Overcoming bias in AI-enabled products isn't just an ethical imperative, it's crucial to products and businesses that succeed across a wide range of users and contexts. In this talk, we review real-world cases of bias in products to date and outline practical steps to overcome the issue in your own product.

3:30 – 4:00 pm

INVITED PRESENTATION

Intelligent Consumer Robots Powering the Smart Home

Mario Munich
SVP Technology, iRobot

The Internet of Things has rapidly developed in the past few years, enabled by affordable electronics components and powerful embedded microprocessors, ubiquitous internet access and WiFi in the household. Connected devices such as the Nest or Canary cameras, the Ring doorbell, the Philips Hue lights, or smart TVs are commonplace in regular houses. The common control of these devices in an effort to improve the comfort of the user has proven challenging.

Several frameworks have been proposed to tackle this problem. Most of them require the user manually to place the devices in a map of the house, rendering the user interface cumbersome for regular consumers. iRobot is focused on mapping and navigation technology development to make our robots smarter, simpler to use, and provide valuable spatial information to the broader ecosystem of connected devices in the home. Robot-built and maintained home maps provide important "spatial context" by capturing the physical space of the home.

In this talk, we will describe iRobot's vision of the Smart Home, a home that maintains itself and just does the right thing in anticipation of occupant needs. This home will be built on an ecosystem of connected and coordinated robots, sensors, and devices that provides the occupants with a high quality of life by seamlessly responding to the needs of daily living–from comfort to convenience to security to efficiency.

BUSINESS INSIGHTS
WEDNESDAY

5:00 – 6:00 pm

Vision Tank Start-up Competition

The Vision Tank is the Embedded Vision Summit's annual start-up competition, showcasing the best new ventures using computer vision in their products or services. Open to early-stage companies, entrants will be judged on four criteria: technology innovation, business plan, team and business opportunity.

2018 Vision Tank judges:

Arun Chhabra
CEO, 8tree

John Feland
CEO and Founder, Argus Insights, Inc.

Liz Gasser
Managing Partner, High Camp Ventures

Vin Ratford
Executive Director,
Embedded Vision Alliance

This year's finalists include:

AiFi
AiFi is building a scalable version of "Amazon Go" to empower stores of the future to be check-out free. AiFi's innovative AI-powered sensor networks also provide retailers with valuable insights about shopping behavior and product preference, as well as improved inventory management. Presented by João Diogo Falcão, Director of Product

Aquifi
Aquifi provides visual inspection services for logistics and manufacturing, based on the combination of 3D reconstruction and deep learning. The company's solution, a trainable virtual inspection system, increases the throughput of human workers and reduces errors due to fatigue and repetition. Presented by Carlo Dal Mutto, CTO

Sturfee
Sturfee is building city-scale Visual Positioning Service (VPS) based on deep learning, computer vision and satellite imaging principles, enabling camera connected devices and machines to precisely locate themselves in the real world, identify where they are looking, and recognize what is around them—all based on visual input data. Cameras need VPS more than GPS. Presented by Anil Cheriyadat, CEO

VirtuSense Technologies
VirtuSense Technologies' product identifies people who are at risk of falls and injuries. The core technology is based on machine vision, using a 3D time-of-flight sensor to track a person's static and dynamic balance, identify sensory and muscular deficits and provide objective data to assess and treat issues. Presented by Deepak Gaddipati, Founder and CTO

Boulder AI
Boulder AI has created an intelligent GPU-enabled deep-learning neural network camera (DNNcam) that is waterproof, dust-proof and runs AI at the image source. This edge processing camera executes AI/machine learning and computer vision algorithms without the cloud, distilling visual information into actionable event data. The end-to-end Boulder AI platform enables users to effortlessly deploy models on the DNNcam and collect edge data events into cloud environments. Presented by Dan Connors, CTO

TUESDAY OVERVIEW

1:30 – 2:00 pm

Machine Learning Inference in Under 5 mW with a Binarized Neural Network on an FPGA

Abdullah Raouf
Senior Marketing Manager, Lattice

The demand for always-on intelligence is rapidly increasing in various applications. You can find cameras that are always watching for anomalies in a manufacturing line, monitoring vehicle speeds on roads or looking for a specific gesture or person. Since these cameras have to be always on, security and power consumption becomes a concern. Users don't want the captured images to be sent to the cloud (available for hackers to access) and therefore item or anomaly detection must occur locally vs. in the cloud. This increases local computational requirements, which potentially increases power consumption—a major issue for battery-powered products.

In this presentation, we provide an overview of how FPGAs, such as Lattice's iCE40 UltraPlus, are able to implement multiple binarized neural networks in a single 2 mm x 2 mm package to provide always-on intelligence without relying on cloud computation.

2:10 – 2:40 pm

EXECUTIVE PERSPECTIVE

Energy-efficient Processors Enable the Era of Intelligent Devices

Ren Wu
Founder and CEO, NovuMind

Artificial intelligence is making waves and headlines. New algorithms, applications and companies are emerging fast. Deep-learning-based systems, trained with massive amounts of data using supercomputers, are more capable than ever before.

The most important opportunity for AI is supercharging the Internet of Things, making the "things" themselves smarter. With AI, edge devices gain the ability to sense, interpret and react intelligently to the world around them—creating the Intelligent Internet of Things (IIoT).

To achieve this goal, it is essential that we make AI much more efficient, so that small, inexpensive, low-power systems can incorporate sophisticated AI to improve people's lives.

In this talk, Dr. Wu will share his perspective on the opportunity for AI at the edge, and explain how NovuMind is tackling this opportunity using domain-specific processor architectures, designed from the ground up for efficient AI, coupled with algorithms tailored for these processors.

2:50 – 3:20 pm

High-end Multi-camera Technology, Applications and Examples

Max Larin
CEO, Ximea

For OEMs and system integrators, many of today's applications in VR/AR/MR, ADAS, measurement and automation require multiple coordinated high-performance cameras. Current generic components are not optimized to achieve the desired traits in terms of resolution, frame rate, latency, imaging quality, reliability, scalability and time to market. XIMEA's xPlatform provides unique solutions to address all of these requirements, utilizing PCI Express and various image sensor technologies.

TUESDAY ENABLING TECHNOLOGIES I

3:30 – 4:00 pm

Mythic's Analog Deep Learning Accelerator Chip: High Performance Inference

Frederick Soo
Head of Product Development, Mythic

Mythic's deep learning accelerator chip uses a unique analog circuit approach to deliver massive power, speed and scalability advantages over current-generation deep learning inference technologies.

This step change in performance and efficiency enables many new high-performance, high-value, embedded deep learning applications.

4:10 – 4:40 pm

Programmable CNN Acceleration in Under 1 Watt

Gordon Hands
Director of Marketing, Solutions/IP, Lattice

Driven by factors such as privacy concerns, limited network bandwidth and the need for low latency, system designers are increasingly interested in implementing artificial intelligence (AI) at the edge. Low-power (under 1 watt), low-cost (under $10) FPGAs, such as Lattice's ECP5, offer an attractive option for implementing AI at the edge. In order to achieve the best balance of accuracy, power and performance, designers need carefully to select the network model and quantization level.

In this presentation, we use two application examples to help system architects better understand feasible solutions. These examples illustrate the trade-offs of network design, quantization and performance.

4:50 – 5:20 pm

A Physics-based Approach to Removing Shadows and Shading in Real Time

Bruce Maxwell
Director of Research, Tandent Vision Science

Shadows cast on ground surfaces can create false features and modify the color and appearance of real features, masking important information used by autonomous vehicles, advanced driver assistance systems, pedestrian guides, or autonomous wheelchairs.

We present a method for generating an illumination-independent image suitable for analysis and classification using a physics-based 2D chromaticity space. To explore its utility, we have implemented a system for removing spatial and spectral illumination variability from roads and pathways that runs at frame rate on embedded processors. The combination of physics-based pre-processing with a simple classifier to identify road features significantly outperforms a more complex classifier trained to do the same task on standard imagery, while using less computation.

Removing illumination variability prior to classification can be a powerful strategy for simplifying computer vision problems to make them practical within the computational and energy budgets of embedded systems.

5:30 – 6:00 pm

At the Edge of AI At the Edge: Ultra Efficient AI on Low-power Compute Platform

Mohammad Rastegari
CTO, XNOR.ai

Improvements in deep learning models have increased the demand for AI models in several domains. Yet these models demand massive amounts of computation, memory and power. Due to these heavy requirements, current AI applications have to resort to cloud-based solutions. However, AI applications cannot scale via cloud solutions, and sending data over the cloud is not always desired for many reasons (e.g. privacy, bandwidth). Therefore, there is a significant demand for running AI models locally on edge devices. These devices are often resource constrained; they have very limited capacity in terms of compute performance and memory. Porting state-of-the-art deep learning algorithms to these resource-constrained compute platforms is extremely challenging.

In this presentation, we introduce a technology that enables deploying state-of-the-art AI models on edge devices with extreme resource constraints. Our solution is rooted in the efficient design of deep neural networks using binary operations and network compression along with optimization algorithms for training under limited resource budgets. We provide optimized software platforms for AI models that can be deployed on a variety of low-power compute platforms.

WEDNESDAY OVERVIEW

10:40 – 11:10 am

Designing Smarter, Safer Cars with Embedded Vision

Fergus Casey
R&D Director, ARC Processors, Synopsys

Consumers, the automotive industry and government regulators are requiring greater levels of automotive functional safety with each new generation of cars. Embedded vision, using advanced neural networks, plays a critical role in bringing these high levels of safety to market.

In this presentation, we provide an overview of the latest safety standards, e.g., ISO 26262, explain how they apply to embedded vision applications and describe the technical features for which system architects should look when selecting an embedded vision processor for their safety-critical automotive ICs/SoCs.

11:20 – 11:50 am

Neural Network Compiler: Enabling Rapid Deployment of DNNs on Low-cost, Low-power Processors

Megha Daga
Senior Technical Marketing Manager, Cadence

The use of deep neural networks (DNN) has accelerated in recent years, with DNNs making their way into diverse commercial products. But DNNs consume vast amounts of computation and memory bandwidth, challenging developers seeking to deploy them in cost- and power-constrained systems. Specialized processors, such as the Cadence Tensilica Vision C5 DSP and Vision P6 DSP, address these performance and efficiency challenges. However, additional challenges arise due to time-to-market pressures, and the complexities of DNN algorithms and specialized processor architectures.

Specialized software tools are key to addressing these challenges, enabling engineers to quickly generate efficient implementations of deep neural networks to run on specialized processors without requiring detailed knowledge of the algorithms and processors. We illustrate how Cadence's Neural Network Compiler enables engineers to start from a neural network description based on a framework such as TensorFlow or Caffe and then quickly generate optimized fixed-point neural network code.

12:00 – 12:30 pm

New Memory-centric Architecture Needed for AI

Sylvain Dubois
Vice President of Marketing & Business Development, Crossbar

Artificial intelligence (AI) will not replace the need for humans any time soon, but it will have a profound impact on everyday lives, transforming industries from transportation to education, medical to entertainment. AI is about data computing, and the more data, the smarter the AI algorithms will be. However, the current bottleneck between data storage and computing cores is limiting the innovation of future AI applications.

We will discuss how non-volatile memory technologies, such as ReRAM, can be directly integrated on-chip with processing logic enabling brand new memory-centric computing architectures. The superior characteristics of ReRAM over legacy non-volatile memory technologies are helping to address the performance and energy challenges posed by AI algorithms. High performance computing applications, such as AI, require high-bandwidth, low latency data accesses across processors, storage, and IOs. ReRAM memory technologies provide significant improvements by reducing the performance gap between storage and computing.

WED ENABLING TECHNOLOGIES I

1:30 – 2:00 pm

Enabling Software Developers to Harness FPGA Compute Accelerators

Bernhard Friebe
Senior Director of Marketing, Intel

Field-programmable gate arrays (FPGAs) play a critical role as flexible, reprogrammable, multi-function accelerators. They enable custom-hardware performance with the programmability of software. The industry trend toward software-defined hardware challenges not just the traditional architectures—compute, memory, network resources—but also the programming model of heterogeneous compute platforms.

Traditionally, the FPGA programming model has been narrowly tailored and hardware-centric. As FPGAs become part of heterogeneous systems and users expect the hardware to be "software-defined," FPGAs must be accessible not just by hardware developers but by software developers.

In this presentation, we focus on a software-centric programming model that enables software developers to harness FPGAs through a comprehensive solutions stack, including optimized libraries, compilers, tools, frameworks, SDK integration and an FPGA-enabled ecosystem. We also show a real-world example using machine-learning inference acceleration on FPGAs.

2:10 – 2:40 pm

Deep Learning in MATLAB: From Concept to Optimized Embedded Code

Girish Venkataramani
Product Development Manager and
Avinash Nehemiah
Product Marketing Manager, Computer Vision, MathWorks

Learn how MATLAB can help you design deep learning based vision applications and re-target deployment to embedded CPUs and GPUs. The workflow starts with algorithm design in MATLAB, which enjoys universal appeal among engineers and scientists due to its expressive power and ease of use. The algorithm may employ deep learning networks augmented with traditional computer vision techniques and can be tested and verified within MATLAB.

Next, those networks are trained using MATLAB's GPU and parallel computing support—either on the desktop, a local compute cluster, or in the cloud. For deployment, code-generation tools automatically create optimized code that can target embedded GPUs like Jetson TX2 or DrivePX2, Intel-based CPU platforms, or ARM-based embedded platforms. The generated code is highly optimized to the chosen target platform. The auto-generated code is ~2.5x faster than mxNet and ~5x faster than Caffe2 on these platforms. We use an example of lidar processing for autonomous driving to illustrate these concepts.

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
Senior Product Manager, Xilinx

Machine learning algorithms, such as convolution neural networks (CNNs), are fast becoming a critical part of image perception in embedded vision applications in the automotive, drone, surveillance and industrial vision markets. Applications include multi-object detection, semantic segmentation and image classification. However, when scaling these networks to modern image resolutions, such as HD and 4K, the computational requirements for real-time systems can easily exceed 10 TOP/s, consuming hundreds of watts of power, which is simply unacceptable for most edge applications.

In this talk, we describe a network/weight pruning methodology that achieves a performance gain of over 10 times on Zynq Ultrascale+ SoCs with very small accuracy loss. The network inference running on Zynq Ultrascale+ has achieved performance equivalent to 20 TOP/s in the original SSD network, while consuming less than 10 W.

3:30 – 4:00 pm

NovuTensor: Hardware Acceleration of Deep Convolutional Neural Networks for AI

Miao (Mike) Li
Vice President of IC Engineering, NovuMind

Deep convolutional neural networks (DCNNs) are driving explosive growth of the artificial intelligence industry. Effective performance, energy efficiency and accuracy are all significant challenges in DCNN inference, both in the cloud and at the edge. All these factors fundamentally depend on the hardware architecture of the inference engine. To achieve optimal results, a new class of special-purpose AI processor is needed – one that works at optimal efficiency on both computer arithmetic and data movement.

NovuMind achieves this by exploiting the three-dimensional data relationship inherent in DCNNs, and by combining highly efficient, specialized hardware with an architecture flexible enough to accelerate all foreseeable DCNN structures. The result is the NovuTensor FPGA and ASIC chip, which puts server-class GPU/TPU performance into battery-powered embedded devices.

4:10 – 4:40 pm

The Journey and Sunrise Processors: Leading-edge Performance for Embedded AI

Kai Yu
Founder & CEO,
Horizon Robotics

As the nature of computation changes from logic to artificial intelligence, there's a revolution happening at the edge. A new type of processor is required for this post-Moore's-law era. Horizon Robotics, a leading technology powerhouse in embedded AI, is dedicated to providing embedded AI solutions, including algorithms, chips and cloud.

After two years of research and development, Horizon Robotics unveiled its world-leading embedded AI computer vision processors, Journey and Sunrise. Based on Horizon Robotics' Brain Processing Unit, these processors power smart cars and smart cameras.

TUESDAY OVERVIEW

1:30 – 2:00 pm

A New Generation of Camera Modules: A Novel Approach and Its Benefits for Embedded Systems

Paul Maria Zalewski
Product Line Manager, Allied Vision Technologies

Embedded vision systems typically rely on low-cost image sensor modules with a MIPI CSI-2 interface. Now, machine vision camera vendors are entering the market with so-called camera modules, some of which also rely on the MIPI CSI-2 interface standard. What is the difference between a camera module and sensor module?

In this presentation, we investigate the different architectures of camera and sensor modules, in particular the built-in image-processing hardware and software capabilities of the brand new Allied Vision 1 Product Line camera modules. By performing advanced image processing inside the camera, the 1 product line cameras minimize CPU requirements on the host side. Another major difference is the flexibility of system integration: These modules will be available with various image sensors. One camera driver installed on the host supports all sensor variants, opening new possibilities for system developers to test various sensors, offer different versions of their system or upgrade to new sensors with minimal programming effort.

TUESDAY ENABLING TECHNOLOGIES II

2:10 – 2:40 pm

Enabling Cross-platform Deep Learning Applications with the Intel CV SDK

Yury Gorbachev
Principal Engineer,
Lead Architect for Computer Vision SDK, Intel

Intel offers a wide array of processors for computer vision and deep learning at the edge, including CPUs, GPUs, VPUs and FPGAs, that allow customers to select the best platform for their specific use cases. Creating cross-platform computer vision and deep-learning applications with a high degree of portability is a challenging and resource-intensive task. Intel's edge computer vision/deep learning software ecosystem, based on the Computer Vision SDK (CV SDK), provides application designers with all of the necessary tools and components to create portable and high-performance edge computer vision/deep learning applications.

We will show examples of how the CV SDK reduces computer vision/deep learning software development time, helps to achieve cross-platform application portability and enables a high degree of code reuse across different types of edge deployments.

2:50 – 3:20 pm

Achieving High-performance Vision Processing for Embedded Applications with Qualcomm SoC Platforms

Shardul Brahmbhatt
Staff Product Manager,
Qualcomm

Advancements in machine learning are making it possible to equip devices, such as connected cameras, robots, drones and smart home solutions, with improved on-device vision processing and analytics capabilities.

We will address how a hybrid approach, combining deep learning with traditional computer vision and utilizing Qualcomm Technologies System-on-Chips (SoC) and solutions can deliver significant performance and power-efficiency improvements for embedded applications requiring vision processing. We will present the benefits and trade-offs, use cases and results from concrete implementations using this hybrid approach.

3:30 – 4:00 pm

Infusing Visual Understanding in Cloud and Edge Solutions Using State-of-the-Art Microsoft Algorithms

Anirudh Koul
Senior Data Scientist and
Jin Yamamoto
Principal Program Manager,
Microsoft

Microsoft offers its state-of-the-art computer vision algorithms, used internally in several products, through the Cognitive Services cloud APIs. With best-in-class results on industry benchmarks, these APIs cover a broad spectrum of tasks, spanning image classification, object detection, image captioning, face recognition, OCR and more.

Beyond the cloud, these algorithms even run locally on mobile or edge devices, or via Docker containers, for real-time scenarios. Join us at this talk to learn more about these APIs and tools, and how you can easily get started today to make both cloud and edge intelligent!

4:10 – 4:40 pm

Rapid Development of Efficient Vision Applications Using the Halide Language and CEVA Processors

Yair Siegel
Director of Business Development, CEVA

Gary Gitelson
VP of Engineering,
mPerpetuo, Inc

Halide is a domain-specific programming language for imaging and vision applications that has been adopted by leading technology companies. Halide provides abstractions that allow for rapid development of highly optimized and portable imaging and vision applications. When using the Halide development environment, programmers can create software with performance comparable to or better than that obtained with hand-coded assembly language or intrinsics, while achieving significantly reduced development times and improved code consistency and maintainability. CEVA and mPerpetuo have partnered to enable Halide on the CEVA-XM family of low-power IP processor cores for imaging and vision.

In this presentation, we explore using Halide with CEVA processors as a rapid and powerful method of expressing, prototyping and optimizing complex vision and imaging processing pipelines.

WEDNESDAY OVERVIEW

10:40 – 11:10 am

Deep Learning on Arm Cortex-M Microcontrollers

Vikas Chandra
Senior Principal Engineer and Director, ML, Arm

Deep learning algorithms are gaining popularity in IoT edge devices due to their human-level accuracy in many tasks, such as image classification and speech recognition. As a result, there is increasing interest in deploying neural networks (NN) on the types of low-power processors found in always-on systems, such as those based on Arm Cortex-M microcontrollers.

In this talk, we introduce the challenges of deploying neural networks on microcontrollers with limited memory and compute resources and power budgets. We introduce CMSIS-NN, a library of optimized software kernels to enable deployment of neural networks on Cortex-M cores. We also present techniques for NN algorithm exploration to develop lightweight models suitable for resource constrained systems, using image classification as an example.

11:20 – 11:50 am

Rethinking Deep Learning: Neural Compute Stick

Ashish Pai
Senior Director, Neural Compute Program, Intel

In July 2017, Intel released the Movidius Neural Compute Stick (NCS)–a first-of-its-kind USB-based device for rapid prototyping and development of inference applications at the edge. NCS is powered by the same low power, high-performance Movidius Vision Processing Unit that can be found in millions of smart security cameras, gesture-controlled drones, industrial machine vision devices and more.

Delivering significant acceleration over previous, compute-constrained platforms, the NCS is enabling all sorts of companies to embrace the revolution brought about by deep learning. We also discuss the development of this unique product, adoption of this completely new tool and share stories of companies–ranging from established multinational tech companies to ambitious startups–that have gone from prototyping on the NCS to developing full-fledged products utilizing embedded neural networks.

12:00 – 12:30 pm

Project Trillium: A New Suite of Machine Learning IP from Arm

Steve Steele
Director of Platforms, ML, Arm

Machine learning processing engines today tend to focus on specific device classes or the needs of individual sectors. Arm's Project Trillium changes that by offering ultimate scalability. While the initial launch focuses on mobile processors, future Arm ML products will deliver the ability to move up or down the performance curve, from sensors and smart speakers, to mobile, home entertainment and beyond.

This talk will cover Project Trillium, a suite of Arm IP, including ML and object detection processors and Arm NN, a SW stack supporting ML across a wide range of Arm and other hardware IP. The Arm ML processor is built specifically for machine learning, based on a highly scalable architecture that can target ML across a wide range of performance points. The Arm OD processor efficiently identifies people and other objects with virtually unlimited detections per frame. And Arm NN serves as a bridge between neural network frameworks and the underlying hardware by leveraging a library of highly optimized NN primitives.

WED ENABLING TECHNOLOGIES II

1:30 – 2:00 pm

Embedding Programmable DNNs in Low-power SoCs

Petronel Bigioi
Senior VP Engineering &
General Manager, FotoNation
(an Xperi company)

We present the latest generation of FotoNation's Image Processing Unit (IPU), an embedded, AI-enabled image processing engine that can be customized and adapted to suit a wide range of imaging tasks. Due to its scalable nature, the IPU can be deployed in low-power applications, such as IoT devices, as well as scaled up to much more powerful configurations suitable for challenging automotive computer vision applications. And, in perhaps the most exciting development, the latest variants of the IPU feature FotoNation's programmable convolutional neural network engine (pCNN), which can implement CNN architectures created using state-of-the-art design tools, such as TensorFlow and Caffe.

The pCNN hardware architecture, optimized for image analytics and state-of-the-art digital bonding interconnect technology (DBI™), can also implement multiple CNNs in parallel, being able to meet most stringent real-time requirements. The combination of the IPU and DBI™ enables advanced artificial intelligence solutions implemented on mid-sized chips, opening the door to powerful AI-driven imaging solutions to carry in your pocket.

2:10 – 2:40 pm

Exploiting Reduced Precision for Machine Learning on FPGAs

Kees Vissers
Distinguished Engineer, Xilinx

Machine learning algorithms, such as convolutional neural networks, have become essential for embedded vision. Their implementation using floating-point computation requires significant compute and memory resources. Research over the last two years shows that reducing the precision of the representations of network inference parameters, inputs and activation functions results in more efficient implementations with a minimal reduction in accuracy.

With FPGAs, it is possible to customize hardware circuits to operate on these reduced precision formats: 16-bit, 8-bit and even lower precision. This significantly reduces the hardware cost and power consumption of inference engine implementations. In this talk, we show detailed results of the accuracy and implementation cost for several reduced-precision neural networks on a set of embedded platforms. From these design points, we extract the pareto-optimal results for accuracy versus precision of both weights and activations, ranging from 16-bit to 8-bit, and down to only a few bits.

2:50 – 3:20 pm

Optimizing Your System Software and BSP for Embedded Vision and AI

Daniel Sun
VP of Intelligent Vision
Business Unit, ThunderSoft

While computer vision and AI algorithms tend to get the most attention, many other software components can have an equally important impact on image quality, algorithm accuracy and performance. For example, to get the best images from the chosen image sensor, the image signal processor must be tuned for the characteristics of the sensor and the environment. Camera device drivers also affect image quality. In addition, software component libraries and frameworks (such as frameworks for accelerating neural network inference) can speed application development, but also have a huge impact on performance and, in some cases, algorithm accuracy.

In this presentation, we explore how carefully selecting and integrating these supporting software components alongside algorithm software modules leads to systems with superior image quality, accuracy and performance.

3:30 – 4:00 pm

Pilot AI Vision Framework: From Doorbells to Defense

Jonathan Su
CEO, Pilot AI

Pilot AI's vision framework has enabled real-time detection, classification and tracking in thousands of devices, from consumer applications to federal contracts. Though diverse in end-user application, these use cases all share a common disadvantage: they are compute-constrained. Small consumer electronics are compute-constrained because BoM cost limits the amount of silicon that can be integrated, whereas federal use cases are compute-constrained since the problem (seeing from thousands of feet in the sky) requires processing a tremendous amount of data in real-time without reliable network connectivity. Scaling a single framework to enable the diverse set of hardware platforms these applications represent—from ultra-low power DSPs and microcontrollers to full-size GPUs—is what differentiates Pilot AI's vision framework.

We introduce Pilot AI's deep-learning based computer vision framework for compute-constrained devices and demonstrate this framework in real-world applications to motivate the drive toward embedded deep-learning.

Does your company provide hardware, software or services that enable computer vision?

Or does your company develop end products that use computer vision?

If so, we should talk.

The Embedded Vision Summit offers you a chance to catch up on the latest trends, technologies and developments in practical computer vision and to connect with vendors, customers, prospects and partners.

But the Summit only happens once a year. What if there was a way to gain access to that network, expertise and insight year round?

There is.

Becoming an Embedded Vision Alliance Member gives you continuous access to the most relevant customers, suppliers, partners and data to fuel the success of your vision-related business.

To learn more about membership in the Embedded Vision Alliance, or email us at membership@embedded-vision.com.

embedded VISION ALLIANCE

Join the Embedded Vision Alliance and be among the industry's leading innovators:

SPECIAL THANKS TO OUR

PREMIER SPONSORS

Allied Vision

intel

LATTICE SEMICONDUCTOR

SYNOPSYS®

GOLD SPONSORS

arm

BDTi

cadence®

地平线 Horizon Robotics

NOVUMIND Making Things Think

QUALCOMM®

XILINX

SILVER SPONSORS

AIMOTIVE

luum

LUXOFT

nextchip

NXP

PARTNERS

tribune

SPECIAL EVENT SPONSORS

Google Cloud

TANDENT.