

TensorFlow Guide

TensorFlow Overview

TensorFlow™ is an artificial intelligence framework that can be used for executing machine learning algorithms. While a computation expressed using TensorFlow can be executed in parallel across heterogeneous systems such as GPUs, support has so far been limited to NVIDIA® processors using CUDA®.

SYCL™ is enabling developers to access a wider range of processors by bringing support for OpenCL™ devices to the TensorFlow framework. OpenCL is a framework for writing programs that execute across heterogeneous platforms, and SYCL is a royalty-free, cross-platform C++ abstraction layer that builds on the underlying concepts, portability and efficiency of OpenCL, while adding the ease-of-use and flexibility of modern C++14.

TensorFlow can be used to build deep neural networks capable of machine learning, and these networks rely heavily on linear algebra where matrix calculations are key to building up predictions based on the input data set. TensorFlow, which consists of tensors (n-dimensional matrices), uses the **Eigen library** that has been built specifically for performing linear algebra, and written in C++. The Eigen library does a lot of heavy lifting by creating and fusing kernels, and it is these kernels that make it possible to run many calculations in parallel.

ComputeCpp converts Eigen's kernels into SYCL kernels so that they can be run on a range of OpenCL devices. In addition to enabling parallelization of the Eigen library on OpenCL devices, the most common TensorFlow operations are being supported such as basic math functions.

Parallelization is important from a processing and power management perspective. Since tensors are n-dimensional vectors, having access to parallelization of your TensorFlow code is important not just at the training stage, but also when performing inference on new data sets.

Find out how to set up TensorFlow with SYCL by following the **TensorFlow Getting Started Guide**.