

labforge

BOTTLENOSE™



Applications

Manufacturing

- Industrial automation
- Food & beverage processing
- 3D machine vision
- Industrial IoT
- Object detection
- Object classification

Robotics

- AgTech
- Service robots
- Delivery robots
- Warehouse robots
- Assistive devices & healthcare robots
- Localization
- 3D scanning

Defense & Security

VRU detection

Human and animal biomechanics

- Gait analysis

Bottleneck™ is a family of high resolution smart cameras for industrial automation and robotics applications. It is available in monocular and stereo versions. The stereo version includes dual image sensors, and both monocular and stereo versions have hardware synchronized triple-axis gyroscopes, accelerometers, and magnetometers. **Bottleneck™** can be connected over standard ethernet using GigE Vision 2.1. It has on-camera simultaneous processing for HDR, feature point detection & matching, dense disparity, and AI.

Summary



Feature	Monocular	Stereo
Baseline	N/A	134mm
Image Sensor	1x Sony IMX577, 1/2.3", 12.3MP	2x Sony IMX577, 1/2.3", 12.3MP
Image Resolution	3840 x 2160 pixels (4K / 8MP). Binning at 2x2 and 4x4 to give 1920x1080 and 960x540 resolutions.	3840 x 2160 pixels (4K / 8MP). Binning at 2x2 and 4x4 to give 1920x1080 and 960x540 resolutions.
Depth Range*	N/A	~0.1-100+m
Shutter	Rolling	Rolling
Image Sensor Features	Binning, flip, analog gain control	Binning, flip, analog gain control
ISP Functions	Color correction, gain, gamma, white balance, black level, hue, sharpness, saturation, pixel & lens shading correction	Color correction, gain, gamma, white balance, black level, hue, sharpness, saturation, pixel & lens shading correction
IMU	TDK InvenSense ICM20689 & MEMSIC MMC5603NJ	TDK InvenSense ICM20689 & MEMSIC MMC5603NJ
Field of View (FOV)	Flexible FOV via interchangeable lenses	Flexible FOV via interchangeable lenses
Power Supply	12V	12V
Horsepower	20.5TOPS	20.5TOPS
Interfaces	GigE Vision 2.1, serial port, digital I/O	GigE Vision 2.1, serial port, digital I/O
AI Processed On-Camera	Segmentation, classification, detection, pose. For example YOLOv3, SSD, and others are supported.	Segmentation, classification, detection, pose. For example YOLOv3, SSD, and others are supported.
Dense Depth Processed On-Camera	Monocular depth estimation via the latest neural network based methods	Rectified disparity images or fully triangulated dense point cloud via Semi-Global Matching (SGM)
High Dynamic Range (HDR) Processed On-Camera	1 HDR compressed image via dual or triple exposure & local tone mapping	2 HDR compressed images via dual or triple exposure & local tone mapping
Feature Point Detection and Matching Processed On-Camera	FAST, GFTT, A-KAZE, Hamming distance. Output as feature point pixel locations, descriptors, or temporal matching distances.	FAST, GFTT, A-KAZE, Hamming distance. Output as feature point pixel locations, descriptors, or a fully triangulated sparse point cloud.
Processor	16nm chip with 10 processors, four DSPs, and eight types of accelerators	16nm chip with 10 processors, four DSPs, and eight types of accelerators
Connectors	RJ45 ethernet with mounting holes for GigE Vision cables, 2.1mm ID & 5.5mm OD barrel connector for power, Molex 0530480510 for serial port and digital I/O.	RJ45 ethernet with mounting holes for GigE Vision cables, 2.1mm ID & 5.5mm OD barrel connector for power, Molex 0530480510 for serial port and digital I/O.
Lens Mount	CS-Mount	CS-Mount
Mounting Mechanism	TBD	TBD
Size	TBD	TBD
Weight	TBD	TBD
IP Rating	TBD	TBD
Temperature Range	Temperature: -20°C to 75°C, pending testing.	Temperature: -20°C to 75°C TBD, pending testing.

* Values are approximate and dependent on lens selection.




Disparity is used to calculate differences between a left image and a right image for determining depth. On **Bottlenose™** this process begins by capturing synchronized images and then undistorting and rectifying them. Each pixel in the two images is compared to find the nearest match. Labforge uses a version of the well known Semi-Global Matching (SGM) algorithm. **Bottlenose™** computes this at up-to 4K resolution with a 200+ MP/s internal throughput. This ground breaking performance allows faster depth, wider angles, and/or longer range coverage. Frame rate is limited by user configuration, settings, and ultimately by the 1Gigabit/s output over Ethernet.

Results are output as rectified disparity images or a fully triangulated dense point cloud.

In a traditional setup, this level of 3D disparity generation would be done on an external host PC with a very large GPU or FPGA. **Bottlenose™** does not need a large PC for operation.

Feature Detection is a method for fingerprinting interesting areas in an image. **Bottlenose™** computes corner points in the left and right images, and also corner points in images from a previous time step. It detects points using FAST and GFTT. Built-in hardware acceleration is used to calculate A-KAZE descriptors for each point and these are output as GigE Vision's chunk-data or contours. Optionally these points can also be output as a fully triangulated sparse 3D point cloud.

Customers can use this front-end processing result to build their own pick & place, automated harvesting, SLAM, navigation, and obstacle avoidance solutions.



Hardware Acceleration & On-Camera Processing



AI in imaging applications usually requires the use of deep neural networks. **Bottlenose™** has hardware acceleration for AI inference for neural network models for example YOLOv3, SSD, and others. Labforge provides pre-trained models that run out of the box. Users can also load their own models into the camera, provided that the layers are supported. Dataset collection, labelling, and model training services are available as part of support.

Bottlenose™ has 4GB of on-board LP-DDR4 memory and 320MB of flash storage to cater to large model sizes. This allows on-camera computation of full scale neural networks without the need for quantization and pruning.

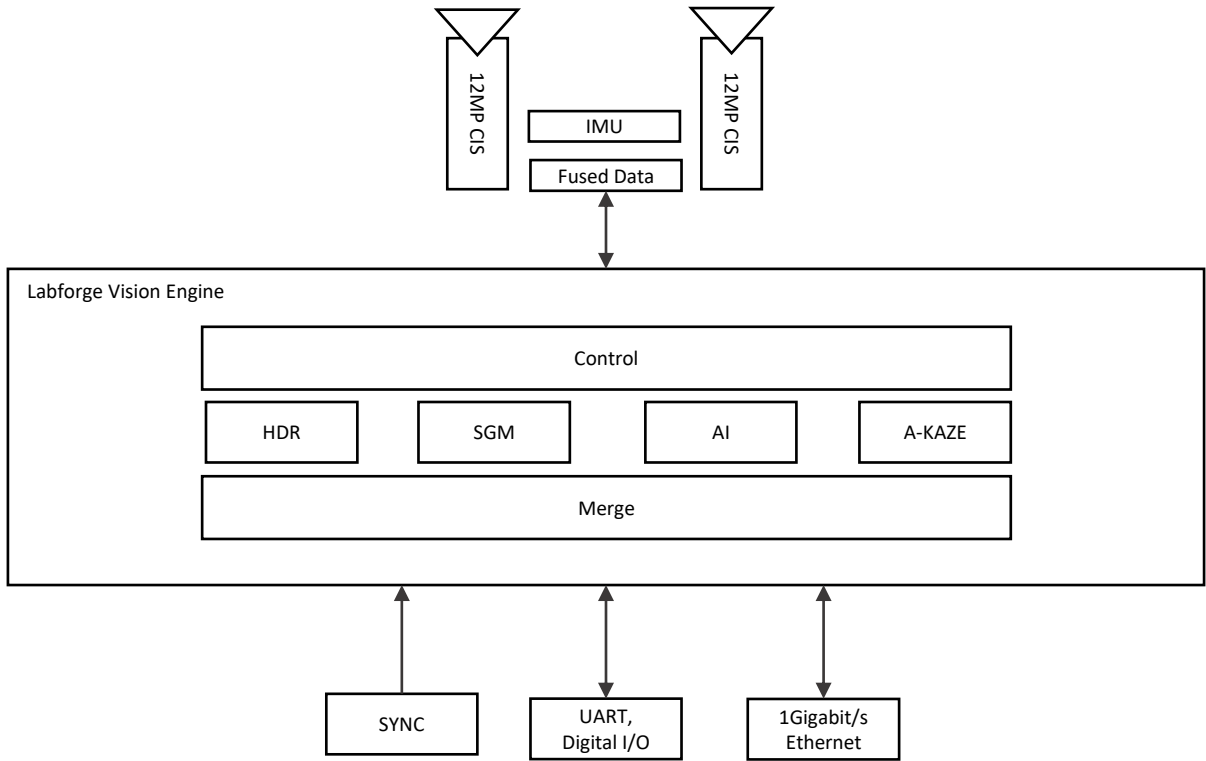
Applications include segmentation, classification, detection, pose detection, and monocular depth estimation. These results are output over GigE Vision 2.1 and can be read via chunk data or as contours.

HDR combines low, mid, and high exposure image frames to create images with larger bit-depth. **Bottlenose™** has hardware acceleration for HDR and local tone mapping. This allows **Bottlenose™** to excel in simultaneous low light and bright light situations. These computations are done on-camera.

HDR processed images can be acquired from **Bottlenose™** using GigE Vision 2.1. Frame rate is limited by user configuration, settings, and ultimately by the 1Gigabit/s output over Ethernet.




Hardware Acceleration & On-Camera Processing



Example Pipeline, Configured via GigE Vision 2.1.

Processing it all simultaneously

Bottlenose's™ Vision Engine does all these great things at the same time, while other smart cameras may be configured to do one or the other in their pipelines. Users can configure the pipeline for their custom applications using the GigE Vision 2.1 parameters. For example, users can turn off AI and change the cross-check settings on SGM. All known parameters for the computer vision algorithms are visible over the GigE interface.



Scalable

In machine vision applications, **Bottlenose's™** front-end vision processing allows users to use smaller industrial computers. These generate less heat, are more readily available, and are better suited to factory environments. The AI and computer vision features on-camera also don't require additional software licenses.

In mobile robotics, front-end processing frees up the central processor for path planning, Simultaneous Localization and Mapping (SLAM), visual inertial odometry, more complex neural networks, and more.

Multiple **Bottlenose™** cameras can be used per assembly line, robot, or heavy machinery, without having to proportionally scale the horsepower on the central processor.

Bottlenose™

GigE Vision 2.1
Images
Chunk data
Contours
Parameters

**Processor, PLC, Industrial
PC, Next Gen PLC, MCU.**

- Hardware synchronized monocular or stereo image sensors + IMU.
- Changeable CS-Mount lenses.
- Digital I/O for serial port and sync.
- High dynamic range combination and local tone mapping.
- Dense depth computation.
- Feature point detection and tracking.
- AI for segmentation, classification, detection, and mono-depth.

Output with On-Camera Compute

- 2x, 3x exposure **HDR** combination up-to **4K** images.
- **Dense disparity / depth.**
- **Feature points** and matches spatial & temporal.
- **Neural network** results.
- Accurate **timing** data.

- No minimum requirements.
- C/C++ API.
- GigE Vision 2.1 protocol compliant: connect to popular vision libraries like Halcon & VisionPro.
- Camera control.
- Use cases: industrial automation, manufacturing, automated harvesting, crop monitoring, phenotyping, apiculture studies, R&D labs, aerial surveys, wildlife monitoring, restaurant robots, cleaning robots, warehouse robots, delivery robots, healthcare robots.

System Integration Overview

Bottlenose™ is connected to a host system via GigE Vision 2.1 or digital I/O. This gives users the freedom to customize the parameters for smart features to best suit their use case.

Hardware




Image Sensor

Sony IMX577, 1/2.3", 12.3MP, back-illuminated, stacked CMOS image sensors, with 1.55 μ m x 1.55 μ m pixel size.

Excellent low-light capability for fast exposure. Minimizes any motion artifacts.

Monocular: 1x with 3840 x 2160 pixels (4K / 8MP) useable.

Stereo: 2x hardware synchronized, with 134mm baseline. 3840 x 2160 pixels (4K / 8MP) usable.

Monocular and stereo both support binning at 2x2 and 4x4 to give 1920x1080 and 960x540 resolutions. This allows use of lower LP/mm lenses. Lower resolutions also allow faster image acquisition.

Timing

IEEE1588 hardware timestamp (in MAC).

Sync input.

Inertial Measurement Unit

TDK InvenSense ICM20689 with advanced Kalman Filters to fuse accelerometer and gyroscope. Synchronized with the image sensors.

TBD MSPS (million samples per second).

MEMSIC MMC5603NJ 3-axis magnetometer.

Raw or quaternion output.

Compute

20.5 TOPS multi-core processor built on a 16 nm process.


10 processors, four DSPs, and eight hardware accelerators including the Neural Network accelerator and ISP.

4GB RAM.

320MB Flash.

Leading performance per Watt for a smart camera.

Intelligence



AI

Inference for neural network models for example YOLOv3, SSD, and others.

Use cases include segmentation, classification, detection, pose estimation, and monocular depth.

Optional: dataset collection, labelling, and training services available.

Feature Detection and Matching

Detect points with FAST and GFTT.

Compute descriptors with A-KAZE at 120, 128, 256, or 486 bits long.

Match in spatial or temporal frames with Hamming distance.

Stereo Dense Depth

134mm baseline. Based on Semi-Global Matching (SGM). Max disparity levels: 128 and up-to 512 with on-camera binning or pyramiding.

Subpixel refinement with 256 steps between each disparity level.

Computes dense disparities at a rate of 200MP/s (internal processing throughput). For example, computes 4K resolution disparities in 40ms. Lower resolutions are computed at a much faster frame rate.

Output limited via 1Gigabit ethernet. Enables wide angle and/or long range depth maps.

Stereo dense depth is only available on the stereo version of **Bottlenose™**.

Format

Color: YUV 4:2:2 unsigned.

Monochrome: Mono8 8-bit unsigned.

Stereo images supported in monochrome and color as multi-part format in either YUV 4:2:2 or Mono8.

Depth: Selectable as 8- or 16-bit unsigned monochrome format.


HDR

Double or triple exposure HDR combination and compression. HDR compression uses a bilateral filter and local tone mapping.


Best suited for uncontrolled or challenging lighting environments and areas where bright and dark conditions will be encountered in the same setting.

Produces 1 (monocular) or 2 (stereo) up-to 4K HDR compressed images.


Mechanical



Stereo



Monocular



Labforge **Bottlenose™** includes mounting holes for easy set up. Locations and format of mounting holes is TBD.

Bottlenose™ uses the CS-Mount lens format which allows the user to choose between CS, C, or S-Mount lenses with an adapter.

Field of view and depth range is configurable via lens selection. GigE Vision 2.1 is used to set extrinsic and intrinsic parameters on the camera for use in un-distortion, rectification, and stereo calculations.

Operating environment is -20C to 75C, pending testing.

Bottlenose™ features a fanless and rugged design.

Electrical



Power Supply

12V input.

Consumption: TBD, less than 15W.

Power connector accepts a barrel with dimensions: 2.1mm ID 5.5mm OD. Widely available.

Power Adapter

TBD – Recommended Mouser part numbers will be added here.

ESD Protection

Built-in.

Optics



Lenses and Field of View

Flexible field of view via interchangeable lenses. CS-Mount lens format.

Adapter for S-mount and C-mount lenses available.

On-camera undistort functions with ability to load tangential and radial parameters.

On-camera stereo rectification with ability to load extrinsic and intrinsic calibration results.

Calibration tutorials provided.

Integration



API and SDK

GigE Vision 2.1 based API. Platform independent. Allows users to configure the processing pipeline and parameters like resolution, AI, feature detection, pyramiding, un-distortion, SGM, etc.

On-camera SDK also available in future, targeted for Q2 2023 that allows developers to create and run their own applications.

Data Output

Via GigE Vision 2.1.

Serial port and digital I/O also available in future, targeted for Q2 2023.

Supported Platforms and Languages

GigE Vision 2.1 for integration into Halcon, VisionPro, and others.
C++.
OpenCV.

Ordering Information

BN-4K-4GB-320MB-CS-M

Monocular, 4K, 4GB RAM, 320MB flash, CS-Mount.

Order via Mouser.


BN-4K-4GB-320MB-CS-ST

Stereo, 4K, 4GB RAM, 320MB flash, CS-Mount.

Order via Mouser.

Networked Tracking System

Labforge's Integrated Classification and Tracking Network (ICTN) product is used to track objects in 3D space without the use of markers. Long range and/or wide coverage system configurations available. Please see ICTN product material for more details.



Google, Sony, TDK, InvenSense, Halcon, VisionPro, CAT, Caterpillar, GigE Vision, and OpenCV are registered trademarks of their respective owners. TM and (R) are not mentioned in each case in this datasheet.

Accelerated-KAZE (A-KAZE) was first proposed by Pablo F. Alcantarilla, Jesús Nuevo, and Adrien Bartoli in "Fast Explicit Diffusion for Accelerated Features in Nonlinear Scale Spaces", BMVC, 2013

YOLOv3: An Incremental Improvement by Joseph Redmon and Ali Farhadi, University of Washington, 2018


SGM: H. Hirschmuller, "Accurate and efficient stereo processing by semi-global matching and mutual information," 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)

Hamming distance: Hamming, Richard W. "Error detecting and error correcting codes." The Bell system technical journal 29.2 (1950): 147-160.

FAST: Förstner, Wolfgang, and Eberhard Gülch. "A fast operator for detection and precise location of distinct points, corners and centres of circular features." Proc. ISPRS intercommission conference on fast processing of photogrammetric data. Vol. 6. 1987.

GFTT: Shi, Jianbo. "Good features to track." 1994 Proceedings of IEEE conference on computer vision and pattern recognition. IEEE, 1994.

SSD: Liu, Wei, et al. "SSD: Single shot multibox detector." European conference on computer vision. Springer, Cham, 2016.



Revision History

Date (MM/DD/YY)	Description
02/02/22	Melixis to MEMSIC typo error.
02/11/22	Removed PoE. Slated for future projects – no part availability currently.
03/14/22	Added info on processor + RAM.
03/21/22	Updated languages and platforms.
04/11/22	Added new photos, edited summary table, removed platform and language logos due to copyright. Part number changes: C->CS, 340->320.