

Advanced AI Vision Processing Using AM68A for Industrial Smart Camera Applications



Do-Kyong Kwon, Karan Saxena, and Manisha Agrawal

ABSTRACT

The advances of deep-learning based artificial intelligence (AI) and embedded processors makes camera-based analytics a crucial technology in many industrial applications where the data from multiple cameras must be processed with high performance at low power and low latency. The AM68A processor provides various ways to optimize the performance of AI applications at the network edge with heterogeneous processing cores and remarkable integrated hardware accelerators. This processor is designed for edge AI with as many as eight cameras. The Edge AI SDK and tools provided make the development of the AI applications on AM68A simpler and faster while fully taking advantage of hardware accelerators for vision and AI processing.

Table of Contents

1 Introduction	2
2 AM68A Processor	2
3 Edge AI Use Cases on AM68A	4
3.1 AI Box.....	4
3.2 Machine Vision.....	5
3.3 Multi-Camera AI.....	6
4 Software Tools and Support	7
4.1 Edge AI Software Development Kit (SDK).....	7
4.2 Edge AI SDK Demonstrations.....	8
4.3 Edge AI Model Zoo.....	8
4.4 Edge AI Studio.....	9
5 Conclusion	9
6 Reference	9

List of Figures

Figure 2-1. AM68A Block Diagram With Subsystems.....	3
Figure 3-1. AI Box Block Diagram With Data Flow on AM68A.....	4
Figure 3-2. Machine Vision Block Diagram With Data Flow on AM68A.....	5
Figure 3-3. Multi-Camera AI Block Diagram With Data Flow on AM68A.....	7
Figure 4-1. Processor SDK Linux for AM68A.....	8

List of Tables

Table 2-1. Inference Performances of Classification, Object Detection, and Semantic Segmentation Models on AM68A.....	3
Table 3-1. AM68A Resource Utilization and Power Consumption for the AI Box Use Case.....	5
Table 3-2. AM68A Resource Utilization and Power Consumption for the Machine Vision Use Case.....	6
Table 3-3. AM68A Resource Utilization and Power Consumption for the Multi-Camera AI Use Case.....	7

Trademarks

Arm® and Cortex® are registered trademarks of Arm Limited.
All trademarks are the property of their respective owners.

1 Introduction

As vision is a primary sensor for human beings, machines also use vision to perceive and comprehend environments around them. Camera sensors provide rich information on surroundings and the advance of deep-learning based AI makes it possible to analyze enormous and complex visual data with higher accuracy. Therefore, in the applications like machine vision, robotics, surveillance, home and factory automation, camera-based analytics has become a more powerful and important tool.

Embedded processors (EP) with AI capability, that is, edge AI processors, are accelerating this trend. EP can process visual data from multiple cameras into actionable insight by mimicking the eyes and brain of a human. In contrast to cloud-based AI, where deep neural network (DNN) inference is running on the central computing devices, edge AI processes and analyzes the visual data on the systems, for example, edge AI processors, directly connected to the sensors. Edge AI technology not only makes existing applications smarter but also opens up new applications that require intelligent processing of large amounts of visual data for 2D and 3D perception.

Edge AI is specifically designed for time-sensitive applications. However, edge AI requires a low-power processor to process multiple vision sensors and execute multiple DNN inferences simultaneously at the edge, which presents challenges in size, power consumption, and heat dissipation. These sensors and processor must fit in a small form factor and operate efficiently under the harsh environments of factories, farm and construction sites, as well as inside vehicles or cameras installed on the road. Moreover, certain equipment such as mobile machines and robots necessitate functionally safe 3D perception. The global market for such edge AI processors was valued at \$2.1 billion in 2021 and is expected to reach \$5.5 billion by 2028⁽¹⁾.

This paper focuses on the highly-integrated AM68A processor and several edge AI use cases including AI Box, machine vision, and multi-camera AI. Optimizing the edge AI systems using the heterogeneous architecture of the AM68A with the optimized AI models and the easy-to-use software architecture is also discussed.

2 AM68A Processor

The AM68A is a dual-core Arm® Cortex® A72 microprocessor. The processor is designed as a high-performance and highly-integrated device providing significant levels of processing power, image and video processing, and graphics capability. Compared with the AM62A⁽²⁾, which is designed for the applications with one or two cameras, the AM68A enables real-time processing of four to eight 2MP cameras with improved AI performance. [Figure 2-1](#) shows the following multiple sub-systems based on the heterogeneous architecture of the AM68A:

- **A dual-core Arm Cortex A72 microprocessor** at 2 GHz provides up to 25K Dhrystone Million Instructions Per Second (DMIPS).
- **Vision Processing Accelerator V3 (VPAC3)** performs image processing in Vision Image Sub-System (VISS) to support raw image sensors through de-mosaic, defective pixel correction, auto exposure, auto white balance, chromatic aberration correction (CAC), and so forth. In addition, VPAC3 includes Lens Distortion Correction (LDC), Multi-Scalar (MSC), and Bilateral Noise Filter (BNF) hardware accelerators (HWAs) to accelerate correction of distorted images, down scaling of images into multiple resolutions and noise filtering, respectively. VPAC3 in the AM68A can process 600 MP per second (MP/s) when assuming 20% system overhead.
- **Digital Signal Processing (DSP) and Matrix Multiplication Accelerator (MMA)** are integrated together for DL acceleration as well as traditional computer vision tasks. The AM68A processor has two 512-bit C7x DSP running at 1 GHz, one of which is tightly coupled with an MMA capable of 4K (64 × 64) 8-bit fixed-point multiply accumulates per cycle. When run at 1 GHz, the AM68A provides 8 dense Trillion Operations per Second (TOPS).
- **H.264, H.265 encoder and decoder** can encode and decode multiple channels simultaneously. This encoder and decoder supports H.264 Baseline, Main, High Profile at L5.2 and H.265 Main Profile at L5.1. The H.264, H.265 encoder and decoder can process 480 MP/s, for example, 8 channels of 2MP at 30 fps.
- **2x 4-lane MIPI CIS-2 RX** are included in the AM68A. Two high-resolution (for example, 12MP) cameras can be directly connected to CSI-2 RX ports, and captured and preprocessed by VPAC. Capturing eight 2MP cameras is possible via MIPI CSI-2 4-to-1 aggregators.
- **BXS-4-64 GPU** offers up to 50 Giga Floating-point Operations per Second (GFLOPS) to enable dynamic 2D and 3D rendering for enhanced viewing applications.

- **Display Sub-System (DSS)** supports multiple displays with the flexibility to interface with different panel types such as eDP, DSP, and DPI.
- **Improved memory architecture and high-speed interfaces** improve the system throughput by enabling high utilization of cores and HWAs. The AM68A supports up to 34 Giga Bytes Per Second (GBps) DDR memory bandwidth.




Figure 2-1. AM68A Block Diagram With Subsystems

Deep learning inference efficiency is crucial for the performance of an edge AI system. As the [Performance and efficiency benchmarking with TDA4 Edge AI processors](#) application note shows, MMA-based deep learning inference is 60% more efficient than a GPU-based one in terms of FPS or TOPS. The optimized network models for C7xMMA are also provided by the TI Model Zoo⁽³⁾, which is a large collection of DNN models optimized for C7xMMA for various computer vision tasks. The models include popular image classification, 2D and 3D object detection, semantic segmentation, and 6D pose estimation models. [Table 2-1](#) shows the 8-bit fixed-point inference performances on AM68A for several models in the TI Model Zoo.

Table 2-1. Inference Performances of Classification, Object Detection, and Semantic Segmentation Models on AM68A

Task	Model	Image Resolution	Frame Rate (fps)	Accuracy (%)
Classification	mobileNetV2-tv	224 × 224	500	70.27 ⁽¹⁾
Object detection	ssdLite-mobDet-DSP-coco	320 × 320	218	34.64 ⁽²⁾
Object detection	yolox-nano-lite-mmdet-coco	416 × 416	268	18.96 ⁽²⁾
Semantic segmentation	deeplabv3lite-mobv2-cocoseq21	512 × 512	120	55.47 ⁽³⁾
Semantic segmentation	deeplabv3lite-regnetx800mf-cocoseq21	512 × 512	58	60.62 ⁽³⁾

(1) TOP-1 accuracy
(2) mAP 50-95
(3) mIoU

The multicore heterogeneous architecture of the AM68A provides flexibility to optimize the performance of edge AI system for various applications by utilizing suitable programmable cores or HWAs for particular tasks. For example, computationally intense deep learning (DL) inference can run on MMA with enhanced DL models, and vision processing, video encoding and decoding can be offloaded to VPAC3 and hardware-accelerated video codec for the best performance. Other functional blocks can be programmed in A72 or C7x. [Section 3](#) describes in detail how edge AI systems can be built on the AM68A for various industrial (non-automotive) use cases.

3 Edge AI Use Cases on AM68A

The popularity of edge AI technology is increasing in many existing and new use cases. The AM6xA scalable processor family is well designed for edge AI owing to a multicore heterogeneous architecture. This section introduces popular edge AI use cases which require varying input requirements, for example, resolution, frame rate, and task and computation requirements. The distribution of each task among multiple cores and HWAs in the AM68A is described to maximize the performance.

3.1 AI Box

AI Box is a cost-effective way of adding intelligence to existing non-analytics-based cameras present in retail stores, traffic roads, factories, and buildings. AI Box is preferred over AI camera because AI Box is more cost effective than replacing legacy cameras with smart cameras that have AI capabilities. Such a system receives live video streams from multiple cameras, decodes them, and does intelligent video analytics at the edge relieving the burden of transferring large video streams back to the cloud for analysis. The applications of AI Box include security surveillance system with anomaly or event detection, workplace safety systems that verifies workers wear personal protective equipment (PPE) such as goggles, safety vests, and hard hats before entering a hazardous zone. In traffic management, AI Box is used for vehicle counting, vehicle type classification, and moving direction predictions for traffic flow measurement and vehicle tracking.




Figure 3-1. AI Box Block Diagram With Data Flow on AM68A

[Figure 3-1](#) shows the data flow for AI Box on AM68A, where six channels of 2MP bitstreams are coming through Ethernet at 30 fps. The HW accelerated H.264 or H.265 decoder decodes the bitstreams and the decoded frames are scaled to smaller resolution by MSC. DL networks are applied to these smaller-resolution frames at a lower frame rate, for example, 12 fps. DL networks are accelerated by MMA. In DL preprocessing, the smaller resolution frames in YUV are converted to RGB, which is the input format to the DL network. In DL post-processing, the outputs (for example, detections) are overlaid on the input frame. Next, the output frames from six channels are stitched together into a single 2MP frame and seven channels, that is, six channels plus one composite channel are encoded by hardware accelerated H.264 or H.265 encoder at lower frame rates and streamed out or saved in storage. [Table 3-1](#) summarizes the resource utilization and estimated power consumption with six and four channels of 2MP bitstreams. An assumption made here is that each channel needs 1 TOPS for inference. The second C7x core is still available for additional vision processing and JPEG image encoding to create snapshots. While both DL pre- and post-processing run on A72 cores in this example,

both processes can run on the second C7x. In such cases, power estimates can be a little higher. The AM68A can enable the AI Box with eight channels of 2 MP bitstreams. However, due to the maximum throughput of video codec, the input frame rate and output frame rate need to be reduced to 24 fps and 4 fps, respectively.

Table 3-1. AM68A Resource Utilization and Power Consumption for the AI Box Use Case

Main IP	Utilization (6 × 2MP at 30 fps)	Utilization (4 × 2MP at 30 fps)
Decoder	6 × 2MP at 30 fps = 360 MP/s (75%)	4 × 2MP at 30fps = 240 MP/s (50%)
Encoder	6 × 2MP at 6 fps + 1 composite × 2MP at 6 fps = 84 MP/s (18%)	4 × 2MP at 6fps + 1 composite × 2MP at 6fps = 60 MP/s (18%)
Decoder + Encoder	360 MP/s + 84 MP/s = 444 MP/s (93%)	240 MP/s + 50 MP/s = 300 MP/s (62.5%)
GPU	20%	20%
VPAC (MSC)	6 × 2MP at 30 fps = 360 MP/s (60%)	4 × 2MP at 30 fps = 240 MP/s (40%)
MMA	6 × 1 TOPS per ch = 6 TOPS (75%)	4 × 1 TOPS per ch = 4 TOPS (50 %)
2 × A72	DL pre- and post-processing, depacketization, and so forth (50%)	DL pre- and post-processing, depacketization, and so forth (40%)
DDR Bandwidth	5.19 GBps (15%)	3.54 GBps (10%)
Power Consumption (85°C)	6.9 W	6.3 W

3.2 Machine Vision

Industrial 4.0 targets the increased automation for production processes within the manufacturing industry, including smart factories, smart manufacture, and so on. Industrial 5.0 emphasizes the human-centric collaboration between human and robots with artificial intelligence, that is, collaborative robot (cobot), to optimize the manufacturing process with improved automation. Machine vision is one of key technologies in Industrial 4.0 and 5.0 and the real-time processing of visual data at the edge is crucial for machine vision. The main use case of machine vision is visual quality inspection, where 2D or 3D vision-based DL is used for various purposes, for example, verifying the presence or absence of parts or ingredients in packaging systems, detecting defects, or identifying the characters on printed circuit board (PCB), gauging the dimension of parts, verifying proper assembly of parts, and the wrapping of labels around containers, detecting tool wear defects as preventive maintenance, and UAV- or drone-based fault detection systems of solar panels, turbines and pipeline, and so forth. The robot arm for pick and place of parts and assembly is another use case of machine vision for the improved collaboration between human and cobots.




Figure 3-2. Machine Vision Block Diagram With Data Flow on AM68A

Figure 3-2 illustrates the data flow for a machine vision use case example on the AM68A, which involves capturing an image sequence at 30 fps using an 8MP camera through a MIPI CSI-2 RX port. The captured raw Bayer image is processed and demosaiced to YUV by VPAC3 VISS, and VPAC3 LDC corrects any lens distortion that can be present. In this machine vision use case, the DL networks are applied to regions of interest (ROI), which are extracted on A72 cores. The number of ROIs and their sizes vary based on the specific use case. The frame rate at which DL networks are applied is also dependent on the use case. The output obtained through DL preprocessing, DL network on MMA, and DL post-processing is displayed via DSS. In the event of any unexpected detection, an alarm can be activated for human attention. The resource utilization and estimated power consumption of AM68A are shown in Table 3-2 for this machine vision use case with a single 8MP input. MMA is assumed to be fully utilized even though the actual MMA utilization can depend on the application. There is still enough room for CSI-2, VPAC, A72, and DDR bandwidth to process higher resolutions of input, for example, 1 × 16MP at 30 fps or another input of 8MP, for example, 2 × 8MP at 30 fps. Therefore, the AM68A can enable the machine vision use case for these camera configurations as long as the MMA can handle the necessary DL inferencing, but at the cost of increased power.

Table 3-2. AM68A Resource Utilization and Power Consumption for the Machine Vision Use Case

Main IP	Utilization (1 × 8MP at 30 fps)
1 × CSI-2 RX	1 × 8MP at 30 fps = 3.84 Gbps (38%)
VPAC (VISS, LDC)	1 × 8MP at 30 fps = 240 MP/s (40%)
MMA	8 TOPS (100%)
2 × A72	ROI extraction, DL pre- and post-processing, and so forth (50%)
DDR Bandwidth	5.13 GBps (15%)
Power Consumption (85°C)	6.6 W

3.3 Multi-Camera AI

Lots of existing and emerging markets for edge AI technology and edge AI processors fall into the multi-camera AI category. Multi-camera AI is similar to AI Box, but multiple cameras are directly connected to the system through MIPI CSI-2 in multi-camera AI while encoded bitstreams from remote cameras are streamed into the system through Ethernet in AI Box. Multi-camera AI use cases include the following:

- **Smart shopping cart** is an emerging end equipment for enhanced shopping experiences. Edge analytics on 2 to 6+ multiple cameras mounted on the shopping cart automatically detects the items placed in the cart, read barcodes on them, calculate order totals, and allow consumers to pay for groceries, bypassing long checkout lines. With the localization of a shopping cart, personalized shopping experience can be provided by identifying the locations of items on the shopping list and recommending new products to customers.
- **Functional safe 3D perception** provides advance assistance to machine operators for preventing collision and protecting the workers and pedestrians around. Multiple cameras are mounted around vehicles and mobile machines such as the ones used for construction, agriculture and mining and AI based 3D perception with functional safety is enabled. When combined with localization and navigation, the technology is enabling mobile machines and cobots to operate fully autonomously yet co-exist alongside humans and properties increasing the task efficiency.
- In **smart farming**, the multi-camera AI systems enable 24 hours per day, 7 days per week cattle identification, behavior monitoring and video data analysis to improve farming operations, discover health and feeding patterns, and assess how farming practices impact livestock. The camera system delivers daily event notifications to farmers through their phones, while also providing remote access to detailed analytics about their herd and farm operations, which is helpful for farmers to turn visual information into actionable insights and make data-driven decisions to maximize productivity and profitability.
- In **smart agriculture**, camera mounted on tractor, robots, and drones maximizes the efficiency in planting, fertilizing, and harvesting of the crops with the help of edge analytics.




Figure 3-3. Multi-Camera AI Block Diagram With Data Flow on AM68A

The data flow for a multi-camera AI use case example on AM68A is depicted in Figure 3-3. In this example, six 2MP cameras are combined using two MIPI CSI-2 aggregators and captured at 30 fps through two MIPI CSI-2 RX ports, which is the main difference from AI Box. The images through de-mosaic and lens distortion correction is scaled to smaller resolution by MSC. The outputs obtained through DL preprocessing, DL network on MMA and DL post-processing are encoded by the HW accelerated H.264, H.265 encoder, and streamed out or saved in storage. Table 3-3 shows the resource utilization and estimated power consumption for the multi-camera AI use case with six and four channels of 2MP camera.

Table 3-3. AM68A Resource Utilization and Power Consumption for the Multi-Camera AI Use Case

Main IP	Utilization (6 × 2MP at 30 fps)	Utilization (4 × 2MP at 30 fps)
2 × CSI-2 RX	6 × 2MP at 30fps = 5.76 Gbps (29%)	4 × 2MP at 30fps = 3.84 Gbps (19%)
VPAC (VISS, MSC, LDC)	6 × 2MP at 30fps = 360 MP/s (60%)	4 × 2MP at 30fps = 240 MP/s (40%)
MMA	8 TOPS (100%)	6 TOPS (75%)
Encoder	6 × 2MP at 30 fps = 360 MP/s (75%)	4 × 2MP at 30fps = 240 MP/s (50%)
2 × A72	DL pre- and post-processing, and so forth (50%)	DL pre- and post-processing, and so forth (40%)
DDR Bandwidth	9.7 GBps (29%)	7.2 GBps (21%)
Power Consumption (85°C)	7.2 W	6.5 W

4 Software Tools and Support

While being such a powerful processor, the AI application programming on AM68A is made simpler and faster with the adoption of Yocto Linux, mainline Kernel, production ready hardware accelerated components invoked via popular industry standard frameworks such as GStreamer, OpenVX, ROS, TensorFlow Lite (TFLite), ONNX and TVM run times. This makes edge AI application development as easy as programming in Python or C++ while still taking advantage of hardware accelerators for vision processing and AI functions.

4.1 Edge AI Software Development Kit (SDK)

The Edge AI SDK⁽⁴⁾ leverages and enables an interplay of multiple open-source components such as GStreamer, OpenVX, OpenCV, and deep learning runtime such as TFLite, ONNX, and TVM. The reference applications showcase perception-based examples such as image classification, object detection, and semantic segmentation in both Python and C++ variants. The Edge AI SDK is mainly comprised of three parts as shown in Figure 4-1. The edge AI application stack is used to run analytics applications with real-time inputs and outputs. The Foundational Linux components provide u-boot, kernel, file system, Linux drivers and more. The firmware package is used for the remote core and hardware accelerator drivers.




Figure 4-1. Processor SDK Linux for AM68A

4.2 Edge AI SDK Demonstrations

The Edge AI SDK provides the building blocks for customers to start development on any use cases discussed in this paper for AM68A. TI supports three ways of running a typical edge AI use case on the AM68A Starter Kit. Documentation for these three methods is found in the [Edge AI sample apps](#) section in the SDK documentation.

- Trying the out-of-box edge AI gallery application
- Developing edge AI applications using Python and C++ reference examples
- Running 'No Code' optimized end-to-end GStreamer applications – OpTIFlow

The following edge AI SDK packages are building blocks for DL demonstrations for AM68A:

- **Image Classification:** Detects top 5 classes from the ImageNet data set for a given input frame
- **Object Detection:** Detects and draws bounding boxes around the objects, also classifies the objects to one of the classes in COCO data set
- **Semantic Segmentation:** Classifies each pixel into class in ade20k data set

Along with the DL demonstrations other supported data flows are available in the [Edge AI dataflows](#) in the SDK documentation.

4.3 Edge AI Model Zoo

To run deep neural networks on embedded hardware, the networks need to be optimized and converted into embedded-friendly formats. TI has converted or exported 100+ models from their original training frameworks in PyTorch, TensorFlow, and MXNet into these embedded friendly formats and is hosted in a public GitHub repository⁽³⁾. In this process TI also makes sure that these models provide optimized inference speed on TI's embedded processors. These models provide a good starting point for our customers to explore high performance deep learning on TI's embedded processors.

4.4 Edge AI Studio

[Edge AI Studio](#) ⁽⁵⁾ is a collection of tools to accelerate the development of edge AI applications on TI's embedded processors including AM68A. The Edge AI studio lets you build, evaluate, and deploy deep learning models. The tool is free to use, and is available now. The Edge AI Studio has the following key features:

- No installation is required because the tool resides in the cloud
- Free access to hundreds of optimized, pretrained models
- Remote access to real evaluation hardware
- Bring your own data (BYOD), bring your own model (BYOM)

There are two tools as a part of the Edge AI Studio: model analyzer and model composer. Model analyzer, which is launching in 2Q 2023, allows connection remotely to real evaluation hardware to deploy and test AI model performance on a TI's embedded processor. Model analyzer helps a lot during the model evaluation phase. With model analyzer, the performance of AI model can be evaluated by:

- Selecting from hundreds of optimized, pretrained models or using a custom model (BYOM).
- Deploying models easily using industry standard APIs.
- Getting benchmarks for latency, frames-per-second processing, DDR bandwidth, and accuracy.

Model composer is a fully integrated solution for creating Edge AI applications. Model composer helps collect, annotate data, train, optimize and compile AI models for TI's embedded processors. This tool supports BYOD, which enables the retraining of the models from TI Model Zoo to fine tune the performance for your unique application requirements with custom data.

5 Conclusion

The AM68A is the best design for the development of various Edge AI applications with the moderate number of cameras in terms of the performance and the ease of development. The heterogeneous architecture with multiple cores and hardware accelerators provides flexible ways to optimize the performance. The Edge AI SDK with Edge AI Studio and Model Zoo streamlines the development process from the training, optimizing and deploying of the model, to the application deployment with GStreamer, OpenVX, and OpenCV.

The AM68A is ready for the developer to start developing Edge AI applications with the [AM68A starter kit EVM](#) and the [Processor SDK Linux for AM68A](#). Documentation, prebuilt images, and demonstrations are available from the [AM68A product page](#). See TI's [Edge AI](#) folder for more information about Edge AI across TI's processors.

6 Reference

1. [Edge AI processor market 2023 industry growth estimate, Key Players Analysis, Gross Margins, Strategy, Application 2031](#), MarketWatch
2. Texas Instruments, [Edge AI smart cameras using energy efficient AM62A processors Technical White Paper](#)
3. Edge AI Model Zoo, <https://github.com/TexasInstruments/edgeai-modelzoo>
4. Processor SDK Linux for AM68A, https://software-dl.ti.com/jacinto7/esd/processor-sdk-linux-edgeai/AM68A/08_06_00/exports/docs/common/sdk_overview.html
5. Texas Instruments, [Edge AI Studio](#)

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATA SHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES "AS IS" AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, regulatory or other requirements.

These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI's products are provided subject to [TI's Terms of Sale](#) or other applicable terms available either on [ti.com](https://www.ti.com) or provided in conjunction with such TI products. TI's provision of these resources does not expand or otherwise alter TI's applicable warranties or warranty disclaimers for TI products.

TI objects to and rejects any additional or different terms you may have proposed.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2023, Texas Instruments Incorporated