

The logo for the 2024 Embedded Vision Summit is centered on the left side of the slide. It features a white octagonal background with a colorful, multi-layered border in shades of purple, blue, green, yellow, and orange. The text "2024" is at the top, "embedded" is below it, "VISION" is in large, bold, dark blue letters with a gradient, and "SUMMIT" is at the bottom.

2024
embedded
VISION
SUMMIT®

OpenCV for High-Performance, Low-Power Vision Applications on Snapdragon

Xin Zhong

Director, Product Management,
Qualcomm Technologies

Qualcomm

- Rich applications using OpenCV
- Accelerating OpenCV on Snapdragon® platforms
- Heterogenous computing delivers optimal performance on Snapdragon
- Leveraging CPU, DSP, GPU and EVA
- OpenCV modules we plan to optimize on Snapdragon
- Performance comparison
- One Example: accelerating optical flow using Snapdragon
- Roll-out schedule
- Benefits for developers on Snapdragon

OPENCV LIBRARY



70k Stars on Github

Extremely popular Github repo

github.com/opencv/opencv



Installs per month

Rivals Tensorflow in Python installs using **pip**

pypi.org/packages/opencv-python



Embedded vision engineers

The most popular computer vision library among embedded vision engineers.

embedded-vision.com/academy/ToolsAndProcessorsForCV_Jan2019_R1.pdf




Algorithms

The largest collection of my computer vision algorithms in a single library

github.com/opencv/opencv

Rich applications using OpenCV

- Object detection and tracking
- Security and video surveillance
- Body and pose estimation
- Face detection, face recognition
- Camera calibration and depth
- Segmentation, background subtraction
- Computational photography
- Scene classification
- Robotics
- Smart City



Main modules

Core math	Image processing	Video I/O	High-level GUI
Video Analysis	Camera calibration/3D reconstruction	2D features	Object detection
Deep neural network (DNN)	Machine learning	Clustering and search	Computational photography
	Image stitching	Graph API	

Extra modules

Alpha matting	Segmentation	Deformable part	Face analysis
Fuzzy mathematics	Optical flow	3D object tracking	RGB-D processing
Saliency	Structured light	Super resolution	Video stabilization
Video encoding & decoding	Image filtering	3D Object detection	

Accelerating OpenCV on Snapdragon® platforms


Math	Image processing	Image transform	Color conversion	Feature detection	Descriptor matching	3D reconstruction	Motion tracking	Shape & drawing	Clustering
Scalar	Image filters	Resizing	RGB-YCC	Corner detector	Gradient based	Camera pose estimation	LK optical flow	Convex fill	K-Means
Vector	Image integrate	Affine transform	Interleave – planar	Contour detector	NCC	Perspective transform fitting	BM optical flow	Find convex hull	
Matrix	Pyramid	Remap	RGB-HSV	Circle/line detector	MSER	Homography fitting	Mean shift	Draw contour	
Linear solvers	Stats	FFT Inv FFT	NV21-YUV	Good feature to track	SAD	Depth estimation	Cam shift		
		Rotate		KD Tree			SGM optical flow		
		Grid warping							

Key


Software acceleration

Hardware acceleration

Heterogenous computing delivers optimal performance



Heterogenous computing delivers optimal performance



OpenCV 4.9.0 modules we plan to optimize on Snapdragon® platforms

Main modules

Extra modules

Core math	Image processing	Video I/O	High-level GUI	Alpha matting	Segmentation	Deformable part	Face analysis
Video analysis	Camera calibration/3D reconstruction	2D features	Object detection	Fuzzy mathematics	Optical flow	3D object tracking	RGB-D processing
Deep neural network (DNN)	Machine learning	Clustering and search	Computational photography	Saliency	Structured light	Super resolution	Video stabilization
	Image stitching	Graph API		Video encoding & decoding	Image warping	Image filtering	Object detection

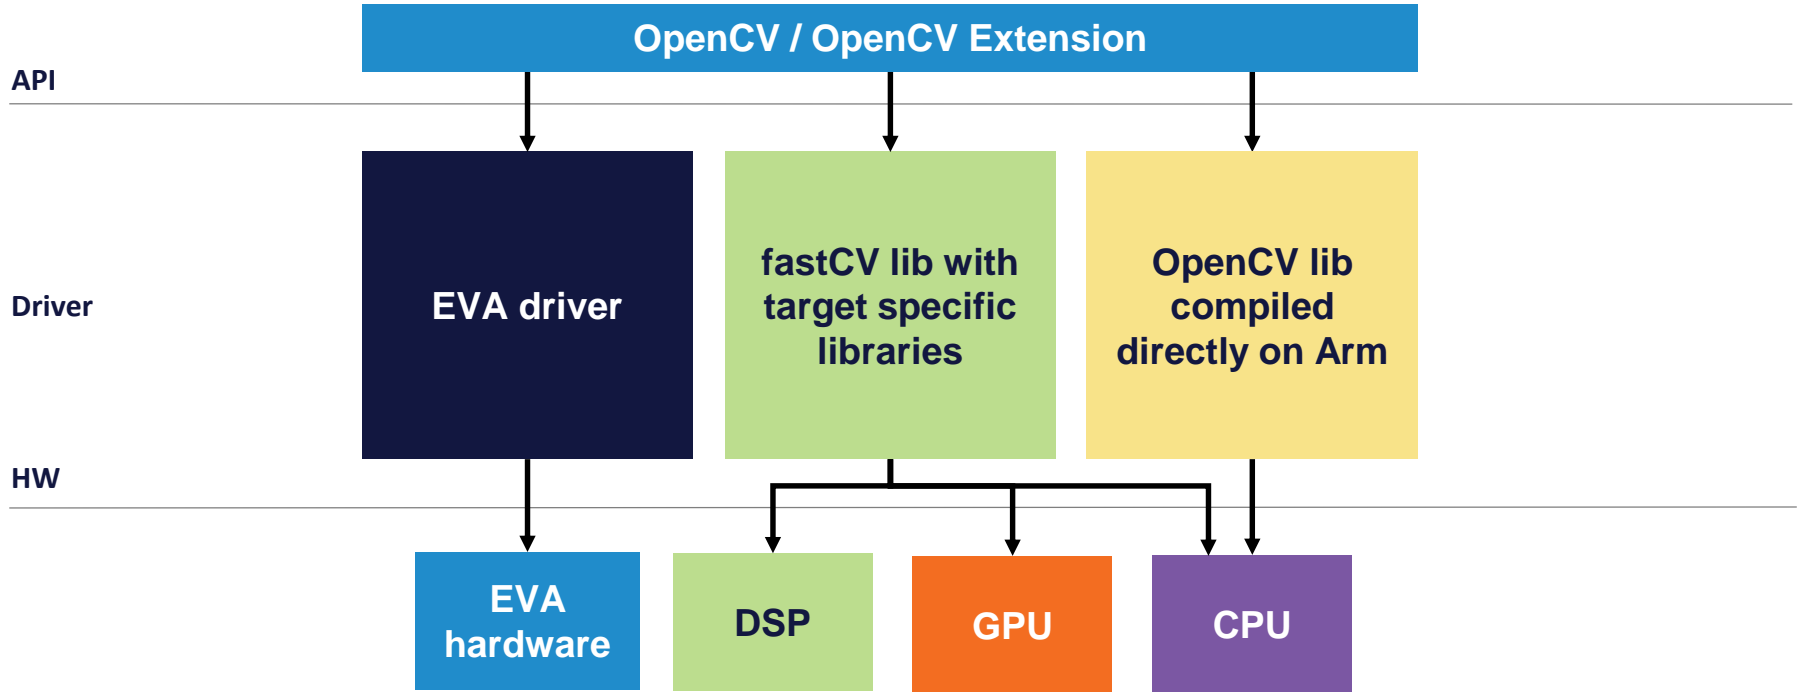
Non-optimized

Optimized or New

Non-optimized

Optimized or New

Leveraging fastCV and EVA hardware



Performance Comparison – summary

	OpenCV	Qualcomm optimized		
	Native implementation	Arm compatible	Neon optimized (Arm compatible)	DSP optimized
Min Speed	1x	0.4x	0.5x	0.2x
Max Speed	1x	34.2x	79x	53x
Median Speed	1x	1.2x	2.3x	1.2x
Average Speed	1x	9.0x	24x	11x

All performance numbers measured on Snapdragon 8 Gen1
Statistics collected on a subset of OpenCV library that's accelerated on Snapdragon

Performance Comparison – selected functions

	OpenCV	Qualcomm optimized			
	Native implementation	Arm compatible	Neon optimized (Arm compatible)	DSP optimized	EVA hardware implementation
Gaussian 3x3	1x	3.2x	6.6x	3.5x	-
Gaussian 5x5	1x	25x	25.4x	24x	-
Mean shift	1x	1.2x	3.6x	3.2x	-
Cholesky solve	1x	9.0x	12.8x	4.1x	25x
Bilinear downscale	1x	1.7x	1.7x	0.6x	5.5x
Remap bilinear	1x	1.0x	1.1x	0.4x	11.8x

One example: accelerating optical flow using Snapdragon

- SGM-based optical flow has been available in Snapdragon EVA hardware for generations
- MV quality is best-in-class
- Plan to add OpenCV extension API so developers can use it in similar ways of other Optical Flow functions
- Performance of EVA-based OF is about **5x faster** than LK optical flow in OpenCV



- **Hardware acceleration**
 - Better support for CPUs and GPUs
 - Arm support and added RISC-V support
 - Major revision to CPU HAL API
 - Optimizations for low powered devices
 - Substantially faster image processing
- **API overhaul**
 - C API will be deprecated
 - C++17 as the minimum required version
 - Revision of basic modules, such as Core, Imgproc, Features2D, Objdetect etc
- **Deep Learning module**
 - Major revision of DNN module with much better ONNX support
 - Support for major transformer-based models
- **3D module**
 - Improved camera calibration including multi-camera calibration.
 - Improved and extended 3D vision in a dedicated 3D module.
- **Better documentation**, improved samples and tutorials

OpenCV roll-out schedule



March 2024

Phase I

OpenCV running
on CPU



June 2024

Phase II

OpenCV optimized
in CPU/DSP/GPU,
leveraging fastCV



June 2025

Phase III

OpenCV accelerated
by EVA HW through
OpenCV or OpenCV
extensions

Benefits for Snapdragon developers

- Enabling OpenCV APIs on Snapdragon bridges the gap between CV developers and steep learning curve of Qualcomm Technologies' complex computing platform
- Easily accessible through OpenCV and simple integration
- Large number of CV algorithms with extensive OpenCV documentation and acceleration
- Commonly used OpenCV functions are accelerated up to **79x** using Snapdragon **heterogeneous computing**, including **CPU, DSP and GPU**
- Most critical CV functions are directly **accelerated by EVA hardware**, to achieve up to **4x power saving** and up to **25x performance**

OpenCV

<https://opencv.org>

Qualcomm CV SDK

<https://developer.qualcomm.com/software/qualcomm-computer-vision-sdk>

Viso.ai

<https://viso.ai/computer-vision>

Thank you!