Coral M.2 Accelerator with Dual Edge TPU

SKU
102110449

The Coral M.2 Accelerator with Dual Edge TPU is an M.2 module that brings two Edge TPU coprocessors to existing systems and products with an available M.2 E-key slot.

PRODUCT DETAILS

Features

Performs high-speed ML inferencing: Each Edge TPU coprocessor is capable of performing 4 trillion operations per second (4 TOPS), using 2 watts of power. For example, it can execute state-of-the-art mobile vision models such as MobileNet v2 at almost 400 FPS, in a power-efficient manner. With the two Edge TPUs in this module, you can double the inferences per second (8 TOPS) in several ways, such as by running two models in parallel or pipelining one model across both Edge TPUs.

Works with Debian Linux and Windows: Integrates with Debian-based Linux or Windows 10 systems with a compatible card module slot.

Supports TensorFlow Lite: No need to build models from the ground up. TensorFlow Lite models can be compiled to run on the Edge TPU.

Supports AutoML Vision Edge: Easily build and deploy fast, high-accuracy custom image classification models to your device with AutoML Vision Edge.

Description

The Coral M.2 Accelerator with Dual Edge TPU is an M.2 module that brings two Edge TPU coprocessors to existing systems and products with an available M.2 E-key slot.

Specifications

ML accelerator	**2x Google Edge TPU coprocessor: 8 TOPS (int8); 2 TOPS per watt**
Hardware Interface	M.2 E-key (M.2-2230-D3-E)
Serial Interface	Two PCIe Gen2 x1
Operating Voltage	3.3V +/- 10%
Storage Temperature	-40 to +85°C
Operating Temperature	-40 to +85°C
Relative Humidity	0 to 90% (non-condensing)

Dimensions

Part List

1 x Coral M2 Accelerator with Dual Edge TPU

ECCN/HTS

HSCODE   8471800000

UPC