# UG10166 i.MX Machine Learning User's Guide Rev. LF6.12.3\_1.0.0 — 31 March 2025

User guide

#### **Document information**

Information	Content
Keywords	i.MX, Linux, LF6.12.3_1.0.0
Abstract	The NXP eIQ Machine Learning Software Development Environment (hereinafter referred to as "NXP eIQ") provides a set of libraries and development tools for machine learning applications targeting NXP microcontrollers and applications processors.



# **1** Software Stack Introduction

The NXP eIQ Machine Learning Software Development Environment (hereinafter referred to as "NXP eIQ") provides a set of libraries and development tools for machine learning applications targeting NXP microcontrollers and application processors. The NXP eIQ is contained in the *meta-imx/meta-ml* Yocto layer. See also the *i.MX Yocto Project User's Guide* (UG10164) for more information.

The following four inference engines are currently supported in the NXP eIQ software stack: TensorFlow Lite, ONNX Runtime, PyTorch, and OpenCV. The following figure shows the supported eIQ inference engines across the computing units.



Figure 1. NXP elQ supported compute vs. inference engines

The NXP eIQ inference engines support multi-threaded execution on Cortex-A cores. Additionally, TensorFlow Lite also supports acceleration on the GPU or NPU. Generally, the NXP eIQ is prepared to support the following key application domains:

- Vision
  - Multi-camera observation
  - Active object recognition
  - Gesture control
- Voice
  - Voice processing
  - Home entertainment
- Sound
  - Smart sense and control
  - Visual inspection
  - Sound monitoring

# 2 TensorFlow Lite

TensorFlow Lite, recently renamed to LiteRT, is an open-source software library focused on running machine learning models on mobile and embedded devices (available at <a href="http://www.tensorflow.org/lite">http://www.tensorflow.org/lite</a>). It enables

on-device machine learning inference with low latency and small binary size. TensorFlow Lite also supports hardware acceleration:

- Using the VX Delegate on i.MX 8 series.
- Using the Ethos-U Delegate on i.MX 93.
- Using the Neutron Delegate and GPU Delegate on i.MX 95.

The TensorFlow Lite source code for this Yocto Linux release is available at this <u>repository</u>, branch If-6.12.3\_1.0.0. This repository is a fork of the mainline <u>https://github.com/tensorflow/tensorflow</u>, and it is optimized for NXP i.MX 8 and i.MX 9 platforms.

## Features:

- TensorFlow Lite v2.18.0
- Multithreaded computation with acceleration using Arm Neon SIMD instructions on Cortex-A cores
- Parallel computation using GPU/NPU hardware acceleration (on shader or convolution units)
- C++ and Python API (supported Python version 3)
- Per-tensor and Per-channel quantized models support

# 2.1 TensorFlow Lite software stack

The TensorFlow Lite software stack is shown in the following picture. The TensorFlow Lite supports computation on the following hardware units:

- CPU Arm Cortex-A cores
- GPU/NPU hardware accelerator using the VX Delegate on i.MX 8 Series. See Section 6.1 for details.
- NPU hardware acceleration using Ethos-U Delegate on i.MX 93 NPU. See Section 6.2 for details.
- NPU hardware acceleration using Neutron Delegate on i.MX 95 NPU. See Section 6.4 for details.
- GPU hardware acceleration using the GPU Delegate on i.MX 95 GPU

See <u>Section 1</u> for some details about supporting of computation on GPU/NPU hardware accelerator on different hardware platforms.

# **NXP Semiconductors**

# UG10166

## i.MX Machine Learning User's Guide



Figure 2. TensorFlow Lite software stack

## Note:

The first execution of the model inference using the delegate takes longer, because of the time required for computational graph compilation and initialization for the hardware accelerator. The following iterations perform much faster. The computational graph is the representation of the operations and their dependencies to perform computation specified by the model. The computation graph is built during the model parsing phase. See <u>Section 6</u> for details.

The VX Delegate implementations use the OpenVX library for computational graph execution on the GPU/NPU hardware accelerator. Therefore, OpenVX library support must be available for the selected device to be able to use the acceleration. For more details on the OpenVX library availability, see the i.MX Graphics User's Guide (UG10159).

Refer to the i.MX Graphics User's Guide (UG10159) for list GPUs with OpenVX support. Note that the GC7000 Lite and GC7000 Ultra Lite GPUs does not support full OpenVX however still capable to run ML workload.

# 2.2 Inference backends and delegates

Inference backend is the compute engine that enables efficient execution of machine learning models on edge devices. Tensorflow Lite comes with the options to enable different backends through the delegate mechanism.

## 2.2.1 Built-in kernels

Default inference backend is the CPU with reference kernels from TensorFlow Lite implementation. Built-in kernels provide full support for TensorFlow Lite operator set.

The built-in kernels are built with RUY matrix multiplication library enabled, which increases the performance of the kernels for floating point and quantized operations.

## 2.2.2 XNNPACK Delegate

XNNPACK library is a highly optimized library of floating-point quantize neural network inference operators for ARM, WebAssembly, and x86 platforms. The XNNPACK library is available through XNNPACK delegate in TensorFlow Lite. The XNNPACK delegate computation is performed on the CPU.

It provides optimized implementation for a subset of TensorFlow Lite operator set. In general, it provides better performance than the built-in kernels.

**Note:** The models are executed via the XNNPACK Delegate by default. XNNPack now also supports quantized operators (since 2021).

## 2.2.3 VX Delegate

VX Delegate enables accelerating the inference on on-chip hardware accelerator on i.MX 8 series. The VX Delegate directly uses the hardware accelerator driver (OpenVX with extension) to fully utilize the accelerator capabilities.

The VX Delegate is available as *external delegate*<sup>1</sup>. The corresponding library is available in /usr/lib/libvx\_delegate.so.

VX Delegate is supported in both C++ and Python API. For using VX Delegate (or any external delegate), see the <u>external\_delegate\_provider</u> implementation in C++ and/or <u>label\_image.py</u> for Python. List of supported operators are available in <u>op\_status.md</u>.

## 2.2.4 Ethos-U Delegate

Ethos-U Delegate is an external delegate on i.MX 93 Linux platforms. It enables accelerating the inference on the on-chip hardware accelerator. The Ethos-U Delegate directly uses the hardware accelerator driver (Ethos-U driver stack) to fully utilize the accelerator capabilities.

The Ethos-U Delegate is available as *external delegate*. The corresponding library is available in /usr/lib/ libethosu delegate.so.

Ethos-U Delegate is supported in both C++ and Python API. For using Ethos-U Delegate (or any external delegate), see the <code>external\_delegate\_provider</code> implementation in C++ and/or <code>label\_image.py</code> for Python. List of supported operators are available in <u>SUPPORTED\_OPS.md</u>.

## 2.2.5 Neutron Delegate

Neutron Delegate is an external delegate on i.MX 9 series Linux platform containing Neutron-S NPU. It captures the operators and aggregates them as a neutron graph node, which can be directly offloaded and accelerated by the Neutron-S NPU.

The delegate library is available in /usr/lib/libneutron\_delegate.so. It can be used in both C+ + and Python API environments. For using Neutron Delegate, see the external\_delegate\_provider implementation in C++ and/or label\_image.py for Python usage.

<sup>1</sup> An external delegate is a special Tensorflow Lite delegate that is simply initialized from loading a dynamic library which encapsulates an actual <u>TensorFlow Lite delegate implementation</u>

## Note:

For the offline compilation, the model should be converted through the eIQ toolkit first. In the converted model, the neutronGraph node is already generated. The neutron-delegate only captures the neutronGraph node and offloads the work to Neutron-S.

## 2.2.6 GPU Delegate

GPU Delegate is an internal TensorFlow Lite delegate enabled on the i.MX 95 platform to leverage inferences through the Arm Mali G310 GPU.

The GPU Delegate supports OpenCL as the main backend, and applies a set of optimizations such as accuracy lowering for performance (FP32 to FP16 lowering) and allows quantize model execution. Further information about the GPU delegate can be found in <a href="https://www.tensorflow.org/lite/performance/gpu">https://www.tensorflow.org/lite/performance/gpu</a>.

GPU delegate is supported in C++ API. Refer to gpu\_delegate\_provider.cc in <a href="https://github.com/nxp-imx/tensorflow-imx/blob/lf-6.6.23\_2.0.0/tensorflow/lite/tools/delegates/gpu\_delegate\_provider.cc">https://github.com/nxp-imx/tensorflow-imx/blob/lf-6.6.23\_2.0.0/tensorflow/lite/tools/delegates/gpu\_delegate\_provider.cc</a>.

#### Note:

- Even if the GPU Delegate supports quantized models, performance might be degraded when compared to other delegates such as XNNPACK using 6 cores.
- As the GPU delegate dynamically loads OpenCL, libOpenCL.so is expected to be present in the system. If libOpenCL.so is missing but libOpenCL.so.1 exists, issues can be solved by adding a symbolic link:

ln -s /usr/lib/libOpenCL.so.1 /usr/lib/libOpenCL.so

# 2.3 Delivery package

The TensorFlow Lite is available using Yocto Project recipes.

The TensorFlow Lite delivery package contains:

- TensorFlow Lite shared libraries
- TensorFlow Lite header files
- Python Module for TensorFlow Lite
- Image classification example application for C++ (label image) and for Python (label image.py)
- TensorFlow Lite benchmark application (benchmark model)
- TensorFlow Lite evaluation tools (coco\_object\_detection\_run\_eval, imagenet\_image\_classification\_run\_eval, inference\_diff\_run\_eval), see TensorFlow Lite Delegates for details.

For application development, the TensorFlow Lite shared libraries and header files are available in the SDK. See <u>Section 2.5</u> for more details.

There are following delegates available in the TensorFlow Lite delivery package:

- XNNPACK Delegate
- VX Delegate
- Ethos-U Delegate
- Neutron Delegate
- GPU Delegate

# 2.4 Build details

TensorFlow Lite uses CMake build system for compilation. Notable remarks to package build are:

UG10166

## i.MX Machine Learning User's Guide

- RUY matrix multiplication library is enabled (TFLITE\_ENABLE\_RUY=On). RUY matrix multiplication library offers better performance compared to kernels build with Eigen and GEMLOWP.
- XNNPACK Delegate support (TFLITE ENABLE XNNPACK=On)
- External Delegate support (TFLITE ENABLE EXTERNAL DELEGATE=On)
- (i.MX 95) GPU Delegate support (TFLITE ENABLE GPU=On)
- The runtime library is built and provided as a shared library (TFLITE\_BUILD\_SHARED\_LIB=On). If static linking of the TensorFlow Lite library to the application is preferred, keep this switch in off state (default settings). This might be convenient if the application is built with CMake as described in the Section <u>Section 2.5.1</u>.
- The package is compiled with the default -O2 optimization level. Some CPU kernels, such as RESIZE\_BILINEAR, are known to perform better with -O3 optimization level. However, some performs better with -O2, such as ARG\_MAX. We recommend to adjust the optimization level, based on the application needs.

Yocto project builds the TensorFlow Lite with these settings. The build configuration can be changed by either updating the TensorFlow Lite Yocto recipe in the meta-imx layer (located in meta-imx/meta-ml/recipes-libraries/tensorflow-lite/), or building the TensorFlow Lite from source code using the CMake and the Yocto SDK.

# 2.5 Application development

This section describes how to use TensorFlow Lite C++ API in the application development.

To start with TensorFlow Lite C++ application development, a Yocto SDK must be generated firstly. See the *i.MX* Yocto Project User's Guide (UG10164) for detailed information on how to generate Yocto SDK environment for cross-compiling.

To build an application that uses the TensorFlow Lite, use the following options:

- Create a CMake project that uses TensorFlow Lite (CMake superbuild pattern).
- Use the Yocto SDK precompiled libraries.

The CMake configuration file of TensorFlow Lite is under tensorflow/lite/CMakeLists.txt from the root repository.

## 2.5.1 Create CMake project which uses TensorFlow Lite

The recommended way is to create a CMake project, which uses TensorFlow Lite as described in <u>Build</u> <u>TensorFlow Lite with CMake</u>. CMake takes care of dependencies preparation, including download, configure and build steps.

To demonstrate this build option, there is a minimal example project available in tensorflow/lite/ examples/minimal. To build it:

## 1. Build the Native flat compiler for TensorFlow Lite.

```
cmake -S tensorflow/lite/tools/cmake/native_tools/flatbuffers \
    -B native-tools \
    -DCMAKE_INSTALL_PREFIX=native-tools
cmake --build native-tools -- -j 4 --keep-going
cmake --install native-tools
```

2. Set up the Yocto SDK as described above. To activate this Yocto SDK environment on your host, use the following command:

```
$ source <Yocto_SDK_install_folder>/environment-setup-aarch64-poky-linux
```

3. Configure the project using CMake:

```
$ mkdir build-minimal-example; cd build-minimal-example
$ cmake -DCMAKE_TOOLCHAIN_FILE=${OE_CMAKE_TOOLCHAIN_FILE} -
DTFLITE_ENABLE_XNNPACK=on \
-DTFLITE_ENABLE_RUY=on \
-DTFLITE_HOST_TOOLS_DIR=native-tools
../tensorflow/lite/examples/minimal
```

4. Build the project:

\$ cmake --build . -j4

5. The minimal example is available in the build directory:

```
$ file minimal
minimal: ELF 64-bit LSB shared object, ARM aarch64, version 1 (GNU/
Linux), dynamically linked, interpreter /lib/ld-linux-aarch64.so.1,
BuildID[sha1]=4a928894439e0b33217ea28790378690ab4ce7cd, for GNU/Linux
3.14.0, with debug info, not stripped
```

6. Optionally you can strip the final binary:

```
$ $STRIP --remove-section=.comment --remove-section=.note --strip-unneeded
<file>
```

This build option has several advantages:

- Automatic dependency resolution based on configure options
- Option to choose between static or dynamic linking (TFLITE BUILD SHARED LIB=on/off)
- Building the whole project (including its dependencies) in the Debug mode (CMAKE\_BUILD\_TYPE=Debug/ Release/...), for enhanced debugging experience

## 2.5.2 Using Yocto SDK precompiled libraries

Another option is to use the precompiled binaries and header files which are directly available in the Yocto SDK. The TensorFlow Lite artifacts are in the Yocto SDK as follows:

- TensorFlow Lite shared library (libtensorflow-lite.so) in /usr/lib
- TensorFlow Lite header files in /usr/include

**Note:** Not all TensorFlow Lite dependencies are installed in the Yocto SDK and it is necessary to download and optionally build them manually. For the required versions see the <code>tensorflow/lite/tools/cmake/modules/folder</code>.

To build the image classification demo (label\_image), located in tensorflow/lite/examples/ label\_image/, follow these steps:

1. Create build directory:

```
$ mkdir build-manual
$ cd build-manual
```

2. Download the Abseil library dependency:

```
$ wget https://github.com/abseil/abseil-cpp/
archive/997aaf3a28308eba1b9156aa35ab7bca9688e9f6.tar.gz -O abseil-cpp.tar.gz
$ tar -xzf abseil-cpp.tar.gz
$ mv abseil-cpp-997aaf3a28308eba1b9156aa35ab7bca9688e9f6 abseil-cpp
```

#### 3. Build the label\_image example:

```
$ $CC ../tensorflow/lite/examples/label_image/label_image.cc ../tensorflow/
lite/examples/label_image/bitmap_helpers.cc ../tensorflow/lite/tools/
evaluation/utils.cc ../tensorflow/lite/tools/delegates/delegate_provider.cc -
Iabseil-cpp -02 -ltensorflow-lite -lstdc++ -lpthread -lm -ldl -lrt -I../
```

## 2.6 Enabling TensorFlow Operators in TensorFlow Lite Runtime

The TensorFlow Lite Operator Set counts more than a hundred of frequently used operators and layers, and majority of Machine Learning models can fit into it. Still the TensorFlow Lite Operator Set is only a subset of TensorFlow Operator Set, so not every model is convertible.

To tackle this limitation, TensorFlow offers an option to use TensorFlow operators inside the TensorFlow Lite runtime. See <a href="https://www.tensorflow.org/lite/guide/ops\_select">https://www.tensorflow.org/lite/guide/ops\_select</a>. It shows how this feature can be used with NXP i.MX devices with Yocto Linux platform.

## 2.6.1 TensorFlow and TensorFlow Lite Operator Set

If the model is not convertible within the standard TensorFlow Lite Operator Set, the TensorFlow Lite converter raises an error, indicating particular operator is not available in TensorFlow Lite, for example:

```
Some ops are not supported by the native TFLite runtime, you can enable TF
kernels fallback using TF Select. See instructions: https://www.tensorflow.org/
lite/guide/ops_select
TF Select ops: Roll
Details: tf.Roll(tensor<?x10xf32>, tensor<i32>, tensor<i32>) -> (tensor<?
x10xf32>) : {device = ""}
```

To Convert such a model, the Select TensorFlow Operators feature needs to be enabled in the Converter:

converter.target\_spec.supported\_ops = [
 tf.lite.OpsSet.SELECT\_TF\_OPS # enable TensorFlow ops.
]

When the model is converted with SELECT\_TF\_OPS enabled, TensorFlow operators are transformed into Flex operators, which are supported by TensorFlow Lite through the Flex Delegate. The Flex Delegate is the TensorFlow Lite counterpart for the **Select TensorFlow Operators** feature and bridges the TensorFlow Lite and TensorFlow runtimes.

## 2.6.2 Building the TensorFlow Lite Library with the Flex Delegate for i.MX Linux platforms

The library can be built using bazel on any supported host or inside the Docker container. It is recommend to use Docker because the environment is ready for TensorFlow compilation. Compilation outside of Docker might fail for multiple reasons. This document focuses on building the TensorFlow Lite Library with Flex Delegate inside the TensorFlow's Docker image.

## Note:

To build the library outside the Docker image, the <code>bazel</code> build system needs to be installed on the machine. The TensorFlow requires an exact version of <code>bazel</code>, which is specific to particular TensorFlow version. Therefore, use <code>bazelisk</code> to handle the <code>bazel</code> version management. Find the <code>bazelisk</code> tool on its GitHub space <u>https://github.com/bazelbuild/bazelisk</u>, with prebuilt executables for multiple platforms available.

It is recommended to have at least 32 GB RAM to build the Flex Delegate, and ensure that enough inodes are available in build-related directories, such as /tmp.

## 2.6.2.1 Checking out the TensorFlow repository

To build the TensorFlow Lite library, check out the TensorFlow sources:

Clone the tensorflow-imx repository from <u>https://github.com/nxp-imx/tensorflow-imx</u> and check out the corresponding release branch:

```
$ git clone https://github.com/nxp-imx/tensorflow-imx.git -b lf-6.12.3_1.0.0
$ cd tensorflow-imx
```

```
S Cd tensorilow-imx
```

## 2.6.2.2 Setting up Docker VM

For more details about the Docker VM setup for TensorFlow, see <u>https://github.com/tensorflow/tensorflow/tree/</u>master/tensorflow/tools/tf\_sig\_build\_dockerfiles.

## Note:

Depending on the host, the Docker may require administrative privileges to run (e.g., *sudo* in Linux). Alternatively, the Docker Daemon can run as a non-root user (Rootless mode), as described here <u>https://</u> <u>docs.docker.com/engine/security/rootless/</u>.

- 1. Download the tensorflow/build:2.18-python3.12 Docker image. This image is aligned with the upstream release and contains all the required tools to build Flex Delegate related components.
  - \$ docker pull tensorflow/build:2.18-python3.12
- 2. Run the Docker VM. During the build process, Bazel downloads various packages from the Internet. Therefore, Internet access inside the instantiated container is required. In case of conflict, a minimal setup is to initialize http\_proxy and https\_proxy environmental variables when launching the Docker image. Particular steps depend on the host configuration.

```
$ docker run -e "http_proxy=<your-http-proxy>" \
    -e "https_proxy=<your-https-proxy>" \
    -e "no_proxy=localhost,127.0.0.1" \
    -it -w /tensorflow -v /<path-to-tensorflow-sources>:/tensorflow
    -e HOST_PERMS="\\((id -u):\\)(id -g)" \
    tensorflow/build:2.18-python3.12
```

## 2.6.2.3 Building the TensorFlow Lite with Flex Delegate

The main CPUs in NXP i.MX 8 and i.MX 9 families are Arm based, implementing AArch64 ISA, and the platform OS is based in the Embedded Linux BSP release. To build TensorFlow Lite with Flex Delegate, the build system supports the described environment combination through the elinux aarch64 option.

1. Configure the project using the configure script:

\$ ./configure

**Note:** More details about the configuration can be found in <code>build\_benchmark\_flex\_test.sh</code> and <code>build\_elinux\_flex\_libs\_from\_models.sh</code> located in the <code>tensorflow/lite/tools</code> directory. The Flex Delegate sources and <code>bazel</code> build recipes are located in <code>/tensorflow/lite/delegates/flex.</code> There are two libraries defined:

- tensorflowlite\_flex\_[full|reduced] TensorFlow Lite Flex Delegate shared library (libtensorflowlite\_flex.so)
- delegate\_[full|reduced] special target to be used for static linking of the TensorFlow Lite Flex Delegate. Similar to object library concept in CMake.

Additionally Bazel targets for building different variants of the benchmark model binary are provided in the /tensorflow/lite/delegates/flex/test/BUILD, for evaluation purposes:

benchmark model plus flex [dynamic] [full|reduced]

2. Build the benchmark model plus flex target with a full TensorFlow Operator Set:

```
$ bazel --output_base=/tensorflow/docker-build/ build --config=monolithic
--config=elinux_aarch64 -c opt --cxxopt='--std=c++17' --
host_crosstool_top=@bazel_tools//tools/cpp:toolchain //tensorflow/lite/
delegates/flex/test:benchmark model plus flex full
```

**Note:** To preserve the Bazel's cache, use the --output base switch to override the default output base. For the build outside the Docker, this switch can be omitted. The directory shall be available prior to running Bazel build.

The output is the benchmark model plus flex binary with statically linked Flex Delegate. This can be directly used on NXP MPU platforms.

The -c (or -- compilation mode) affects code generation option. It can be set to fastbuild, dbg or opt. See https://bazel.build/docs/user-manual#build-semantics. To build the Flex Delegate for debugging purposes, use the -c dbg option.

The following table lists the benchmark model plus flex \* build configurations.

Operator set	Static linkage	Dynamic linkage
Full Flex Delegate	benchmark_model_plus_flex_full	benchmark_model_plus_flex_dynamic_full
Reduced Flex Delegate	benchmark_model_plus_flex_reduce	edenchmark_model_plus_flex_dynamic_reduce

#### Table 1. benchmark model build configurations

## 2.6.3 Reducing the size of the Flex Delegate library

The previous section describes how to build a TensorFlow Lite Library with a complete TensorFlow Operator Set. The approach is useful for quick evaluation, but for practical use, it generates an oversized binary. Moreover, typically only a small subset of TensorFlow operators are required.

To minimize the size, there is a model-dependent build option which extracts the required operators from the models, and selectively includes them in the deployed TensorFlow Lite library. For more details, see https:// www.tensorflow.org/lite/guide/reduce binary size.

For example, there are targets with the reduced suffix, which builds the TensorFlow Lite library for the / tensorflow/lite/delegates/flex/test/simple flex model int8.tflite example model. The model contains a single TensorFlow operation: tf.roll().

The reduced size binaries are built in the following process:

- 1. Generate the headers and listings for the Flex operators contained in the provided TensorFlow Lite models. This is managed by the rule fetch flex files. Note that the command is "run" and executed natively in the host platform (lack of the elinux aarch64 option).
- 2. Build the target Flex delegate (using the rules contained in Table 1).

```
$ bazel --output base=/tensorflow/docker-build/ run //tensorflow/lite/delegates/
flex/test:fetch_flex_files
$ bazel --output_base=/tensorflow/docker-build/ build --config=monolithic
--config=elinux_aarch64 -c opt --cxxopt='--std=c++17' --
host crosstool top=@bazel tools//tools/cpp:toolchain
```

```
//tensorflow/lite/delegates/flex/test:benchmark model plus flex reduced
```

To build the TensorFlow Lite library for custom model, use a bazel function tflite\_flex\_cc\_library (for static library) or tflite\_flex\_shared\_library (for shared library), and list the models into the models attribute. Remember regenerating the headers and listings through the fetch\_flex\_files rule for each additional model containing Flex operators. The kernel\_headers attribute uses the files generated by the fetch\_flex\_files rule through the load\_flex\_kernel\_header rule. For more details, see the tensorflow/lite/delegates/flex/test/BUILD file.

```
tflite_flex_cc_library(
    name = "delegate_reduced",
    models = [
        "simple_flex_model_int8.tflite",
    ],
    visibility = ["//visibility:public"], kernel_headers = loaded_headers,
)
```

This library can be used inside the bazel to link to a custom TensorFlow Lite binary, like this:

**Note:** To ease the build process for Reduced Size Flex Delegate containing user models, refer to build elinux flex libs from models.sh **located in the** tensorflow/lite/tools directory.

## 2.6.4 Flex Delegate deployment on NXP i.MX Linux platform

For the statically linked binary (in this usecase, benchmark\_model\_plus\_flex\_[full|reduced]), copy the binary to target the device rootfs.

For the dynamically linked binary (in this usecase, benchmark\_model\_plus\_flex\_dynamic\_[full| reduced]), copy both libtensorflowlite\_flex.so and the binary to target the device rootfs. The libtensorflowlite\_flex.so should be copied to /usr/lib/, or alternatively the LD\_LIBRARY\_PATH should be set, to load the library by the dynamic linker or loader.

1. Copy the simple\_flex\_model\_int8.tflite example model on the i.MX platform, e.g., to /usr/bin/ tensorflow-lite-2.18.0/examples/.

```
$ scp tensorflow/lite/delegates/flex/test/simple_flex_model_int8.tflite
root@<imx-board>:/usr/bin/tensorflow-lite-2.18.0/examples/
```

2. Run the example application benchmark model:

```
$ ./benchmark_model_plus_flex_dynamic_full --graph=/usr/bin/tensorflow-
lite-2.18.0/examples/simple flex model int8.tflite
```

======================================						
KBl [tim	[Node type]	[count]	[avg ms]	[avg %]	[cdf %]	[mem
TfLiteXNNP	ackDelegate	3	2.821	81.250%	81.250%	
TfLiteFlexDelegate		1	0.640	18.433%	99.683%	
0.000	RESHAPE	1	0.008	0.230%	99.914%	
0.000	SOFTMAX	1	0.003	0.086%	100.000%	
0.000	1					

With --enable\_op\_profiling=true, the FlexDelegate invocation is displayed:

## 2.6.5 Using hardware accelerators

The TensorFlow Operators are not part of the TensorFlow Lite Operators Set, so the hardware accelerator on i.MX platforms does not support these operators, though the acceleration of the TensorFlow Lite operators present in the model is supported.

For the hardware Acceleration on i.MX8 Linux platforms, use the VX Delegate (external delegate). The <code>benchmark\_model\_plus\_flex</code> already includes support for external delegates, so the <code>-external\_delegate\_path</code> CLI option can be used for inference acceleration:

```
$ ./benchmark_model_plus_flex_dynamic_full --graph=/usr/bin/tensorflow-
lite-2.15.0/examples/simple_flex_model_int8.tflite --enable_op_profiling=true --
external delegate path=/usr/lib/libvx delegate.so
```

For the hardware acceleration on the i.MX 9 Linux platform, use the Ethos-U Delegate for i.MX 93 or Neutron Delegate for i.MX 95.

Alternativelly, convert the model with the Arm Vela Compiler as described in <u>Section 6.2.3</u>, and also use the Ethos-U Delegate.

\$ vela /usr/bin/tensorflow-lite-2.18.0/examples/simple flex model int8.tflite

Run benchmark model plus flex\* with the Ethos-u Delegate.

```
$ ./benchmark_model_plus_flex_dynamic_full --graph=/usr/bin/tensorflow-
lite-2.18.0/examples/simple_flex_model_int8.tflite --enable_op_profiling=true --
external_delegate_path=/usr/lib/libethosu_delegate.so
```

## 2.6.6 Flex Delegate limitations

The Flex Delegate has the following limitations:

## **CPU support only for TensorFlow Operators**

Flex Delegate operators are not supported on the hardware accelerators of i.MX platforms. The TensorFlow Operators fall back to CPU. The acceleration of supported TensorFlow Lite Ops in the model is not impacted. The model can freely combine TensorFlow Lite and TensorFlow Operators. Supported TensorFlow Lite operators of the model will be accelerated.

## 2.7 Running image classification example

A Yocto Linux BSP image with machine learning layer included by default contains a simple pre-installed example called 'label\_image' usable with image classification models. The example binary file is located at:

/usr/bin/tensorflow-lite-2.18.0/examples



Figure 3. TensorFlow image classification input Demo instructions:

To run the example with mobilenet model on the CPU, use the following command:

```
$ ./label_image -m mobilenet_v1_1.0_224_quant.tflite -i grace_hopper.bmp -l
labels.txt
```

The output of a successful classification on the i.MX 8MPlus SoC for the 'grace\_hopper.bmp' input image is as follows:

```
Loaded model mobilenet_v1_1.0_224_quant.tflite
resolved reporter
invoked
average time: 39.271 ms
0.780392: 653 military uniform
0.105882: 907 Windsor tie
0.0156863: 458 bow tie
0.0117647: 466 bulletproof vest
0.00784314: 835 suit
```

Note: For floating point layers, the TensorFlow Lite uses XNNPACK delegated by default.

## 2.7.1 Running the example on the i.MX 8 platform hardware accelerator

To run the example application on the CPU without using the XNNPACK delegate, use the -- use\_xnnpack=false switch.

To run the example with the same model on the GPU/NPU hardware accelerator, add -- external\_delegate\_path=/usr/lib/libvx\_delegate.so (for VX Delegate) command line argument.

To differentiate between the 3D GPU and the NPU, use the USE\_GPU\_INFERENCE environmental variable. For example, to run the model accelerated on the NPU hardware using VX Delegate, use this command:

```
$ USE_GPU_INFERENCE=0 ./label_image -m mobilenet_v1_1.0_224_quant.tflite
-i grace_hopper.bmp -l labels.txt --external_delegate_path=/usr/lib/
libvx delegate.so
```

The output of the NPU acceleration on the i.MX 8MPlus processor is as follows:

```
INFO: Loaded model ./mobilenet v1 1.0 224 quant.tflite
INFO: resolved reporter
Vx delegate: allowed builtin code set to 0.
Vx delegate: error during init set to 0.
Vx delegate: error during prepare set to 0.
Vx delegate: error during invoke set to 0.
EXTERNAL delegate created.
INFO: Applied EXTERNAL delegate.
W [HandleLayoutInfer:257]Op 18: default layout inference pass.
INFO: invoked
INFO: average time: 2.567 ms
INFO: 0.768627: 653 military uniform
INFO: 0.105882: 907 Windsor tie
INFO: 0.0196078: 458 bow tie
INFO: 0.0117647: 466 bulletproof vest
INFO: 0.00784314: 835 suit
```

## 2.7.2 Running the example on the i.MX 93 platform with Ethos-U

To use the hardware acceleration on i.MX 93, convert the model using the Vela compiler first, and run the model with the Ethos-U delegate. Alternatively, directly run the model with the Ethos-U delegate. The model is then converted in the delegate. For details, see Section 6.2.3.

To run the example with the model on the NPU hardware accelerator, add the -external\_delegate\_path=/usr/lib/libethosu\_delegate.so command line argument. For example, to run the model accelerated on the NPU hardware using Ethos-U Delegate, use this command:

```
$ ./label_image -m mobilenet_v1_1.0_224_quant.tflite
-i grace_hopper.bmp -l labels.txt --external_delegate_path=/usr/lib/
libethosu_delegate.so
```

## 2.7.3 Running the example on the i.MX 9 platform with Neutron-S

To use the hardware Acceleration on i.MX 9 with Neutron-S NPU, convert the model using the neutronconverter using the eIQ Toolkit. For details, see the <u>eIQ Toolkit documentation</u>.

• Run the example with option:

```
--external delegate path=/usr/lib/libneutron delegate.so
```

• Run the Python example with option:

```
-e /usr/lib/libneutron delegate.so
```

## 2.7.4 Running the Python example

Alternatively, the example using the TensorFlow Lite interpreter-only Python API can be run. The example file is located at:

/usr/bin/tensorflow-lite-2.18.0/examples

To run the example using the predefined command line arguments, use the following command:

\$ python3 label image.py

The output should be as follows:

```
Warm-up time: 159.1 ms
Inference time: 156.5 ms
0.878431: military uniform
0.027451: Windsor tie
0.011765: mortarboard
0.011765: bulletproof vest
0.007843: sax
```

The Python example supports external delegates also. The switch --ext\_delegate <PATH> and -- ext\_delegate\_options <EXT\_DELEGATE\_OPTIONS>, can be used to specify the external delegate library and optionally its arguments.

For example, to run the model accelerated on the NPU hardware using Neutron Delegate on i.MX 95, run this command:

\$ python3 label image.py --ext delegate /usr/lib/libneutron delegate.so

The output should be as follows:

```
Loading external delegate from /usr/lib/libneutron_delegate.so with args: {}
INFO: NeutronDelegate delegate: 29 nodes delegated out of 31 nodes with 1
partitions.
Warm-up time: 1.9 ms
Inference time: 1.7 ms
0.850980: military uniform
0.058824: Windsor tie
0.011765: bulletproof vest
0.007843: bow tie
0.007843: mortarboard
```

## 2.7.5 Running the example on the i.MX 95 platform using GPU

To run label\_image using the GPU, execute the following command:

```
$ ./label_image -m mobilenet_v1_1.0_224_quant.tflite -i grace_hopper.bmp -l
labels.txt --use_gpu=true --gpu_backend=cl --gpu_precision_loss_allowed=true --
gpu_experimental_enable_quant=true --gpu_inference_for_sustained_speed=false
```

## 2.8 Running benchmark applications

A Yocto Linux BSP image with machine learning layer included by default contains a pre-installed benchmarking application. It performs a simple TensorFlow Lite model inference and prints benchmarking information. The application binary file is located at:

/usr/bin/tensorflow-lite-2.18.0/examples

Benchmarking instructions are as follows:

To run the benchmark with computation on CPU, use the following command:

\$ ./benchmark model --graph=mobilenet v1 1.0 224 quant.tflite

You can optionally specify the number of threads with the --num\_threads=X parameter to run the inference on multiple cores. For highest performance, set X to the number of cores available.

The output of the benchmarking application should be similar to:

```
STARTING!
Log parameter values verbosely: [0]
Graph: [mobilenet_v1_1.0_224_quant.tflite]
Loaded model mobilenet_v1_1.0_224_quant.tflite
Going to apply 0 delegates one after another.
The input model file size (MB): 4.27635
Initialized session in 3.051ms.
Running benchmark for at least 1 iterations and at least 0.5 seconds but
terminate if exceeding 150 seconds.
count=4 first=160408 curr=155384 min=155384 max=160408 avg=156869 std=2076
Running benchmark for at least 50 iterations and at least 1 seconds but
terminate if exceeding 150 seconds.
count=50 first=155586 curr=155424 min=155274 max=155622 avg=155443 std=81
Inference timings in us: Init: 3051, First inference: 160408, Warmup (avg):
156869, Inference (avg): 155443
Note: as the benchmark tool itself affects memory footprint, the following is
only APPROXIMATE to the actual memory footprint of the model at runtime. Take
the information at your discretion.
Peak memory footprint (MB): init=4.49219 overall=10.6133
```

To run the inference without the XNNPACK delegate, add the --use xnnpack=false switch.

To run the inference using the GPU/NPU hardware accelerator, use the --external delegate path switch:

- For VX Delegate on i.MX 8: --external delegate path=/usr/lib/libvx delegate.so
- For Ethos-U Delegate on i.MX 93: --external delegate path=/usr/lib/libethosu delegate.so
- For Neutron Delegate on i.MX 95: --external\_delegate\_path=/usr/lib/ libneutron delegate.so

To run the inference using the Arm Mali G310 GPU using the GPU Delegate on i.MX 95:

```
--use_gpu=true --gpu_backend=cl --gpu_precision_loss_allowed=true --
gpu_experimental_enable_quant=true --gpu_inference_for_sustained_speed=false
```

The output with GPU/NPU module acceleration enabled (for VX Delegate) should be similar to:

```
STARTING!
Log parameter values verbosely: [0]
Graph: [mobilenet v1 1.0 224 quant.tflite]
```

# UG10166

## i.MX Machine Learning User's Guide

```
External delegate path: [/usr/lib/libvx delegate.so]
Loaded model mobilenet_v1_1.0_224_quant.tflite
Vx delegate: allowed builtin code set to 0.
Vx delegate: error during init set to 0.
Vx delegate: error during prepare set to 0.
Vx delegate: error during invoke set to 0.
EXTERNAL delegate created.
Going to apply 1 delegates one after another.
Explicitly applied EXTERNAL delegate, and the model graph will be completely
 executed by the delegate.
The input model file size (MB): 4.27635
Initialized session in 13.437ms.
Running benchmark for at least 1 iterations and at least 0.5 seconds but
terminate if exceeding 150 seconds.
W [HandleLayoutInfer:257]Op 18: default layout inference pass.
count=1 curr=4586473
Running benchmark for at least 50 iterations and at least 1 seconds but
 terminate if exceeding 150 seconds.
count=398 first=2541 curr=2419 min=2419 max=2549 avg=2467.87 std=13
Inference timings in us: Init: 13437, First inference: 4586473, Warmup (avg):
 4.58647e+06, Inference (avg): 2467.87
Note: as the benchmark tool itself affects memory footprint, the following is
 only APPROXIMATE to the actual memory footprint of the model at runtime. Take
 the information at your discretion.
Peak memory footprint (MB): init=7.24609 overall=34.0117
```

The delegates are not required to support the full set of operators defined by the TensorFlow Lite runtime. If the model contains such an operation, which is not supported by the particular delegate, this operation execution falls back to CPU using the TensorFlow Lite reference kernels. This way the computational graph represented by the model gets divided into segments and each segment is executed . The graph segmentation or also called graph partitioning is the process, where the computational graph defined by the model is divided into smaller segments (or partitions) and each of them is executed via the delegate or on the CPU using reference kernels (CPU fallback), based on operation supported by the delegate.

The benchmark application is also useful to check the optional segmentation of the models if accelerated on GPU/NPU hardware accelerator. For this purpose, the combination of the --enable\_op\_profiling=true and --max\_delegated\_partitions=<br/>tig number> (e.g., 1000) options can be used.

Which generates detailed profiling information, such as:

Profiling Info for Bench	mark Initiali	zation:				
[node type] ModifyGraphWithDelegate AllocateTensors	[start] 0.000 4.528	[first] 4.597 0.198	[avg ms] 4.597 0.101	[%] 95.791% 4.209%	[cdf%] 95.791% 100.000%	
[node type] ModifyGraphWithDelegate AllocateTensors Number of nodes executed	Top by Compu [start] 0.000 4.528 : 2	(1111) [first] 4.597 0.198	ne ======== [avg ms] 4.597 0.101	[%] 95.791% 4.209%	[cdf%] 95.791% 100.000%	
[Node type] ModifyGraphWithDelegate AllocateTensors Timings (microseconds): Memory (bytes): count=0 2 nodes observed Operator-wise Profiling	=== Summary k [count][avg m 1 4.5 1 0.2 count=1 curr= Info for Regu	by node ty ns] [avg % 597 95.791 202 4.209 =4799	pe ======== ] [cdf %] [r % 95.791% ( % 100.000% mark Runs:	nem KB] [tir 584.000 0.000	nes called] 1 2	
[node type] [ Vx Delegate 0.000 RESIZE_BILINEAR	start] [fi 14.890 14.896 1	14.894	avg ms] 11.349% 1.331	[%] 11.34 1.014%	[cdf%] 12.363%	
G10166	All information p	provided in this documer	nt is subject to legal disclain	ners.	© 2025 NXP B.V.	All rights

Vx Delegate 16.2	27 2.944	2.909	2.216%	14	.579%	
RESIZE BILINEAR	19.137	0.279	0.277	0.211%	14.790%	
RESIZE BILINEAR	19.415	44.316	44.496	33.905%	48.695%	
- ARG MAX	63.912	67.438	67.332	51.305%	100.000%	
	===== Top by (	Computation '	Time ======			
[node type]	[start]	[first] [	avg ms]	[%]	[cdf%]	
ARG MAX	63.912	67.438	67.332	51.305%	51.305%	
RESIZE BILINEAR	19.415	44.316	44.496	33.905%	85.210%	
Vx Delegate 0.	000 14.890	14.894	11.349%	96	.559%	
Vx Delegate 16.	227 2.944	2.909	2.216%	98	.775%	
RESIZE BILINEAR	14.896	1.331	1.331	1.014%	99.789%	
RESIZE BILINEAR	19.137	0.279	0.277	0.211%	100.000%	
Number of nodes exec	uted: 6					
	===== Summary	y by node ty	pe ========	==========		
[Node type]	[count] [avg n	ns] [avg %]	[cdf %] [	mem KB] [	times called]	
ARG MAX	1 67.3	32 51.306%	51.306%	0.000	1	
RESIZE BILINEAR	3 46.1	.02 35.129%	86.435%	0.000	3	
Vx Delegate 2	17.802 13.5	565% 100.00	0.000		2	
Timings (microsecond	ls): count=8 fi	rst=131198	curr=130580	min=13058	0 max=132766 a	vg=131238
std=616						
Memory (bytes): coun	t=0					
6 nodes observed						

Based on section "Number of nodes executed" in the output, it can be determined which part of the computation graph was executed on GPU/NPU hardware accelerator. Every node except Vx Delegate falls back to CPU. In the example above, the ARG\_MAX and RESIZE\_BILINEAR nodes fall back to CPU.

## 2.9 Post training quantization using TensorFlow Lite converter

TensorFlow offers several methods for model quantization:

- Post training quantization with TensorFlow Lite Converter
- Quantization aware training using Model Optimization Toolkits and TensorFlow Lite Converter
- Various other methods available in previous TensorFlow releases

## Note:

The model quantization is also supported by the "eIQ Toolkit". See also eIQ Toolkit User's Guide (EIQTUG).

Covering all of them is beyond the scope of this documentation. This section describes the approach for the post training quantization using the TensorFlow Lite Converter.

The Converter is available as a part of standard TensorFlow desktop installation. It is used to convert and optionally quantize TensorFlow models into TensorFlow Lite model format (\*.tflite). There are two options how to use the tool:

- The Python API (recommended)
- Command line script

The post training quantization using the Python API is described in this chapter. The documentation useful for model conversion and quantization is available here:

- Python API documentation: <u>https://www.tensorflow.org/versions/r2.15/api\_docs/python/tf/lite/TFLiteConverter</u>
- Guide for model conversion: <u>www.tensorflow.org/lite/convert</u>
- Guide for model quantization: https://www.tensorflow.org/lite/performance/post\_training\_quantization
- Guide for model optimization: <u>https://www.tensorflow.org/model\_optimization</u>

## Note:

The guides on TensorFlow page usually covers the most up to date version of TensorFlow, which might be different from the version available in the NXP eIQ. To see what features are available, check the corresponding API for the specific version of the TensorFlow or TensorFlow Lite.

The current version of the TensorFlow Lite available in the NXP eIQ is 2.15.0. It is recommended to use the TensorFlow Lite converter from corresponding TensorFlow version. The TensorFlow Lite runtime should be compatible with models generated by previous version of TensorFlow Lite Converter, however this backward compatibility is not guaranteed. Usage of successive version of TensorFlow Lite converter shall be avoided.

The 2.15.0 version of the converter has the following properties:

- In the post training quantization regime, the per-channel quantization is the only option. The per-tensor quantization is available only in connection with quantization aware training.
- Input and output tensors quantization is supported by setting the required data type in inference\_input\_type and inference\_output\_type.
- TOCO or MLIR based conversions are available. This is controlled by the experimental\_new\_converter attribute. As TOCO is becoming obsolete, MLIR-based conversion is already set by default in the 2.15.0 version of the converter.

MLIR converter uses dynamic tensor shapes, what means the batch size of the input tensor is unspecified. Dynamic tensor shapes are not supported by the GPU and NPU hardware accelerators and this shall be turned off. Standard installation of TensorFlow does not provide API to control the dynamic tensor shape feature, but can be deactivated in the tensorflow installation, as follows. Locate the cpython-installdir>/site-packages/tensorflow/lite/python/lite.py file and change the private method TFLiteConverterBase. is unknown shapes allowed (self) to return False value, as follows:

```
def _is_unknown_shapes_allowed(self):
    # Unknown dimensions are only allowed with the new converter.
    # Return self.experimental_new_converter
    # Disable unknown dimensions support.
    return False
```

## Note:

MLIR is a new NN compiler used by TensorFlow, which supports quantization. Before MLIR, quantization was performed by TOCO (or TOCO Converter), which is now obsolete. See <u>https://www.tensorflow.org/api\_docs/</u>python/tf/compat/v1/lite/TocoConverter. For details about MLIR, see <u>https://www.tensorflow.org/mlir</u>.

#### Note:

Do not use the dynamic range method for models being run on NN accelerators (GPU or NPU). It converts only the weights to 8-bit integers, but retains the activations in fp32, which results in the inference running in fp32 with an additional overhead for data conversion. In fact, the inference is even slower compared to a fp32 model, because the conversion is done on the fly.

For the full-integer post training quantization, a representative dataset is needed. The proper choice of samples in representative dataset highly influences the accuracy of the final quantized model. The best practices for creating the representative dataset are:

- Use train samples for which the original floating points model has very good accuracy, based on metrics the model used (e.g., SoftMax score for classification models, IOU for object detection models, etc.).
- There shall be enough samples in representative dataset.
- The size of representative dataset and the specific samples available in it are considered as hyperparameters to tune, with respect of the required model accuracy.

# 2.10 TensorFlow Lite for Microcontrollers on Xtensa HiFi4 core

TensorFlow Lite for Microcontrollers (TFLM) is a lightweight re-implementation of the TensorFlow Lite library for microcontroller CPU cores and NN accelerators (like the Xtensa HiFi4 core on i.MX 8ULP or Arm Ethos-U on i.MX 93). Compared to TensorFlow Lite, it uses less memory, has no C/C++ library dependencies and uses only static memory allocation. On the other hand, the list of supported operators is more limited and optimized kernels are available only for Cortex-M and Xtensa cores or the ARM Ethos-U accelerator. The main purpose of TFLM on the i.MX platform is low-power applications.

To use TFLM on the Xtensa HiFi4 core, the DSP firmware has to be rebuilt with the TFLM library and a TensorFlow Lite model included. As the Xtensa HiFi4 core is also used for audio encoding/decoding, the TFLM library has to be wrapped into an Xtensa Audio Framework (XAF) component to allow simultaneous audio and model inference execution. Moreover, the XAF client/server protocol implements input and output buffer passing to and from the CPU core via the Linux XAF API. The DSP firmware and usage example source codes are available at <a href="https://github.com/NXP/imx-audio-framework">https://github.com/NXP/imx-audio-framework</a>. See the DSP User's Guide in the docs subfolder for information on toolchain setup and build instructions.

To build the DSP firmware with the TFLM library (after the toolchain is installed), use the following Makefile options:

make PLATF=imx8ulp TFLM=1 DSP FIRMWARE

The command produces a hifi4\_tflm\_imx8ulp.bin file which has to be copied to the /lib/firmware/imx/dsp folder of the Yocto Linux BSP image.

To build the TFLM usage example for Linux, use the following Makefile options:

make PLATF=imx8ulp TFLM=1 UNIT TEST

The command compiles the unit\_test/src/dsp\_tflm\_test.c source file and produces a dsp\_tflm\_test.out binary executable file which demonstrates a simple keyword detection application processing a built-in static audio buffer with "yes" and "no" speech command data samples.

By default, TFLM included in the DSP firmware is compiled with reference kernel implementations due to licensing. To improve the library performance on Xtensa HiFi4 cores, the library has to be built with proprietary licensed optimized kernel implementations provided by Cadence at <a href="https://github.com/foss-xtensa/nnlib-hifi4">https://github.com/foss-xtensa/nnlib-hifi4</a> (see the license file in the GitHub repository). Add the OPTIMIZED\_KERNEL\_DIR=xtensa option into the <a href="https://github.com/foss-xtensa/nnlib-hifi4">dsp\_framework/tensorflow\_lite\_micro.inc file to automatically download the Cadence library and build TFLM with the optimized kernels:</a>

```
cd $(SRC_DIR)/tflite-micro && make -f tensorflow/lite/micro/tools/make/
Makefile TARGET=xtensa TARGET_ARCH=hifi4 OPTIMIZED_KERNEL_DIR=xtensa
XTENSA USE LIBC=true microlite
```

A DSP firmware file with the same name as previously is produced, which has to be copied to the /lib/ firmware/imx/dsp folder of the Yocto Linux BSP image.

# 3 ONNX Runtime

ONNX Runtime is an open-source inference engine to run ONNX models, which enables the acceleration of machine learning models across all of your deployment targets using a single set of API. Source codes are available at <a href="https://github.com/nxp-imx/onnxruntime-imx">https://github.com/nxp-imx/onnxruntime-imx</a>.

#### Note:

For the full list of the CPU supported operators, see the 'operator kernels' documentation section: <u>OperatorKernels</u>.

Features:

- ONNX Runtime 1.17.1
- Multithreaded computation with acceleration using Arm Neon SIMD instructions on Cortex-A cores provided by the CPU execution provider
- C++ and Python API (supported Python version 3)
- ONNX Runtime 1.17.1 supports ONNX 1.15 and Opset version 20.

# 3.1 ONNX Runtime software stack

The ONNX Runtime software stack is shown in the following figure. The ONNX Runtime supports computation on the following HW units:

· CPU Arm Cortex-A cores using CPU execution provider



# 3.2 ONNX model test

ONNX Runtime provides a tool that can run the collection of standard tests provided in the ONNX Model Zoo. The tool named onnx\_test\_runner is installed in /usr/bin/onnxruntime-1.17.1.

ONNX models are available at <u>https://github.com/onnx/models</u> and consist of models and sample test data. Because some models require a lot of disk space, it is advised to store the ONNX test files on a larger partition, as described in the SD card image flashing section.

Here is an example with the steps required to run the mobilenet version 2 test:

• Download and unpack the mobilenet version 2 test archive to some folder, for example to/home/root:

```
$ cd /home/root
$ wget https://github.com/onnx/models/raw/main/vision/classification/mobilenet/
model/mobilenetv2-7.tar.gz
$ tar -xzvf mobilenetv2-7.tar.gz
$ ls ./mobilenetv2-7
mobilenetv2-7.onnx test_data_set_0
```

# UG10166

## i.MX Machine Learning User's Guide

• Run the onnx\_test\_runner tool providing mobilenetv2-7 folder path and setting the execution provider:

```
$ /usr/bin/onnxruntime-1.17.1/onnx_test_runner -j 1 -c 1 -r 1 -e cpu ./
mobilenetv2-7/
result:
Models: 1
Total test cases: 3
Succeeded: 3
Not implemented: 0
Failed: 0
Stats by Operator type:
Not implemented(0):
Failed:
Failed Test Cases:
$
```

## Note:

Use onnx test runner -h for the full list of supported options.

# 3.3 ONNX performance test

To run model benchmarks, ONNX Runtime provides a tool that measures performance. The tool named onnxruntime\_perf\_test is installed in /usr/bin/onnxruntime-1.17.1. In order to run it, the user must provide an *.onnx* model file together with test data. To benchmark the SqueezeNet model running a single iteration using the CPU execution provider, run to the following command:

```
/usr/bin/onnxruntime-1.17.1/onnxruntime_perf_test /usr/bin/onnxruntime-1.17.1/
squeezenet/model.onnx -r 1 -e cpu
```

## Note:

Use onnxruntime perf test -h for the full list of supported options.

# 4 PyTorch

PyTorch is a scientific computing package based on Python that facilitates building deep learning projects using power of Graphics Processing Units (GPUs).

Features:

- PyTorch 2.3.0
- Python version 3 supported
- Deep neural networks built on a tape-based autograd system

## Note:

Only the CPU is supported. By default, the PyTorch runtime is running with floating point model. To enable quantized model, the quantized engine should be specified explicitly as follows:

torch.backends.quantized.engine = 'qnnpack'

# 4.1 Installing PyTorch

PyTorch is available on the PyPI registry. To install PyTorch on the BSP, run the following command:

\$ pip install torch torchvision

# 4.2 Running image classification example

There is an example located in the examples folder, which requires urllib, PIL, and maybe some other Python3 modules depending on your image. You may install the missing modules using pip3.

\$ cd /usr/bin/pytorch/examples

To run the example with inference computation on the CPU, use the following command. There are no arguments and the resources will be downloaded automatically by the script:

\$ python3 pytorch mobilenetv2.py

The output should be similar as follows:

```
File does not exist, download it from
https://download.pytorch.org/models/mobilenet_v2-b0353104.pth
... 100.00%, downloaded size: 13.55 MB
File does not exist, download it from
https://raw.githubusercontent.com/Lasagne/Recipes/master/examples/resnet50/
imagenet_classes.txt
... 100.00%, downloaded size: 0.02 MB
File does not exist, download it from
https://s3.amazonaws.com/model-server/inputs/kitten.jpg
... 100.00%, downloaded size: 0.11 MB
('tabby, tabby cat', 46.34805679321289)
('Egyptian cat', 15.802854537963867)
('lynx, catamount', 1.1611212491989136)
('lynx, catamount', 1.1611212491989136)
('tiger, Panthera tigris', 0.20774540305137634)
```

# 5 TVM

Apache TVM is an open source machine learning compiler framework for CPUs, GPUs, and NPUs. It aims to enable machine learning engineers to optimize and run computations efficiently on any hardware backend.

Features:

- TVM 0.7.0
- Compilation of deep learning models into minimum deployable modules
- Infrastructure to automatic generate and optimize models on more backend with better performance
- Support for i.MX 8M Plus platforms with OpenVX library
- TVM builder supported for Ubuntu 18.04, x86\_64 platform

## Note:

For more detailed information, see <u>TVM Documentation</u>.

## 5.1 TVM software workflow

The pre-trained model will be transformed into the Relay IR and passed through to the TVM model optimizations like constant-folding, memory planning, and finally passed to a codegen phase. In this phase, the operators supported by the target device are transformed as intrinsic calls into the offloading library which connects the model accelerator devices such as GPU/NPU.



# 5.2 Getting started

## 5.2.1 Running example with RPC verification

TVM provides the Remote Procedure Call (RPC) capability to run a model on the remote device.

User can run examples at tests/python/contrib/test\_vsi\_npu with RPC verification. The model running result on device will be verified against the result on host with same input.

• Launch the RPC server on the device

\$ python3 -m tvm.exec.rpc server --host 0.0.0.0 --port=9090

• Export the system variables:

```
$ export TVM_HOME=/path/to/tvm
$ export PYTHONPATH=$TVM_HOME/python
```

```
• Run the specified models on the host PC:
```

```
$ python3 tests/python/contrib/test_vsi_npu/test_rpc_tflite_models.py -i
{device ip} -m mobilenet v2 1.0 224 quant
```

• Run all supported TensorFlow Lite models on the host PC:

```
$ python3 tests/python/contrib/test_vsi_npu/test_rpc_tflite_models.py -i
{device_ip}
```

**Note:** This test will download the model automatically, please be sure the network can access the public internet. Example scripts may import additional Python libraries. Please check scripts and make sure they are installed correctly.

To test pytorch/onnx/keras model, additional python packages needs to be installed on the host PC:

```
$ python3 -m pip install torch==1.7.0 torchvision==0.8.1
$ python3 -m pip install onnx=1.8.1 onnxruntime==1.8.1
$ python3 -m pip install tensorflow==2.5.0
```

## 5.2.2 Running example individually on device

In this mode, the model is compiled on the host offline and saved as model.so. Please refer tests/python/ contrib/test vsi npu/compile tflite models.py to compile a TensorFlow Lite model on the host.

Below script snippet shows how to load and run a compiled model at the device:

```
ctx = tvm.cpu(0)
# load the compiled model
lib = tvm.runtime.load_module(args.model)
m = graph_runtime.GraphModule(lib["default"](ctx))
# set inputs
data = get_img_data(args.image, (args.input_size, args.input_size),
args.data_type)
m.set_input(args.input_tensor, data)
# execute the model
m.run()
# get outputs
tvm_output = m.get_output(0)
```

Please refer tests/python/contrib/test\_vsi\_npu/label\_image.py to a complete label image example with pre-processing of image decoding and post-processing to generate label.

## 5.3 How to build TVM stack on host

Conceptually, TVM can be split into two parts:

- TVM build stack: compiles the deep learning model at host
- TVM runtime: loads and interprets the model at device

This build stack is using the LLVM to cross-compile the generated source as a deployable dynamic library for device. Please, follow the <u>LLVM Doc</u> to install LLVM on the host. If installed successfully, llvm-config should be found under /usr/bin.

To build the tvm, please be sure below dependence packages installed on the host:

- cmake
- python3-dev
- build-essential
- Ilvm-dev
- g++-aarch64-linux-gnu

- libedit-dev
- libxml2-dev
- python3-numpy
- python3-attrs
- python3-tflite

For Ubuntu 18.04, the user could use below commands to install all dependences:

```
$ sudo apt-get update
$ sudo apt-get install -y python3 python3-dev python3-setuptools
$ sudo apt-get install -y cmake llvm llvm-dev g++-aarch64-linux-gnu gcc-aarch64-
linux-gnu
$ sudo apt-get install -y libtinfo-dev zlib1g-dev build-essential libedit-dev
libxml2-dev
$ python3 -m pip install numpy decorator scipy attrs six tflite
```

Follow below instructions to build TVM stack on the host:

```
$ export TOP_DIR=`pwd`
$ git clone --recursive https://github.com/nxp-imx/eiq-tvm-imx/ tvm-host
$ cd tvm-host
$ mkdir build
$ cp cmake/config.cmake build
$ cd build
$ cd build
$ sed -i 's/USE_LLVM\ OFF/USE_LLVM\ \/usr\/bin\/llvm-config/' config.cmake
$ cmake ..
$ make tvm -j4 # make tvm build stack
```

# 5.4 Supported models

The following models are verified with TVM.

Table 2. TVM models ZOO

Model	float32	int8	Input size
mobilenet_v1_0.25_128	mobilenet_v1_0.25_128	mobilenet_v1_0.25_128_quan	128
mobilenet_v1_0.25_224	mobilenet_v1_0.25_224	mobilenet_v1_0.25_224_quan	224
mobilenet_v1_0.5_128	mobilenet_v1_0.5_128	mobilenet_v1_0.5_128_quant	128
mobilenet_v1_0.5_224	mobilenet_v1_0.5_224	mobilenet_v1_0.5_224_quant	224
mobilenet_v1_0.75_128	mobilenet_v1_0.75_128	mobilenet_v1_0.75_128_quan	128
mobilenet_v1_0.75_224	mobilenet_v1_0.75_224	mobilenet_v1_0.75_224_quan	224
mobilenet_v1_1.0_128	mobilenet_v1_1.0_128	mobilenet_v1_1.0_128_quant	128
mobilenet_v1_1.0_224	mobilenet_v1_1.0_224	mobilenet_v1_1.0_224_quant	224
mobilenet_v2_1.0_224	mobilenet_v2_1.0_224	mobilenet_v2_1.0_224_quant	224
inception_v1	N/A	inception_v1_224_quant	224
inception_v2	N/A	inception_v2_224_quant	224
inception_v3	inception_v3	inception_v3_quant	299
inception_v4	inception_v4	inception_v4_299_quant	299
deeplab_v3_257_mv_gpu	deeplab_v3_256_mv_gpu	N/A	257

Table 2 TVM models 700 continued

## i.MX Machine Learning User's Guide

Table 2. TVW models 200continued						
Model	float32	int8	Input size			
deeplab_v3_mnv2_pascal	N/A	deeplab_v3_mnv2_pascal	513			
ssdlite_mobiledet	ssdlite_mobiledet_cpu_320x32		320			

#### **NN Execution on Hardware Accelerators** 6

## 6.1 Hardware acceleration on i.MX 8 Series

## 6.1.1 Hardware accelerator description

The i.MX 8 class devices are deployed with two kind of NN accelerators (see also the figure below):

- Neural Processing Unit (NPU)
- Graphics Processing Unit (GPU)

Neural processing unit is optimized for fixed point arithmetic, in 8-bit and 16-bit width. For optimal performance on the NPU, quantized models shall be used.

Graphics processing unit is optimized for fixed point arithmetic and half precision floating point arithmetic. For optimal performance on the GPU, quantized models or floating-point models with half precision shall be used.

### Note:

The TensorFlow Lite framework enables to compute the floating-point models directly in 16-bit half precision arithmetic.



Figure 6. NN accelerator SW stack

Interface to NPU/GPU HW accelerator is provided via the OpenVX v1.3 with NN Extensions. OpenVX is an open, royalty-free standard for cross platform acceleration of computer vision applications. It provides:<sup>2</sup>

- · A library of predefined and customizable vision functions
- A graph-based execution model to combine function enabling both task and data independent execution
- · A set of memory objects that abstract the physical memory

Open VX defines a C-application programming interface for building, verifying and coordinating graph execution and accessing memory objects. More information about OpenVX can be find on the OpenVX home page.

<sup>2</sup> OpenVX 1.3 specification: https://registry.khronos.org/OpenVX/specs/1.3/html/OpenVX\_Specification\_1\_3.html

## Note:

*In the current OpenVX driver implementation, the maximum number of nodes supported in OpenVX graph is 2048.* 

## 6.1.2 Profiling on hardware accelerators

This section describes how to enable profiler on the GPU/NPU, and how to capture logs.

- 1. Stop the EVK board in the U-Boot by pressing Enter.
- 2. Update mmcargs by adding galcore.showArgs=1 and galcore.gpuProfiler=1.

```
u-boot=> editenv mmcargs
edit: setenv bootargs ${jh_clk} ${mcore_clk} console=${console} root=
${mmcroot} galcore.showArgs=1 galcore.gpuProfiler=1
u-boot=> boot
```

- 3. Boot the board and wait for the Linux OS prompt.
- 4. The following environment flags should be enabled before executing the application.

VIV\_VX\_DEBUG\_LEVEL and VIV\_VX\_PROFILE flags should always be 1 during the process of profiling. The CNN\_PERF flag enables the driver's ability to generate per layer profile log. NN\_EXT\_SHOW\_PERF shows the details of how compiler estimates performance and determines tiling based on it.

export CNN\_PERF=1 NN\_EXT\_SHOW\_PERF=1 VIV\_VX\_DEBUG\_LEVEL=1 VIV\_VX\_PROFILE=1

- 5. Capture the profiler log. We use the sample ML example part of standard NXP Linux release to explain the following section.
  - TensorFlow Lite profiling

Run the TensorFlow Lite application with GPU/NPU backend as follows:

```
$ cd /usr/bin/tensorflow-lite-2.18.0/examples $ ./label_image -
m mobilenet_v1_1.0_224_quant.tflite -t 1 -i grace_hopper.bmp -l
labels.txt --external_delegate_path=/usr/lib/libvx_delegate.so -v 0 >
viv_test_app_profile.log 2>&1
```

The log captures detailed information of the execution clock cycles and DDR data transmission in each layer.

#### Note:

The average time for inference might be longer than usual, as the profiler overhead is added.

## 6.1.3 Hardware accelerators warmup time

For TensorFlow Lite, the initial execution of model inference takes longer time, because of the model graph initialization needed by the GPU/NPU hardware accelerator. The initialization phase is known as warmup. This time duration can be decreased for subsequent application that runs by storing on disk the information resulted from the initial OpenVX graph processing. The following environment variables should be used for this purpose:

VIV\_VX\_ENABLE\_CACHE\_GRAPH\_BINARY: flag to enable/disable OpenVX graph caching

VIV VX CACHE BINARY GRAPH DIR: set location of the cached information on disk

For example, set these variables on the console in this way:

```
export VIV_VX_ENABLE_CACHE_GRAPH_BINARY="1"
export VIV_VX_CACHE_BINARY_GRAPH_DIR=`pwd`
```

By setting up these variables, the result of the OpenVX graph compilation is stored on disk as network binary graph files (\*.nb). The runtime performs a quick hash check on the network and if it matches the \*.nb file hash, it loads it into the NPU memory directly. These environment variables need to be set persistently, for

example, available after reboot. Otherwise, the caching mechanism is bypassed even if the \*.nb files are available.

The iterations following the graph initialization are performed many times faster. When evaluating the performance of an application running on GPU/NPU, the time should be measured separately for warmup and inference. Warmup time usually affects only the first inference run. However, depending on the machine learning model type, it might be noticeable for the first few inference runs. Some preliminary tests must be done to make a decision on what to consider warmup time. When this phase is well delimited, the subsequent inference runs can be considered as pure inference and used to compute an average for the inference phase.

*Note:* OpenVX graph caching is not available on *i*.MX 8QuadMax platform.

## 6.1.4 Switching between GPU and NPU

Some platforms are deployed with both 3D GPU and NPU hardware accelerators. Both can be used for execution of the OpenVX graph (i.e. for ML inference). To differentiate between the GPU and the NPU, there is an environmental variable USE\_GPU\_INFERENCE. The variable is directly read by the HW acceleration driver.

The behavior is as follows:

- If USE GPU INFERENCE=1, the graph is executed on the GPU
- Otherwise, the graph is executed on the NPU (if available)

By default, the NPU is used for OpenVX graph execution.

Example with TensorFlow Lite:

```
$ USE_GPU_INFERENCE=1 ./label_image -m mobilenet_v1_1.0_224_quant.tflite
-i grace_hopper.bmp -l labels.txt --external_delegate_path=/usr/lib/
libvx delegate.so
```

## 6.1.5 Per-tensor vs. per-channel quantization

Both per-tensor and per-channel quantizations are supported by the GPU and NPU hardware accelerators on i.MX 8 Series.

The NPU on i.MX 8M Plus is optimized for per-tensor quantization. When running per-channel quantized models, the accelerator must involve additional compute. Therefore, the performance might be lower compared to per-tensor quantized models. The additional compute might also introduce a small accuracy error. The actual impact for both accuracy and latency depends on the model used.

# 6.2 Hardware acceleration with Ethos-U on i.MX 93 platform

Ethos-U65 is a neural processing unit (NPU) designed to accelerate ML inference in area-constrained embedded and IoT devices from Arm. This NPU is integrated with NXP i.MX 93 processor and works in concert with the Cortex-M core and on-chip SRAM of the SoC. Currently, it provides the following main features:

- Running at 1 GHz and providing 0.5 Tops computation power (256 MAC/cycle).
- Targets quantized Convolutional Neural Networks (CNN) and supports 8 bit weights and 8/16 bit activations.
- · Supports TensorFlow Lite (TFLite) inference with fallback to Cortex-A.
- Supports TensorFlow Lite Micro (TFLite-Micro) inference with fallback to Cortex-M.
- Supports inference API to offload the entire model to TFLite-Micro and NPU on Cortex-M.
- Supports TFLite API to offload the customized "ethos-u" operator to NPU on Cortex-M.
- Provides Vela model tool to optimize the model performance and memory usage for the Ethos-U65 target.

## 6.2.1 Ethos-U subsystem overview

This i.MX 93 machine learning system involves several HW components working collaboratively to support the acceleration of the tensor computation of an ML model: Cortex-A, Cortex-M, Messaging Unit (MU), and Ethos-U NPU.



Figure 7. Ethos-U subsystem overview

The Cortex-A55 is responsible for loading the ML model, capture and pre-process the inputs with Linux OS and rich libraries. The Cortex-M is the controller of the attached Ethos-U NPU and it prepares the offloading descriptor for the NPU and triggers the NPU execution. It also provides the un-supported kernels execution for NPU. The MU is the message unit IP to facilitate the core communication between Cortex-A and Cortex-M.

## 6.2.2 Ethos-U software architecture

The software for Ethos-U support includes three main components, as shown in the following figure.



Figure 8. Ethos-U software architecture

- Vela model compiler: offline tool to compile the TensorFlow Lite model graph for Ethos-U. The compiler replaces supported operators in the model with custom "ethos-u" operator containing the command stream for Ethos-U NPU. The output of the compiler is a modified TensorFlow Lite model graph for TensorFlow Lite/ TensorFlow Lite-Micro inference engines. This is only required for Inference API.
- Cortex-A SW stack for Linux: containing MPU inference engine (TensorFlow Lite), driver library, and Linux kernel driver.
- Cortex-M SW stack: containing MCU inference engine SW (TensorFlow Lite-Micro, CMSIS-NN) and NPU driver.

The typical inference workflow is as follows:

- Converts the TensorFlow Lite model into Vela model using the Vela model compiler and generates the optimized version for Ethos-U NPU.
   Note: For TensorFlow Lite inference engine with Ethos-U delegate, this step is optional. The Ethos-U delegate supports both TensorFlow Lite model and Vela compiled model. Using the Ethos-U delegate increases the warm-up time for model execution. The model needs to be compiled at runtime. Models
- *precompiled with Vela brings warm-up time decrease.* 2. The optimized model is inferred either by:
  - a. TensorFlow Lite Ethos-U delegate, which recognizes the custom "ethos-u" operator in Vela compiled model, allocates the buffer for input/output feature map (IFM/OFM) and executes the operator via Ethos-U Linux driver.
  - b. TensorFlow Lite Ethos-U delegate, which recognizes the supported operators in TensorFlow Lite model, compiles the operators to "ethos-u" operator and allocates the buffer for input/output feature map (IFM/ OFM) and executes the operator via Ethos-U Linux driver.
  - c. Inference API, which allocates the buffer for input/output feature map and sends entire model via Ethos-U driver.
- 3. The Ethos-U driver composes the inference task message and sends it over RPMSG to Cortex-M.
- 4. The Ethos-U Runner on Cortex-M dispatches the task to TensorFlow Lite-Micro or Ethos-U driver directly according to the task type.
  - a. If the task type is accelerating the "ethos-u" operator (using the TensorFlow Lite), the Runner calls the Ethos-U driver directly.
  - b. If the task type is accelerating the entire model (using the Inference API), the Runner dispatches the model to TensorFlow Lite-Micro and further calls Ethos-U driver for processing.
- 5. After the Ethos-U driver completes the inference task, it writes the result into the OFM buffer and sends the response back to Cortex-A via RPMSG.

## Note: The model is loaded from Cortex-A and shared with Cortex-M over RPMSG.

The Cortex-M SW is pre-built with both the model and Ethos-U operator acceleration capabilities in a singlebinary firmware. This firmware is integrated into Yocto rootfs and will be loaded automatically when the user starts an inference task using the TensorFlow Lite or Inference API by opening the Ethos-U device.

## 6.2.3 Getting started

In the Yocto rootfs, there are several examples provided to show how to use different APIs to interact with Ethos-U NPU with an image classification inference.

1. Go to the example folder and copy the label.txt and input picture from TensorFlow.

```
$ cd /usr/bin/ethosu/examples
$ cp ../../tensorflow-lite-2.18.0/examples/labels.txt ./
$ cp ../../tensorflow-lite-2.18.0/examples/grace_hopper.bmp ./
```

### 2. Compile the model for Ethos-U using Vela tool, reusing the model

mobilenet\_v1\_1.0\_224\_quant.tflite from /usr/bin/tensorflow-lite-2.18.0/examples/.
If running successfully, an optimized vela model mobilenet\_v1\_1.0\_224\_quant\_vela.tflite is
generated in the output folder.

```
$ vela ../../tensorflow-lite-2.18.0/examples/
mobilenet_v1_1.0_224_quant.tflite
```

3. Run the model with the Inference API (offloads the entire model to TFLite-Micro).

```
$ ./inference_runner -n ./output/mobilenet_v1_1.0_224_quant_vela.tflite -i
grace_hopper.bmp -l labels.txt -o output.txt
```

The following will be printed if no error occurs:

```
Send capabilities request
Capabilities:
    version_status:1
    version:{ major=0, minor=0, patch=0 }
    product:{ major=6, minor=0, patch=6 }
    architecture:{ major=1, minor=0, patch=6 }
    driver:{ major=0, minor=16, patch=0 }
    macs_per_cc:8
    cmd_stream_version:0
    custom_dma:false
Create network
Create inference
Wait for inferences
Inference status: success
Detected: military uniform, confidence:70
```

 Run the converted model with TFLite inference engine using the Ethos-U Delegate (offload the "ethos-u" operator to Ethos-U NPU).

```
$ cd /usr/bin/tensorflow-lite-2.18.0/examples
$ ./label_image -m
../../ethosu/examples/output/mobilenet_v1_1.0_224_quant_vela.tflite
--external_delegate_path=/usr/lib/libethosu_delegate.so
```

The following is printed if no error occurs:

```
INFO: Loaded model ../../ethosu/examples/output/
mobilenet v1 1.0 224 quant vela.tflite
INFO: resolved reporter
Ethosu delegate: device name set to /dev/ethosu0.
Ethosu delegate: timeout set to 6000000000.
Ethosu delegate: enable cycle counter set to 0.
Ethosu delegate: pmu event0 set to 0.
Ethosu delegate: pmu_event1 set to 0.
Ethosu delegate: pmu_event2 set to 0.
Ethosu delegate: pmu event3 set to 0.
EXTERNAL delegate created.
INFO: EthosuDelegate delegate: 1 nodes delegated out of 1 nodes with 1
partitions.
INFO: Applied EXTERNAL delegate.
INFO: invoked
INFO: average time: 3.903 ms
INFO: 0.780392: 653 military uniform
INFO: 0.105882: 907 Windsor tie
INFO: 0.0156863: 458 bow tie
INFO: 0.0117647: 466 bulletproof vest
INFO: 0.00784314: 835 suit
```

## 6.2.4 Vela tool

The vela tool is used to compile a <u>TensorFlow Lite for Microcontrollers</u> neural network model into an optimized version that can run on an embedded system containing an <u>Arm Ethos-U NPU</u>. The optimized model contains TFLite Custom operators for those parts of the model that can be accelerated by the Ethos-U NPU. Parts of the model that cannot be accelerated are left unchanged and run on CPU (Cortex-A or Cortex-M) using an appropriate kernel (such as the <u>Arm</u> optimized <u>CMSIS-NN</u> kernels). After compilation, the optimized model can only be run on an Ethos-U NPU embedded system. The tool also generates performance estimates for the compiled model.

To deploy the neural network (NN) model on Ethos-U, the first step is to use Vela to compile the prepared model. To be accelerated by the Ethos-U NPU, the network operators must be quantized to either 8-bit (unsigned or signed) or 16-bit (signed).

NXP Vela is based on Arm <u>ethos-u-vela</u>. Compared to <u>ethos-u-vela</u>, NXP added more OPs support and reduced some OP constraints.

**Note:** A specific version of Vela is tightly coupled with a specific version of the Ethos-U driver. The compatibility between different Vela versions is not guaranteed.

## 6.2.4.1 Installing the Vela tool

The Vela tool can be run on the i.MX 93 board or Linux PC. It is already available in NXP Yocto rootfs. This section describes how to install it on the X86 Linux PC. The steps are as follows.

1. Get the vela source code.

```
$ git clone https://github.com/nxp-imx/ethos-u-vela.git
```

2. Install with python pip.

```
$ cd ethos-u-vela
$ git checkout lf-6.12.3_1.0.0
$ pip3 install .
```

After all the commands are successful, we can use vela --help to check if Vela is installed successfully.

```
$ vela --version
4.1.0
```

## 6.2.4.2 Compiling the TFLite model

After Vela is installed, the following commands can be used to compile a TFLite model to the optimized version for Ethos-U NPU. The optimized model is stored into the OUTPUT\_DIR ("./output" by default). The output file has the suffix \_vela.tflite. It is also a TFLite model. After the compilation, Vela outputs the detailed log into the console.

**Note:** The Vela expects that the TFLite model is quantized already. Vela supports asymmetric quantization to 8 bit (signed and unsigned) and 16 bit (signed), as defined by TFLite. To accelerate the model operators with Ethos-U NPU, the input model to Vela has to be quantized. Non-quantized operators will fall back to CPU.

The following provides an example for how to compile a model and shows the corresponding output log:

```
$ vela mobilenet_v1_1.0_224_pb_int8.tflite
```

```
Output log:
```

```
Network summary for mobilenet v1 1.0 224 pb int8
Accelerator configuration
                                         Ethos U65 256
System configuration
                                       internal-default
                                        internal-default
Memory mode
Accelerator clock
                                                  1000 MHz
Design peak SRAM bandwidth
                                                 16.00 GB/s
Design peak DRAM bandwidth
                                                  3.75 GB/s
Total SRAM used
                                                381.08 KiB
Total DRAM used
                                               4293.34 KiB
CPU operators = 0 (0.0%)
NPU operators = 60 (100.0\%)
```

# UG10166

## i.MX Machine Learning User's Guide

Average Input Weight Output Total Average Input Weight Output Total	SRAM bandwidth SRAM bandwidth SRAM bandwidth SRAM bandwidth SRAM bandwidth DRAM bandwidth DRAM bandwidth DRAM bandwidth DRAM bandwidth DRAM bandwidth	per input	<pre>4.28 GB/s 7.95 MB/batch 12.61 MB/batch 0.00 MB/batch 20.67 MB/batch 20.67 MB/inference (batch size 1) 3.00 GB/s 5.53 MB/batch 3.92 MB/batch 5.06 MB/batch 14.52 MB/batch 14.52 MB/batch</pre>
Neural r Network NPU cycl SRAM Acc DRAM Acc On-chip Off-chip	DRAM Bandwidth network macs Tops/s les cess cycles cess cycles Flash Access cycles o Flash Access cycles	per input	14.52 MB/Interence (batch Size I) 572406226 MACs/batch 0.24 Tops/s 3937697 cycles/batch 719415 cycles/batch 0 cycles/batch 0 cycles/batch 4821570 cycles/batch
Batch Ir	nference time	4.83 ms,	206.97 inferences/s (batch size 1)

The following is the computational graph after the model (mobilenet\_v1\_1.0\_224\_pb\_int8.tflite) is compiled. Here, Vela encapsulates all supported OPs into one Ethos-U OP.



#### Note:

The Vela tool takes a lot of memory when converting big models, so a swap file is needed when converting big models to avoid the Out Of Memory issue. Use the following commands to add a swap file:

```
dd if=/dev/zero of=/swapfile count=2048 bs=1M
mkswap /swapfile
chmod 0600 /swapfile
swapon /swapfile
```

## 6.2.5 Inference with Ethos-U inference API

The Ethos-U inference API provides the methods to use the Ethos-U NPU on Linux OS without the TensorFlow Lite inference engine. It takes the compiled model and IFM/OFM as inputs, composes an inference task, and dispatches the inferences to the Cortex-M with Ethos-U.

## 6.2.5.1 Ethos-U driver library

The Ethos-U Driver provides a C++ APIs for dispatching the inference to the Ethos-U Linux kernel driver. The library and the corresponding header file is available on Yocto rootfs and SDK:

- /usr/include/ethosu.hpp
- /usr/lib/libethosu.so

The following is the component diagram of Ethos-U Driver library:

- The Device class represents the instance of Ethos-U unit.
- The Buffer class is used to store any data, including the model.
- The Network class represents a model instance bind to specific Device.
- The Inference class represents the inference, which is computation of the computation graph (model) on input data. Notice, the Network class is separated from the Inference class, allowing multiple inferences to share the same network.



Figure 10. Component diagram of Ethos-U Driver library

The inference runner demonstrates how to dispatch inferences to the Ethos-U Linux kernel driver. All the steps described in the sequence diagram below are executed by the inference runner application.

- 1. The Device class obtains a file descriptor handle for the device node (/dev/ethosu<nr>) using the open() system call. The Device class uses ioctl() system calls to manipulate with the underlying device, like buffer and network creation.
- 2. The Network class uses the Device and buffer handles to create a new network object. The model is stored in the Buffer that the network parses to discover the input and output shapes of the network model.
- 3. The Inference class uses the Network object to create an inference. The array of IFM Buffers need to be populated with data before the inference object is created.

The inference object must poll the file descriptor waiting for the inference to complete.
i.MX Machine Learning User's Guide



## 6.2.5.2 Ethos-U Linux kernel driver interface

The Ethos-U Linux kernel driver exposes User-space API (UAPI) for Ethos-U subsystem, and to communicate with the Cortex-M in the Ethos-U subsystem.

The communication with the Ethos-U subsystem is based on message passing in shared memory, and the Linux kernel mailbox APIs for triggering IRQs on the remote CPU, what is the Cortex-M in this case.

The address of the message queues is hard coded in the Cortex-M application, and configured in the DTB for the Ethos-U Linux kernel driver.

When the Linux kernel driver allocates dynamic memory for the Ethos-U subsystem, it must be able to map a physical address to a bus address. The DTB contains a dma-ranges, which define how to remap physical addresses to the Cortex-M address space.

## 6.2.5.3 Device and Buffer class

The device driver creates a device node at /dev/ethosu<nr> that a user space application can open and issues IOCTL requests to. This is how buffers and networks are created.

Creating a new buffer returns another file descriptor that can be memory mapped for reading and/or writing.



## 6.2.5.4 Network class

Creating a network assumes that the device node has already been opened, and that a buffer has been allocated and populated with the network model.

A new network is created by issuing an IOCTL command on the device node file descriptor. A file descriptor to a buffer, containing the network model, is passed in the IOCTL data. The network class increases the reference count on the buffer, preventing the buffer from being freed before the network object has been destructed.



## 6.2.5.5 Inference class

Creating an inference assumes that a network has already been created, IFM buffers have been allocated and populated with data, and OFM buffers have been allocated.

A new inference is created by issuing an IOCTL command to the network file descriptor. An array of IFM and OFM buffers are passed in the IOCTL data, which reference counts will be increased.

As the inference object has been created an inference request message is sent to the Cortex-M application. The inference request message is written to a ring buffer in shared memory, cache maintenance is executed if necessary, and an IRQ is raised using the Linux mailbox APIs.

On success, a valid file handle is returned to user space. The file handle is used to read the results when the inference completes. Note this is a blocking call.

Once the inference task has finished on the Ethos-U subsystem, the message process writes an inference response message into the response queue in shared memory, executes cache maintenance if needed, and raises an IRQ.

On the Linux side the IRQ is handled and cleared. The IRQ bottom handler is a separate Linux kernel thread responsible for reading the message queue. When the inference response message is received it updates the status of the inference and unblocks any waiting user space processes.

i.MX Machine Learning User's Guide



Figure 14. Inference class

# 6.2.5.6 How to use the inference API

The following steps show how to run a Vela model from Cortex-A:

1. Create the inference device.

device = Device("/dev/ethosu0")

2. Load the model into a buffer from the Vela model file.

shared\_ptr<Buffer> model\_buf = allocAndFill(device, vela\_model);

3. Create the Network instance with the model buffer.

shared\_ptr<Network> network = make\_shared<Network>(device, model\_buf);

4. Load the input feature map (IFM) from the input file (such as a picture for image classification app) into a buffer. If there are multiple inputs, create the buffers one by one and push back to a vector.

```
vector<shared_ptr<Buffer>> ifm;
ifm_size = network->getIfmDims()[0];
ifm_buf = make_shared<Buffer>(device, ifm_size);
memcpy(ifm_buf ->data(), input_data, input_size);
ifm.push_back(ifm_buf)
```

5. Create the output feature map (OFM) buffers according to the output dimensions in the model. If there are multiple outputs, create the buffer one by one and push back to a vector.

```
vector<shared_ptr<Buffer>> ofm;
ofm size = network->getOfmDims()[0]
```

## i.MX Machine Learning User's Guide

```
ofm_buf = make_shared<Buffer>(device, ofm_size);
ofm.push back(ofm buf);
```

6. Create an inference instance with the Network buffer, IFM buffer, and OFM buffer.

```
inf = make_shared<Inference>(net, ifm.begin(), ifm.end(), ofm.begin(),
    ofm.end());
```

7. Call Inference->invoke() to trigger and wait for the completion of the inference task.

Inf->invoke()

8. Access the OFM buffers to get the inference result.

```
Outputs = inf->getOfmBuffers()
```

## 6.2.5.7 Interpreter class

In addition to low-level APIs described above, the Ethos-U driver also provides the Interpreter class. The Interpreter handles the steps mentioned above (device, network, and buffer initialization) internally with class Interpreter.

#### Constructor:

```
Interpreter(const char *model,
const char *device = "/dev/ethosu0",
int64_t arenaSizeOfMB = 16);
model: vela model file
device: ethos-u device name, default: "/dev/ethosu0"
arenaSizeOfMB : shared DDR memory size between Cortex-A and Cortex-M, default:
16(MB)
```

Inference blocking API:

```
void Invoke(int64_t timeoutNanos = 6000000000);
timeoutNanos: timeout for the inference, default value is 60s.
```

Input/Output tensor buffer address helper:

```
template <typename T>
T* typed_input_buffer(int index) {
    int32_t offset = network->getInputDataOffset(index);
    return (T*)(arenaBuffer->data() + offset);
}
template <typename T>
T* typed_output_buffer(int index) {
    int32_t offset = network->getOutputDataOffset(index);
    return (T*)(arenaBuffer->data() + offset);
}
```

Given the tensor index in a model, returns the tensor address and type information.

Input/Output information query:

```
std::vector<TensorInfo> GetInputInfo();
std::vector<TensorInfo> GetOutputInfo();
```

Rev. LF6.12.3\_1.0.0 — 31 March 2025

i.MX Machine Learning User's Guide

These two provides the interface to query inputs and outputs information from a model, including shape and type information (int8/uint8/f32...).

# 6.2.5.8 Interpreter Python wrapper

In addition to C++ API, the Ethos-U Driver also provides the Python API.

It is installed into Yocto rootfs: /usr/lib/python3.10/site-packages/ethosu.

# Example usage:

```
import ethosu.interpreter as ethosu
# loading the vela model file into interpreter
interpreter = ethosu.Interpreter(args.vela model file)
# get the input and output dimensions
inputs = interpreter.get input details()
outputs = interpreter.get output details()
# resize the input according to the model input dimensions
w, h = inputs[0]['shape'][1], inputs[0]['shape'][2]
img = Image.open(args.image).resize((w, h))
data = np.expand dims(img, axis=0)
# associcate the input data with interpreter
interpreter.set input(0, data)
# invoke the inference, this is a blocking API, timeout is 60s
interpreter.invoke()
# get back the inference results, different models have different results.
# Check the model output dimensions and get all the outputs with index.
output_data = interpreter.get output(0)
```

# 6.2.6 Inference with TensorFlow Lite

# 6.2.6.1 Ethos-U Delegate

See Section 2.2.4.

# 6.2.6.2 Delivery package

The Ethos-U support is built into shared library: /usr/lib/libtensorflow-lite.so. When the user loads the Vela model with TFLite API, the engine calls the Ethos-U Linux driver and dispatches the customized Ethos-U operator to Ethos-U firmware on Cortex-M.

The Ethos-U Delegate shared library: /usr/lib/libethsou\_delegate.so.

# 6.2.6.3 Running image classification example

See <u>Section 6.2.3</u> to try the example.

See <u>TensorFlow Lite</u> for how to build and use the Tensorflow Lite API with an application.

## 6.2.6.4 Hardware accelerators warmup time

For TensorFlow Lite, the initial execution of model inference takes longer time, because of the model graph initialization needed by the NPU hardware accelerator. The initialization phase is known as warmup. In this phase, the delegate calls the Vela tool to compile the TensorFlow Lite models.

• This time duration can be decreased for subsequent application that runs by storing on disk the information resulted from the initial Vela processing. The Ethos-U delegate option "cache\_file\_path" should be used for this purpose. For example, set this option in your application in this way:

```
# the external delegate accepts the option "cache_file_path",
ext_delegate = [tflite.load_delegate("/usr/lib/libethosu_delegate.so",
    {"cache_file_path":"your_path"})]
interpreter = tflite.Interpreter(model_path=model_file,
    experimental_delegates=ext_delegate)
```

By setting up this option, the result of the Vela compilation is stored on disk. The runtime performs a quick check on the network. If it matches, it loads it into the NPU memory directly.

• This time duration can also be decreased by compiling the model by Vela beforehand. In this way, you should compile the model with the Vela tool and pass the Vela optimized model file to the TensorFlow Lite application.

## 6.2.6.5 Ethos-U performance enhancement with memory zero-copy

In addition to the memory space of TensorFlow Lite, Ethos-U also has its own memory. Therefore, there is a large number of memory copies in the inference process, which takes a lot of CPU time. To improve the performance, in Ethos-U delegate, some of TensorFlow Lite memory is mapped into the memory of the NPU to avoid memory copies.

**Note:** When doing inference with the Ethos-U delegate, do not modify the TensorFlow Lite Tensor data pointers. This may cause unpredictable issues.

## 6.2.7 Building and deploying the Ethos-U firmware

## 6.2.7.1 Getting the source

The ethos-u-core-software is part of the i.MX 93 Ethos-U NPU machine learning software package, which is an optional middleware component of MCUXpresso SDK. The ethos-u-core-software is integrated into the MCUXpresso SDK Builder delivery system available on mcuxpresso.nxp.com. To include Ethos-U NPU machine learning into the MCUXpresso SDK package, the ethos-u-core-software middleware component is selected in the software component selector on the SDK Builder page when building a new package. See the following figure.

## i.MX Machine Learning User's Guide

SDK Dashboard BUILD SDK Select Board / Processor Middleware (0) Examples (0)	Host OS III III IIII IIIIIIIIIIIIIIIIIIIIII					
Processor Parametrics	Search	h		Q SELECT ALL	UNSELECT ALL	
(OT) ADMINISTRATION		Name	Category	Description	Dependencies	
Notifications		CMSIS DSP Library	CMSIS DSP Lib	CMSIS DSP Software Library		
Preferences  DOWNLOADS		elQ 🖸	Middleware	elQ machine learning SDK containing: - ARM CMSIS-NN library (more)	CMSIS DSP Library	~
MCUXpresso IDE		ethos-u-core-software	Middleware	ethos-u-core-software Controlled Access	CMSIS DSP Library	
Config Tools		lwIP	Middleware	Lightweight IP open-source TCP/IP stack		
MCUXpresso		multicore	Middleware	Multicore Software Development Kit		
Secure Provisioning Tool		FreeRTOS		Real-time operating system for microcontrollers from Amazon		
1 Deployed Releases				DOWNLOAD SDK		

Figure 15. SDK Builder page

Once the MCUXpresso SDK package is downloaded, it can be extracted on a local machine or imported into the MCUXpresso IDE. For more information on the MCUXpresso SDK folder structure, see the *Getting Started with MCUXpresso SDK User's Guide* (document: MCUSDKGSUG). The package directory structure is similar as follows.

```
<MCUXpresso-SDK-root>
|-- boards
   -- <board>
      -- demo apps - Example build projects
          -- ethosu_apps_rpmsg - Ethos-U default firmware with rpmsg
          -- ethosu apps - Ethos-U standalone app example
|-- middleware/ethos-u-core-software
   -- applications - The inference process APIs
   -- boards - The board related initialization and configuration files
   -- core driver - Ethos-U core driver which includes reading/writing
registers
                 - Ethos-U example applications
   -- examples
       -- ethosu_apps_rpmsg - Ethos-U default firmware with rpmsg
       -- ethosu apps
                      - Ethos-U standalone app example
```

# 6.2.7.2 Ethos-U example applications

## 6.2.7.2.1 Introduction

There are two Ethos-U apps available:

- ethosu apps rpmsg: firmware for Yocto Linux BSP
- ethosu apps: standalone example for Cortex-M

The ethosu\_apps\_rpmsg is the firmware for Ethos-U subsystem for Linux OS. It contains core message handling, inference request processing from Cortex-A core, NPU's registers configuration, inference execution, and inference result providing to Cortex-A core. The supported inference engine is TFLite or TFlite-Micro (if Inference API is used).

The example <code>ethosu\_apps</code> is a Cortex-M standalone application that demonstrates the inference execution entirely on Cortex-M which can be used in the low power scenario with the Cortex-A sleeping. The example uses conv2d op model. There is no core message handling and only supports TFLite-Micro.

The apps are available in the /boards/<board>/demo\_apps/ethosu\_apps\* folders.

## 6.2.7.2.2 Toolchains supported

• IAR Embedded Workbench for Arm When the project is opened in IAR, press the "Make" button to build the project in IAR as follows.



Figure 16. IAR Embedded Workbench for Arm

#### ArmGCC - GNU Tools Arm Embedded

Run the following command to build the project.

```
$ cd mcu-sdk-2.16/boards/mcimx93evk/demo_apps/ethosu_apps_rpmsg/armgcc
$ export ARMGCC_DIR=${YOUR_TOOLCHAIN_LOC}/gcc-arm-none-eabi-10-2020-q4-major
$ export PATH=$PATH:${YOUR_TOOLCHAIN_LOC}/gcc-arm-none-eabi-10-2020-q4-major/
bin
$ ./build_release.sh
```

## 6.2.7.3 Deploy procedure

1. Deploy the ethosu\_apps\_rpmsg firmware.

Example ethosu\_apps\_rpmsg is built as .out or .elf and installed in rootfs as the name of "ethosu\_firmware". The pre-built binary is integrated in the rootfs and loaded by Linux Ethos-U driver upon an inference request.

If the user rebuilds the firmware, the rebuilt ethosu\_apps\_rpmsg.out or ethosu\_apps\_rpmsg.elf
should be copied to /lib/firmware/ in rootfs and renamed as the name of "ethosu\_firmware" as
follows:

```
$ cp ethosu apps rpmsg.elf ./lib/firmware/ethosu firmware
```

2. Deploy the ethosu\_apps with U-Boot.

The ethosu\_apps is built as .bin. In U-Boot terminal, users can run the following command to do inference for the conv2d op model.

```
=> tftp 0x80000000 ethosu_apps.bin
=> cp.b 0x80000000 0x201e0000 0x20000;
=> bootaux 0x1ffe0000 0
```

When the example runs, the log and inference result would be displayed on the Cortex-M terminal as follows:

```
Initialize Arm Ethos-U
Inference status: success
```

## Note:

The default firmware ethosu\_apps\_rpmsg contains the following operators support with TFLite-micro on Cortex-M33: Ethos-U, TFLite\_Detection\_PostProcess, and Dequantize. If an operator is supposed to fall back on Cortex-M33 but not included, rebuild the source code and deploy the firmware.

The *ethosu\_apps* is a standalone Cortex-M application running without Cortex-A interacted, so it is deployed at the U-Boot stage.

## 6.2.7.4 Using the Ethos-U on Cortex-M

The Ethos-U NPU on i.MX 93 is accessible by the TFLite-Micro library. The TFLite-Micro interprets the optimized Vela model and delegates the kernels to different execution providers.

Currently, there are 3 types of execution provider supported:

- NN Kernel: default kernel implementation provided by TFLite-Micro for Cortex-M CPU.
- **CMSIS-NN kernel**: optimized kernel implementation by Arm using the CMSIS-NN library. The CMSIS-NN library executes the kernel on Cortex-M CPU or <TBD>.
- Ethos-U Kernel: kernel implementation for the custom Ethos-U operator. This operator registered in TFLite-Micro framework and executes the computation on Ethos-U using the NPU driver.

## 6.2.7.4.1 Running Vela model with TFLite-Micro

The following provides the steps to run the Vela model on Cortex-M directly:

1. Get the flatbuffer Vela model.

const tflite::Model\* model = tflite::GetModel(vela model);

2. Configure/Allocate the inputs, outputs tensors statically.

```
constexpr int kTensorArenaSize = 1024 * 1024;
static uint8 t tensorArena[kTensorArenaSize];
```

3. Build the TFLite-Micro interpreter for the inference.

```
static tflite::MicroInterpreter interpreter(
model, //the flatbuffer model
microOpResolver, //resolve to kernel implementers
tensorArena, // tensor memory address
kTensorArenaSize, //tensor memory length
microErrorReporter); //error reporter
```

4. Set the input tensors.

// Get access to the input tensor data

```
TfLiteTensor* inputTensor = interpreter->input(0);
```

// Copy the input tensor data from an application buffer

```
for (int i = 0; i < inputTensor->bytes; i++)
inputTensor->data.int8[i] = input data[i];
```

5. Run the inference and get the output.

// Invoke the inference

```
interpreter->Invoke();
```

// Get access to the output tensor data

TfLiteTensor\* outputTensor = interpreter->output(0);

// Copy the output tensor data to an application buffer

```
for (int i = 0; i < outputTensor->bytes / sizeof(float32); i++)
    output_data[i] = outputTensor->data.f[i];
```

TFLite-Micro does not depend on dynamic memory allocation, so it requires users (application developers) to supply a memory arena when an interpreter is created. In practice, the user usually allocates this memory arena as a static buffer when the program starts, for example:

```
#define TENSOR_ARENA_SIZE (1024 * 1024 * 16)
uint8 t tensorArena[TENSOR_ARENA_SIZE];
```

TFLite-Micro framework uses this memory arena as inputs/outputs/intermediate tensors store. This memory size "TENSOR\_ARENA\_SIZE" must be adjusted according to the practical usage to consider the following points:

- Model used for the application
- · Size of the input/output data
- Memory needed for intermediate result
- Memory arena mapping to SRAM or TCM, considering the effective usage of memory hierarchy

## 6.2.8 Memory hierarchy for Cortex-M

For Cortex-M, there are several types of memory media with different capacity, speed and cost which can be accessed by CPU. On i.MX 93, the memory hierarchy looks like below with speed decreasing order:

TCM (128 kB + 128 kB)

OCRAM (256 kB TF - A + 384 kB NPU data)

DRAM (default 16 MB, dynamically allocated from dma pool)

aaa-053541

TCM size is 256 KB, usually used for Cortex-M runtime data. By design, this memory space is not allocated for system purpose after booting. How to use it effectively is left for user decision.

OCRAM size is 640 KB. By design, the first 256 KB is allocated for ATF (Arm Trusted Firmware) which used to bootstrap the Cortex-A before the DRAM is available. The rear 384 KB is reserved for NPU data: the weight/ bias of an ML model.

DRAM size is 2 GB on i.MX 93 EVK board. However, only shared DMA region between Cortex-A and Cortex-M can be used. Ethos-U Linux driver requests DMA buffers for tensorArena dynamically from DMA pool and passes the buffer address to Ethos-U firmware on Cortex-M. If not explicitly specified, by default 16 MB DMA buffer is requested.

Ethos-U can only access the DRAM and OCRAM memory by design. The current memory mapping for Ethos-U firmware is as follows:

# **NXP Semiconductors**

i.MX Machine Learning User's Guide



With this configuration, the model data and tensor arena is allocated in DRAM and the OCRAM is used as NPU cache. "Dedicated\_Sram" memory mode has to be used for model compilation with Vela:

```
vela --accelerator-config ethos-u65-256 --system-config Ethos_U65_High_End
--memory-mode Dedicated_Sram --config vela.ini {tflite-model}
```

For standalone Cortex-M app, the memory mapping is as follows:



With this configuration, No DRAM is used. All the model data and tensorArena memory for NPU is allocated in OCRAM. "Sram\_Only "memory mode has to be used for model compilation with Vela:

```
vela --accelerator-config ethos-u65-256 --system-config Ethos_U65_High_End
--memory-mode Sram Only --config vela.ini {tflite-model}
```

## 6.2.9 Supported ML operators and constraints

See the <u>supported operator list</u> on the Ethos-U NPU. When the operator is not supported by the NPU, the Vela compiler displays the constraints information on the console. They are left untouched and scheduled on the CPU. Use the eIQ toolkit to view what operators are merged into Ethos-U operator for a model.

## 6.2.10 Profiling on hardware accelerators

This section describes how to enable profiler on the NPU, and how to capture logs.

Two environment variables are prepared for PMU profiling, ETHOSU\_ENABLE\_CYCLE\_COUNTER and ETHOSU PMU CONFIG.

Set ETHOSU\_ENABLE\_CYCLE\_COUNTER to 1 to enable cycle counter, 0 to disable cycle counter. The default value is 0.

```
$ export ETHOSU ENABLE CYCLE COUNTER="1"
```

Set ETHOSU\_PMU\_CONFIG to enable PMU counter. Up to 4 PMU event IDs can be added in this variable, e.g.,

```
$ export ETHOSU PMU CONFIG="3 4 5 6"
```

or

\$ export ETHOSU\_PMU\_CONFIG="3 4"

## The following table shows all the event IDs supported by Ethos-U.

#### Table 3. Event IDs supported by Ethos-U

Event type	Event ID			
CYCLE	1			
NPU_IDLE	2			
CC_STALLED_ON_BLOCKDEP	3			
CC_STALLED_ON_SHRAM_RECONFIG	4			
NPU_ACTIVE	5			
MAC_ACTIVE	6			
MAC_ACTIVE_8BIT	7			
MAC_ACTIVE_16BIT	8			
MAC_DPU_ACTIVE	9			
MAC_STALLED_BY_WD_ACC	10			
MAC_STALLED_BY_WD	11			
MAC_STALLED_BY_ACC	12			
MAC_STALLED_BY_IB	13			
MAC_ACTIVE_32BIT	14			
MAC_STALLED_BY_INT_W	15			
MAC_STALLED_BY_INT_ACC	16			
AO_ACTIVE	17			
AO_ACTIVE_8BIT	18			
AO_ACTIVE_16BIT	19			
AO_STALLED_BY_OFMP_OB	20			
AO_STALLED_BY_OFMP	21			
AO_STALLED_BY_OB	22			
AO_STALLED_BY_ACC_IB	23			
AO_STALLED_BY_ACC	24			
AO_STALLED_BY_IB	25			
WD_ACTIVE	26			
WD_STALLED	27			
WD_STALLED_BY_WS	28			
WD_STALLED_BY_WD_BUF	29			
WD_PARSE_ACTIVE	30			
WD_PARSE_STALLED	31			
WD_PARSE_STALLED_IN	32			

UG10166 User guide

i.MX Machine Learning User's Guide

Event type	Event ID				
WD_PARSE_STALLED_OUT	33				
WD_TRANS_WS	34				
WD_TRANS_WB	35				
WD_TRANS_DW0	36				
WD_TRANS_DW1	37				
AXI0_RD_TRANS_ACCEPTED	38				
AXI0_RD_TRANS_COMPLETED	39				
AXI0_RD_DATA_BEAT_RECEIVED	40				
AXI0_RD_TRAN_REQ_STALLED	41				
AXI0_WR_TRANS_ACCEPTED	42				
AXI0_WR_TRANS_COMPLETED_M	43				
AXI0_WR_TRANS_COMPLETED_S	44				
AXIO_WR_DATA_BEAT_WRITTEN	45				
AXI0_WR_TRAN_REQ_STALLED	46				
AXI0_WR_DATA_BEAT_STALLED	47				
AXI0_ENABLED_CYCLES	48				
AXI0_RD_STALL_LIMIT	49				
AXI0_WR_STALL_LIMIT	50				
AXI_LATENCY_ANY	51				
AXI_LATENCY_32	52				
AXI_LATENCY_64	53				
AXI_LATENCY_128	54				
AXI_LATENCY_256	55				
AXI_LATENCY_512	56				
AXI_LATENCY_1024	57				
ECC_DMA	58				
ECC_SB0	59				
AXI1_RD_TRANS_ACCEPTED	60				
AXI1_RD_TRANS_COMPLETED	61				
AXI1_RD_DATA_BEAT_RECEIVED	62				
AXI1_RD_TRAN_REQ_STALLED	63				
AXI1_WR_TRANS_ACCEPTED	64				
AXI1_WR_TRANS_COMPLETED_M	65				
AXI1_WR_TRANS_COMPLETED_S	66				
AXI1_WR_DATA_BEAT_WRITTEN	67				
AXI1_WR_TRAN_REQ_STALLED	68				

Table 3. Event IDs supported by Ethos-U...continued

UG10166 User guide

AXI1_WR_DATA_BEAT_STALLED	69					
AXI1_ENABLED_CYCLES	70					
AXI1_RD_STALL_LIMIT	71					
AXI1_WR_STALL_LIMIT	72					
ECC_SB1	73					

After setting one or two environment variables, you can run the TensorFlow Lite application. The PMU counter result is displayed on the console.

```
$ export ETHOSU ENABLE CYCLE COUNTER="1"
$ export ETHOSU PMU CONFIG="1 3 4 5"
$ ./label image -m mobilenet v2 1.0 224 quant vela.tflite --
external delegate path=/usr/lib/libethosu delegate.so -1 labels.txt -i
stopwatch.bmp
Ethos u PMUs : [ 201717971 34712 41893 3457751 ]
Ethos-u cycle counter: 237501410
```

## 6.3 NPU transition guide from i.MX 8M Plus to i.MX 93

This section describes how to port Machine Learning application from i.MX 8M Plus to i.MX 93 with NPU acceleration.

## 6.3.1 Tensorflow Lite difference between i.MX 8M Plus and i.MX 93 NPU acceleration

See Figure 3 for Tensorflow Lite software stack. Both i.MX 8M Plus and i.MX 93 support Tensorflow Lite with NPU acceleration. i.MX 93 also supports TensorFlow Lite external delegate mechanism.

From the development perspective of the Machine Learning application, users can use the same Tensorflow API to develop the Machine Learning application. The only difference is that users need to use the Ethos-U Delegate instead of VX Delegate.

## 6.3.2 NPU supported operator list

While porting the Machine Learning application from i.MX 8M Plus to i.MX 93, check whether the NPU supported operators in your model are supported on the i.MX 93 NPU. This ensures that you leverage i.MX 93 NPU acceleration.

See supported operator list for i.MX 93 NPU operator support status and Table 11 for i.MX 8M Plus NPU operator support status.

## 6.4 Hardware acceleration with eIQ Neutron NPU on i.MX 9 series platform

eIQ Neutron NPU is a Neural Processing Unit (NPU) developed by NXP. It is designed to accelerate the Machine Learning inference. The Neutron-S version of the eIQ Neutron NPU comprises of 3 main blocks:

- Neutron computation core (Neutron) doing MACs, can be pipelined with multiple instances.
- Neutron controller (RISC-V core) to program Neutron registers and control the Neutron block.
- DMA like memory controller (Data Mover) to exchange data between host DDR and Neutron dedicated TCM.

Neutron-S NPU main features:

- Targets quantized Convolutional Neutral Networks (CNN) and supports 8 bit weights and 8/16 bit activations.
- Supports TensorFlow Lite (TFLite) inference with fallback to Cortex-A for unsupported operations.
- Supports TFLite API to offload a custom TFLite node neutronGraph, to Neutron-S NPU.
- Provides model converter tool (through elQ toolkit) to optimize the model performance and memory usage for Neutron-S NPU target.

## 6.4.1 Neutron-S NPU overview

The Neutron-S NPU involves several hardware blocks working together to support the acceleration of the tensor computation defined by the Machine Learning model:

- SoC main CPU (Cortex-A55)
- RISC-V Controller
- Data Mover
- Neutron compute block

The SoC main CPU runs the software under Linux OS, like the TFLite inference engine. It is responsible for loading the Machine Learning model, capturing and pre-processing the inputs and handing over the tensor computation to the NPU. The RISC-V controller in the Neutron-S NPU orchestrates the Neutron compute blocks and Data Mover. The Data Mover is a DMA like engine used for moving data between the SoC DDR and the NPU TCM.

## 6.4.2 Neutron-S software architecture

The software for Neutron-S NPU includes three main components, as shown in the following figure.



 The Neutron model converter is an offline tool to compile the TFLite model for Neutron-S. The converter replaces supported operators in the model with a custom neutronOp node containing a model-specific firmware binary, static data, like weights, and inputs/outputs memory areas for Neutron-S. The output of the converter is a modified TFLite model graph for TFLite inference engine. For inference, it needs to use the corresponding TFLite Neutron Delegate.

• The Cortex-A software stack for Linux contains the TFLite inference engine, Neutron delegate library, user space Neutron driver library, and Neutron device driver for the Linux kernel.

The Neutron-FW stack contains code for the RISC-V controller, interpreting the microcode in the neutronGraph node.

## 6.4.3 NPU performance tuning

NXP Yocto Linux enables power-saving technology by default. To boost Machine Learning Performance on the i.MX 95 EVK board, perform the following tuning operations:

1. Disable DDR clock gating.

DDR clock gating is the hardware feature applied to the i.MX 9 family. It is enabled by default in Yocto Linux for maximum power saving with performance trade-off.

On i.MX 95, send the following command on the System Manager console (the serial port right after the Linux console) to disable DDR clock gating.

>\$ mm 0x4e010010 0

2. Set the CPU to performance mode.

In Yocto Linux, the CPU is in ondemand mode by default, which makes the CPU work at low frequency when there is low workload. Make the CPU in performance mode before eIQ benchmarking.

```
# cpufreq-set -g performance
# cat /sys/devices/system/cpu/*/cpufreq/cpuinfo_cur_freq
1800000
```

3. Disable CPU idle in U-Boot.

For the Linux system, when disabling CPU Idle management, it can reduce context switch overhead.

```
# setenv cpuidle 'cpuidle.off=1'
```

4. Enable the Zero-copy feature.

To reduce memory transaction between the CPU and NPU, i.MX 95 Machine Learning software stack implements zero-copy feature by setting the following environment in Yocto Linux before running the inference task.

```
# export NEUTRON ENABLE ZERO COPY=1
```

# 7 Vision Pipeline with NNStreamer

<u>NNStreamer</u> is an efficient and flexible stream pipeline framework for complex neural network applications. It was initially developed by Samsung and then transferred to LF AI Foundation as an incubation project.

It is a set of <u>GStreamer plugins</u> that allows both GStreamer developers to adopt neural network models and neural network developers to manage neural network pipelines and their filters easily and efficiently.

The project is well documented through its dedicated <u>github documentation site</u>, but the main takeaways are described below for convenience.

In addition to the standard GStreamer data types, NNStreamer adds new data types "other/tensor" and "other/ tensors" using a dedicated converter element. This data type represents a stream of multidimensional array and a stream of a container of multiple instances of such arrays, respectively.

NNStreamer provides a set of stream filters applying multiple operations on tensors:

• tensor\_converter converts audio, video, text, or arbitrary binary streams to others/tensor streams.

- tensor decoder converts other/tensor(s) to video or text stream with assigned sub-plugins.
- tensor\_filter invokes a neural network model with the given model path and neural network framework name.
- tensor\_transform applies various operators to tensors including typecast, add, mul, transpose, and normalize. For faster processing, it supports SIMD instructions and multiple operators in a single filter.
- tensor\_crop crops the regions of incoming tensor.
- tensor rate controls a frame rate of tensor streams.
- tensor\_mux, tensor\_demux, tensor\_merge, tensor\_split, tensor\_if, and tensor\_aggregator support tensor stream path controls.
- tensor sink is a sink plug-in for making an application to get a buffer of other/tensor(s).
- tensor\_source allow non GStreamer standard input sources, such as sensors, to supply other/ tensor(s) stream.
- tensor\_reposink and tensor\_reposrc implement recurrent path helpers, cutting GStreamer pipeline cycle through a dedicated shared repository. The tensor\_reposink pushes data to the repository, this latter reinjecting data upstream through a tensor reposrc element.

The following figure shows the general architecture of a NNStreamer pipeline.



#### Figure 18. NNStreamer pipeline

There are two elements allowing adding user created features in run-time: <u>tensor\_filter</u> and <u>tensor\_decoder</u>.



Figure 19. NNStreamer filter and decoder flow

While instantiating the *tensor\_filter* and *tensor\_decoder*, the framework and mode options respectively specify the target implementation through a dedicated shared library loaded at runtime. NNStreamer supplies a set of filters and decoders which are described briefly below, and APIs to implement customized user sub-plugins. Hence, it is possible to use a proprietary inference engine sub-plugin as tensor filter, or a specialized NN decoder.

NNStreamer supports the most popular inference engines (open source or not). On this release, TensorFlow Lite and TVM engines are supported.

Framework/Tool	i.MX 95	i.MX 93	i.MX 8M Plus	i.MX 8M Quad/8M Nano/8QuadMax/8Qu	i.MX 8M Mini/8ULP
TensorFlow Lite	CPU/GPU/NPU	CPU/NPU	CPU/NPU/GPU	CPU/GPU	CPU
TVM	-	-	CPU/NPU/GPU	-	-
Custom C++	CPU	CPU	CPU	CPU	CPU
Custom Python	CPU	CPU	CPU	CPU	CPU

#### Table 4. NNStreamer supported features

Framework/Tool	i.MX 95	i.MX 93	i.MX 8M Plus	i.MX 8M Quad/8M Nano/8QuadMax/8Qu	i.MX 8M Mini/8ULP
NNShark	-	-	CPU	-	-

#### Table 4. NNStreamer supported features...continued

In case an inference engine might be supported on multiple hardware backend, one can specify the device mapping the neural network.

Even though Tensor decoder element might not be appropriate for building an application which usually does not consume the neural network outputs for display purpose only, it is especially useful for implementing a prototype during the development phase which might focus on the neural network model or optimizing the data path. Indeed, most neural networks topologies are supported for classical computer vision use cases: classification, object detection, pose estimation or segmentation.

NNStreamer tensor filter element has to be configured to use specific engine and hardware accelerator. Available options are listed in the following tables.

Delegate	Tensor filter properties	USE_GPU_INFERENCE env variable
No delegate	<pre>framework=tensorflow-lite model=<path .tflite="" file="" model="" to=""> custom=NumThreads:<cpu cores=""> Note: <cpu core=""> values: • 2 for i.MX 93 and i.MX 8ULP • 4 for others • 6 for i.MX 95</cpu></cpu></path></pre>	-
XNNPACK Delegate	<pre>framework=tensorflow-lite model=<path .tflite="" file="" model="" to=""> custom=Delegate:XNNPACK,NumThreads:<cpu cores=""> Note: <cpu core=""> values: • 2 for i.MX 93 and i.MX 8ULP • 4 for others • 6 for i.MX 95</cpu></cpu></path></pre>	-
Neutron Delegate (i.MX 95 only)	<pre>framework=tensorflow-lite model=<path .tflite="" file="" model="" to=""> custom=Delegate:External,ExtDelegateLib:</path></pre>	- libneutron_delegate.so
VX Delegate (applicable for supported i.MX 8)	<pre>framework=tensorflow-lite model=<path .tflite="" file="" model="" to=""> custom=Delegate:External,ExtDelegateLib:</path></pre>	0:NPU 1:GPU libvx_delegate.so
Ethos-U Delegate (i.MX 93 only)	<pre>framework=tensorflow-lite model=<path .tflite="" file="" model="" to=""> custom=Delegate:External,ExtDelegate Lib:libethosu_delegate.so</path></pre>	-

Table 5. TensorFlow Lite engine

#### Table 6. TVM engine

Tensor filter properties	USE_GPU_INFERENCE env variable
framework=tvm model= <path .so="" model<="" td="" to=""><td>0: NPU</td></path>	0: NPU
<pre>library&gt; custom=num_input_tensors:<number of<="" pre=""></number></pre>	1: GPU
input tensors>	Relevant for models compiled to use OpenVX
where <number input="" of="" tensors=""> is typically 1.</number>	,

# 7.1 Object detection pipeline example

This section provides implementation details for an object detection pipeline running on i.MX 8M Plus. Additional pipeline examples targeting more use-cases and i.MX platforms can be found in <u>Section 7.2</u>.

In this example, the following pipeline will be implemented leveraging all the compute backend available on i.MX 8M Plus to build an object detection scenario.



On the target, download the trained neural network from google coral github site, and export the filenames to bash environment variables:

```
root:~# wget https://github.com/google-coral/test_data/raw/master/
ssd_mobilenet_v2_coco_quant_postprocess.tflite
root:~# wget https://github.com/google-coral/test_data/raw/master/coco_labels.txt
root:~# export MODEL=$(pwd)/ssd_mobilenet_v2_coco_quant_postprocess.tflite
root:~# export LABELS=$(pwd)/coco_labels.txt
```

#### Then builds and executes the GStreamer pipeline:

```
root:~# gst-launch-1.0 --no-position v4l2src device=/dev/video3 ! \
video/x-raw,width=640,height=480,framerate=30/1 ! \
tee name=t t. ! queue max-size-buffers=2 leaky=2 ! \
imxvideoconvert_g2d ! \
video/x-raw,width=300,height=300,format=RGBA ! \
videoconvert ! video/x-raw,format=RGB ! \
tensor_converter ! \
tensor_filter framework=tensorflow-lite model=${MODEL} \
custom=Delegate:External,ExtDelegateLib:libvx_delegate.so ! \
tensor_decoder mode=bounding_boxes option1=mobilenet-ssd-postprocess option2=${LABELS} \
option3=0:1:2:3,50 option4=640:480 option5=300:300 ! \
```

```
mix. t. ! queue max-size-buffers=2 ! \
imxcompositor_g2d name=mix latency=30000000 min-upstream-latency=30000000
sink_0::zorder=2 sink_1::zorder=1 ! waylandsink
```

*Note: Hit* CTRL+C *keystroke to halt the execution if necessary.* 

# 7.2 NXP NNStreamer pipeline examples

Pipelines targeting i.MX platforms are published to provide working examples for different use cases and implementation options.

Those examples are hosted on the GitHub server in a dedicated tree:

https://github.com/nxp-imx/nxp-nnstreamer-examples

Refer to the included README documentation for pipelines descriptions and instructions for dependencies download (models, metadata) and execution.

The following table lists the features covered by pipeline examples.

Table 7.	Features	of NXP	NNStreamer	examples
		•••••		

Category	Engine	Platform	Implementation
Object detection: MobileNet SSD V2, Yolov4- tiny	TensorFlow Lite	i.MX 8M Plus i.MX 93	Shell script (gst-launch) Python, C++
Image classification: MobileNet V1		i.MX 95	Custom Python tensor_filter
Image segmentation: DeepLab V3			
Pose detection: MoveNet			
Face detection: UltraFace			
Face recognition: FaceNet512			
Emotion detection: DeepFace			
Depth estimation: Midas V2			

# 7.3 Pipeline profiling

NNStreamer team developed <u>NNShark</u>, a profiling tool based on <u>GstShark</u>, to monitor several pipeline metrics useful to assess the SoC hardware usage.

NNShark can be used on the i.MX 8M Plus only, where specific metrics were added:

- 2D GPU (GC520L) utilization load
- 3D GPU (GC7000UL) utilization load
- NPU (GC8000) utilization load
- SoC masters bandwidth, as reported by Linux kernel perf tool
- Additionally, power domain consumption, as reported by <u>power measurement tool (PMT)</u> if the <u>power</u> <u>measurement evaluation kit</u> is available to the user.

Considering the complex GPU/NPU architecture involving concurrent stages, their reported utilization loads shall be considered as an order of magnitude and might not precisely reflect each individual stage's status.

## Note:

For the source code demo location see the <u>nnshark</u> repository.

## 7.3.1 Enable profiling with NNShark

It is recommended to connect to the target through SSH as the NNShark UI refresh rate might not render well on the serial console.

Enable NNShark profiling through environment variables:

```
root:~# export GST_DEBUG="GST_TRACER:7"
root:~# export GST_TRACERS="live"
```

To get GPU usage measurements, disable power saving in the GPU driver (galcore) using command line Linux kernel parameters. You can manually edit the bootargs U-Boot variable before executing the boot command. Add the following parameters:

galcore.gpuProfiler=1 galcore.powerManagement=0

Then run the previous gst-launch command line, and the following screen should now be displayed on your terminal screen. You can scroll through all the pipeline elements with up/bottom direction key to select the desired element and display its connections with other pipeline elements.

You can select the element pads with left/right direction keys to highlight its connection to other elements' pads.

On this example, the tensor filter has an average processing time of 21.64 ms and its sink orange highlighted pad is connected to source pad of tensorconverter0 element (green highlighted).

Press 'q' or 'Q' to exit the profiling tool and return to the shell terminal. You can quit the application as previously explained through CTRL+C.

press q or Q	to quit	Key -00	00001	2021-08-10 15:12:					
CPU Usage	GPU Usage	DDR Usage		PWR Measurement					
CPU 0 6.2%	GC7000 0.0%	all-rd	1676.88 MB/s	VDD_ARM	128.30 mW				
CPU 1 26.5%	GC8000 22.4%	all-wr	1298.89 MB/s	NVCC_DRAM_1V1	173.35 mW		>		>
CPU 2 12.0%	GC520 17.6%	npu-rd	937.64 MB/s	VSYS_5V	4049.07 mW		from tensorconverter0		to tensordec0
CPU 3 28.3%		npu-wr	779.56 MB/s	VDD_SOC	1351.81 mW		bufrate: 30.02		bufrate: 30.13
		gpu3d-rd	36.01 MB/s	LPD4_VDDQ	14.62 mW		bufsize: 270000		bufsize: 484
		gpu3d-wr	0.00 MB/s	LPD4_VDD2	215.77 mW				
		gpu2d-rd	253.45 MB/s	LPD4_VDD1	10.10 mW				
		gpu2d-wr	337.19 MB/s						
		a53-rd	237.94 MB/s					Proctime: 21169969	
		a53-wr	162.57 MB/s					Average: 21643163.605127	
		isi1-rd	0.00 MB/s					Queue_level: 0/0	
		isi1-wr	17.72 MB/s						
ElementName		roctime(ns)	Avg_proctime(ns)	queuelev	el Buffer	rate(bps)			
videoconvert0		1108327	1270261,910	0/	0				
sink						30.03			
src						30.02			
tensorfilter0		21169969	21643163.605	0/					
sink						30.02			
src						30.13			
waylandsink0			0.000	0/					
sink						30.01			
tensordec0		507910	503151.770	0/					
sink						30.13			
src						30.13			
capsfilter1		11501	17773.887	0/					
sink						30.03			
src		20502				30.03			
capsfilter2		29502	28117.355	0/	0				
sink						30.02			
SIC		22627	27447 450			30.02			
capsiliterø		22627	2/44/.458	07		20.02			
STIR						30.03			
miv		0	0 000	a)	0	50.05			
sink 1			0.000	07		30 03			
src_1						30.01			
sink 0						30.13			
queue1		26377	41594_301	0/	2				
sink						30.03			
src						30.03			
imxvideoconvert	_g2d0	591291	952140.561	0/					
sink	(82)					30.03			
src						30.03			
queue0		39003	53774.845	0/					
sink						30.03			
src						30.03			
v4l2src0			0.000	0/					
src						30.03			
tensorconverter	0	17626	16524.782	0/					
sink						30.02			
STC						30.02			
t			0.000	0/					
src_0						30.03			
sink						30.03			
src 1						30.03			

Figure 21. NNShark i.MX 8M Plus example screenshot

## 7.3.2 Adding power measurement to NNShark

On the desktop PC connected to the power measurement evaluation kit, execute <u>the power measurement tool</u> (<u>PMT</u>) in server mode such as the power measurements are collected and available on 65432 TCP/IP port.

user@localhost:pmt# python3 main.py server -b imx8mpevkpwra0 -p 65432

On the target, export the desktop PC ip address (192.168.1.99 for this example):

root:~# export GST TRACERS PWR SERVER IP=192.168.1.99

Note: The user can run the NNShark without the power measurement kit.

#### 7.3.3 Known issues and limitations

In case perf reports inconsistent high numbers, this means that a perf process is still running in background of the previous run. If so, you must terminate manually their execution.

For your convenience, the below command can be used:

```
root:~# kill -9 $(ps -ef | grep nnshark-perf-ddr.sh | grep -v grep | tr -s ' ' |
cut -d ' ' -f 2)
```

# 8 elQ Demos

## 8.1 TensorFlow Lite Demos for i.MX 93

This section provides implementation details for several TensorFlow Lite demos running on i.MX 93.

TensorFlow Lite demos (binaries) are located at: /usr/bin/eiq-examples-git.

Binary models are not located in the image because of the size. Before running the demos, these files should be downloaded to the device:

```
$ cd /usr/bin/eiq-examples-git
$ python3 download_models.py
```

**Note:** This script is downloaded from GitHub and Google drive. Make sure the device network is correctly configured and can access the Internet.

#### 8.1.1 Image classification demo

Note: All the demos require X11 to display, so use the XWayland distro images.

This demo performs image classification using a pretrained mobilenet-v1 network. Demo dependencies are from:

/usr/bin/eiq-examples-git/image classification.

- grace\_hopper.bmp
- label image.py
- labels.txt

The demo network model dependencies:

• mobilenet\_v1\_1.0\_224\_quant.tflite

Run the Python example with the image input from the default location:

```
$ cd /usr/bin/eiq-examples-git/image_classification
$ python3 label_image.py -i grace_hopper.bmp -l labels.txt
0.874510: military uniform
0.031373: Windsor tie
0.015686: mortarboard
0.011765: bulletproof vest
0.007843: bow tie
time: 4.126ms
```

## 8.1.2 SSD object detection demo

The SSD object detection demo performs object detection using the Single-Shot multibox Detection (SSD) detector. It detects objects on camera, video, or image. Demo dependencies are from: /usr/bin/eiq-examples-git/object\_detection.

- cars0.bmp
- labels.py
- main.py

The demo network model dependencies:

ssd\_mobilenet\_v1\_quant.tflite

Run the Python example with the image input from the default location:

```
$ cd /usr/bin/eiq-examples-git/object_detection
$ python3 main.py -i cars0.bmp
rectangle:(640,493),(1756,881) label:car
rectangle:(1470,466),(1947,694) label:car
rectangle:(803,462),(846,502) label:car
rectangle:(733,451),(788,493) label:car
rectangle:(573,473),(705,565) label:car
rectangle:(608,465),(679,519) label:car
rectangle:(203,455),(271,596) label:car
rectangle:(910,461),(956,500) label:car
rectangle:(1020,453),(1076,497) label:person
```



Run the Python example with the live camera connected to port 0.

\$ python3 main.py -i /dev/video0

**Note:** Choose the right port where the camera is currently connected. Use the v412-ct1 --list-devices command to check it.

## 8.1.3 Hand gesture detection demo

This application demonstrates hand detection and gesture detection. It detects objects on camera, video, or image. Demo dependencies are from: /usr/bin/eiq-examples-git/gesture detection.

- anchors.csv
- hand0.bmp
- hand tracker.py
- main.py

The demo network model dependencies:

- palm\_detection\_builtin\_256\_integer\_quant.tflite
- hand\_landmark\_3d\_256\_integer\_quant.tflite

Run the Python example with the image input from the default location:

- \$ cd /usr/bin/eiq-examples-git/gesture\_detection
- \$ python3 main.py -i hand0.bmp



Run the Python example with the live camera connected to port 0.

\$ python3 main.py -i /dev/video0

*Note:* Choose the right port where the camera is currently connected. Use the v4l2-ctl --list-devices command to check it.

## 8.1.4 Face recognition demo

This application is a demonstration for real-time face recognition. It uses pretrained yoloface model for face detection, and facenet model to calculate face landmark. The demo supports the live camera input only.

Demo dependencies are from: /usr/bin/eiq-examples-git/face\_recognition.

- face\_database.py
- face\_detection.py
- face\_recognition.py
- main.py

#### The demo network model dependencies:

- yoloface\_int8.tflite
- facenet\_512\_int\_quantized.tflite

Before running the demo, connect a keyboard to the board.

- 1. Run the Python example with the live camera connected to port 0.
  - \$ cd /usr/bin/eiq-examples-git/face\_recognition
  - \$ python3 main.py -i /dev/video0



**Note:** Choose the right port where the camera is currently connected. Use the v412-ct1 --listdevices command to check it.

Add a name to the face database.
 Face the camera and press 'a' on the keyboard, which is connected to the board, and then input a new name.

## i.MX Machine Learning User's Guide



3. Delete the name from the face database. Press 'd' on the keyboard, which is connected to the board, and then input the name.

# 9 Release Notes

# 9.1 Known issues and limitations

- Inline model compilation on i.MX 95 is not available.
- Building i.MX 93 Ethos-U Firmware from the source will be supported in the MCU SDK 25.03.00 release.
- Hardware Accelerators on i.MX 8 does not support layers with dynamic shapes.
- The NPU on i.MX 8M Plus is not optimized for models with dynamic weights. The layers with dynamic weights (for example, in the FullyConnected layer) are computed significantly slower.
- Some of the links for the models in the download\_models.py script from <u>Section 8.1</u> are no longer available.

# 9.2 Release notes for LF6.6.52\_2.2.0

i.MX 95:

- Neutron Software Stack upgraded.
- Not support inline compilation, only the converted TensorFlow Lite model is interpreted by Neutron Runtime.
- Pre-compiled model specific binary in NeutronOp with significant performance improvement for inference.

# 9.3 Release notes for LF6.6.36\_2.1.0

General:

• Arm Compute Library is removed since this release.

TensorFlow Lite:

- Upgraded to 2.16.2.
- Fixed missing PIL module when running the label\_image Python example (label\_image.py)
- Modified Flex delegate build onto a 2-stage process.

**ONNX** Runtime

• Fixed compilation with Yocto SDK.

i.MX 8M Plus:

- VX Delegate update and bug fixes
- TIM-VX update, and internal OVXLIB updated to 1.2.14.

i.MX 93:

- Arm Vela Compiler updated to version 3.12.
- Ethos-U software updated to 24.05.

i.MX 95:

• Neutron Software Stack upgraded.

# 9.4 Release notes for LF6.6.23\_2.0.0

General:

• Added support for i.MX 91.

TensorFlow Lite:

- Upgraded to 2.15.0.
- Added GPU Delegate for i.MX 95. The GPU delegate is available in the C++ API.

## ONNX Runtime

• Upgraded to 1.17.1.

PyTorch

• PyTorch framework not available by default in the BSP. Can be deployed using pip from the PyPI registry.

i.MX 8M Plus:

- VX Delegate update and bug fixes
- TIM-VX update and bug fixes.

i.MX 93:

- Arm Vela Compiler updated to version 3.11.
- Ethos-U software updated to 24.02.

i.MX 95:

- Added support to offload ML workload on the on-chip Arm Mali G310 GPU with TensorFlow Lite, using the GPU Delegate.
- Added support for eIQ Neutron Neural Processing Unit using the inline compilation.

• Full NPU acceleration for mobilenetv1 and mobilenetv2 models or models with similar operators. Expect other popular CNN models in the upcoming releases.

# 9.5 Release notes for LF6.6.3\_1.0.0

General:

• Initial support for the i.MX 95 platform.

## TensorFlow Lite:

- Upgraded to 2.14.0.
- Added helper script to generate reduced-size Flex Delegate Bazel artifacts.

i.MX 8M Plus:

- VX Delegate update and bug fixes
- TIM-VX update and bug fixes.

i.MX 93:

- Arm Vela Compiler updated to version 3.10.
- Ethos-U software updated to 23.11.

i.MX 95:

- Added eIQ for i.MX 95.
- Added support for eIQ Neutron Neural Processing Unit using offline compilation. The compiler is available in the eIQ Toolkit.

## 9.6 Release notes for LF6.1.55\_2.2.0

#### General:

• Model Runner was removed from Linux BSP. The elQ Toolkit deploys the compatible Model Runner instance automatically.

TensorFlow Lite:

• Upgraded to 2.12.1.

**ONNX Runtime:** 

- Upgraded to 1.16.1.
- NNAPI execution provider support was removed.

i.MX 8M Plus:

- VX Delegate update and bug fixes.
- TIM-VX update and bug fixes.

i.MX 93:

- Arm Vela Compiler updated to version 3.9.
- Ethos-U software updated to 23.08.

elQ Demos:

• Removed the support for AWS end-to-end SageMaker demo.

# 9.7 Release notes for LF6.1.36\_2.1.0

TensorFlow Lite

- Upgraded to 2.11.1.
- Bug fixes.
- Added Flex Delegate support, including the binary size reduction described here: <u>www.tensorflow.org/lite/</u> <u>guide/reduce\_binary\_size</u>

## VX Delegate

- Synchronized with TensorFlow 2.11.1.
- Bug fixes.

DeepViewRT

• DeepViewRT inference engine was removed.

# 9.8 Release notes for LF6.1.22\_2.0.0

## VX Delegate

- Bug fixes.
- Added support for EmbeddingLookup, Cast, and BroadcastTo.
- Fixed performance on MobilenetV1, MobileNetV2, VGG16, VGG19, and NasNet Mobile.

## ONNX Runtime

- Upgraded to 1.13.1.
- VSI-NPU Execution provider is obsolete and was removed from ONNX Runtime.
- Added support to run dynamic-shape models using NNAPI Execution Provider.

PyTorch

• Upgraded to 2.0.0.

DeepViewRT

• DeepViewRT inference engine is deprecated and will be removed in the future.

i.MX 93

• Arm Vela Compiler: Updated to version 3.7.

# 9.9 Release notes for LF6.1.1\_1.0.0

## TensorFlow Lite

- Upgraded to 2.10.0.
- Deprecated Ethos-U Custom operator on i.MX 93. The preferred way for models with Ethos-U Operator is using the Ethos-U Delegate.

## VX Delegate

• Bug fixes.

UG10166

User guide

- Added support for UnidirectionalSequenceLSTM, BidirectionalSequenceLSTM, Shape, HashtableLookup operators.
- Updated C++ Standard to C++17.
- Fixed TransposeConv2d operator.

• Known issue: Decreased performance on MobilenetV1, MobileNetV2, VGG16, VGG19, and NasNet Mobile.

i.MX 93

- Arm Vela Compiler: Updated to version 3.6.
- Introduced Ethos-U Delegate for i.MX 93.

elQ Demos

• Added TensorFlow Lite demo application for i.MX 93.

# 9.10 Release notes for LF5.15.71\_2.2.0

TensorFlow Lite

- Added option to inference diff tool to compare the inference to reference model. This enables validation of the model on i.MX 93 accelerated by Ethos-U NPU.
- Ethos-U: Enables getting PMU counters from the NPU.
- Ethos-U: Uses one flash and Arena buffer for multiple Ethos-U Operators.

## VX Delegate

- Bug fixes.
- Fixed failures with TensorFlow Lite kernel tests: expand\_dims, LRN, strided-slice, resize, maximum, minimum, and conv3d.

## i.MX 93

- NPU profiling support.
- Ethos-u-driver-stack: Updated to version 22.08.
- Arm Vela Compiler: Updated to version 3.5.

# 9.11 Release notes for LF5.15.52\_2.1.0

- General
  - Added support for i.MX 93 platform, including NN acceleration on Ethos-U NPU.
- TensorFlow Lite
  - TensorFlow Lite updated from version from 2.8.0. to 2.9.1. For details, see RELEASE.md in the source code repository.
  - Added support for Ethos-U HW acceleration for i.MX 93 platform.
- VX Delegate
  - Added support for ReverseV2, UnidirectionalSequenceLSTM and Unpack operators.
  - Fixed bug in reshape for inception\_v1\_224\_quant model.
  - Fixed Yolo-V4-tiny.
  - Other minor bug fixes.
- TIM-VX
  - **–** TIM-VX updated from 1.1.42 to 1.1.50.
- Arm Compute Library
- Arm Compute Library updated from 21.08 to 22.05.
- DeepViewRT
  - DeepViewRT updated from 2.4.42. to 2.4.46.
- elQ Examples
  - Resolved dependency issue due to Yocto BSP upgrade: AWS end-to-end SageMaker demo can be built with latest Yocto BSP (LF5.15.52\_2.1.0).

# 9.12 Release notes for LF5.15.32\_2.0.0

- ArmNN inference engine was removed from eIQ.
- TensorFlow Lite
  - TensorFlow Lite was updated from version 2.6.0 to 2.8.0. For details, see RELEASE.md in the source code repository.
  - Features and improvements:
    - Fixed evaluation tools build with Yocto SDK. Prior to build of the evaluation tools with CMake, it is necessary to build and install the required tooling (protobul compiler protoc). Use the CMakeLists.txt from tensorflow/lite/tools/cmake/native tools/.
- ONNX Runtime
  - Features and improvements:
    - ArmNN and ACL Execution providers were removed from eIQ.
    - VSI\_NPU backend is deprecated and will be removed in the future.
    - NNAPI execution provider is experimental feature.
- TIM-VX
- TIM-VX was updated from 1.1.37 to 1.1.42.
- DeepViewRT
  - DeepViewRT was updated from 2.4.37 to 2.4.42.

## 9.13 Release notes for LF5.15.5-1.0.0

- Arm NN inference engine is deprecated in this release and will be removed in the future.
- NNAPI Delegate of TensorFlow Lite and NNAPI Execution Provider of ONNX Runtime is deprecated and will be removed in the future. For leveraging ML model acceleration use VX Delegate instead.
- TensorFlow Lite:
  - Features and improvements:
    - Fixed unit test build with TensorFlow Lite static library.
    - Support FullyConnected layer with implicit bias in VX Delegate.
    - Fix bug in stride slice if end dim set as -1 in VX Delegate.
    - Other minor fixes.
- ONNX Runtime:
  - Features and improvements:
    - Version update from 1.8.2 to 1.10.0.
    - Updated to GCC11 toolchain.
    - NNAPI Execution Provider is ported from 1.5.3 (does not contain latest 1.10.0 updates) and it is considered experimental. We do not suggest using it in production.
    - Arm NN and ACL Execution providers are deprecated and will be removed in the future
- PyTorch upgraded to version 1.9.1.
- TIM-VX:
  - Features and improvements:
    - Version update from 1.1.34 to 1.1.37.
    - DMA Buffer support.
    - Support for additional operators (SVDF, GlobalPool2D, AdaptivePool2D, Erf, grouped Conv1D, Signal Frame, RNN Cell, One Hot).

- Support Layout inference for additional operators (Batch Norm, Transpose, Fully Connected with no explicit bias).
- DeepViewRT:
  - Features and improvements:
    - Version update from 2.4.36 to 2.4.37
    - C and Python API for NPU support are available.
    - Align modelrunner plugin with TFLite/Arm NN/ONNX Runtime inference engine.
  - Issues and limitations:
    - Bug fix for deepview-rt library and example codes.

# 9.14 Release notes for LF5.10.72-2.2.0

- TensorFlow Lite:
  - Upgraded to version 2.6.0.
  - VX Delegate changed to external delegate.
  - Optimization of the PCQ Transpose Convolution operator on the NPU hardware accelerator.
  - Python API support external Delegates:
    - With this change, the label\_image.py Python example support the use of external delegates with arguments. See the help for more information.
    - Python API supports using external delegate via the tflite.load\_delegate() call.
    - NNAPI delegate not available in Python API. For the model acceleration on the HW accelerator, the VX delegate can be used:

- Arm Compute Library:
  - Features and improvements:
    - Major version update from <u>21.02</u> to <u>21.08</u>.
  - Issues and limitations:
    - Only the CPU-accelerated NEON backend is being built. Use Arm NN with the VSI NPU backend to leverage acceleration on the GPU or the NPU.
- Arm NN:
  - Features and improvements:
    - Major version update from 21.02 to 21.08.
    - TensorFlow Parser, Caffe Parser and Quantizer were removed and are no longer available. Only ONNX
      Parser, TensorFlow Lite Parser and Arm NN Delegate for TF Lite are now available to load .*tflite* and
      .onnx models.
    - See full list of changes added by the <u>community</u>.
  - Issues and limitations:
    - Only ACL NEON backend is being built. Use the VSI NPU Backend instead of ACL OpenCL to leverage acceleration on the GPU or the NPU.
    - There are significant performance optimizations for the NPU to TransposeConv2D which are not supported in the VSI NPU backend. If your model uses TransposeConv2D heavily try to use TF Lite with VXDelegate instead.
- ONNX Runtime:
  - Features and improvements:
    - Minor version update from 1.8.1 to 1.8.2.

- Experimental Python API enablement including support for all available Execution Providers (CPU, ACL, Arm NN, NNAPI, VSI NPU).
- Added /usr/bin/onnxruntime-1.8.2/onnxruntime\_peft\_test. Use this instead of onnx\_test\_runner to measure performance of your model.
- Fixed verbose logging during inference on NPU.
- Updated ACL and Arm NN Backends to leverage ACL and Arm NN 21.08.
- All ONNX Runtime artifacts are being installer to /usr/bin/onnxruntime-1.8.2 instead of /usr/ bin.
- See full list of changes added by the community.
- Issues and limitations:
  - There are significant performance optimizations for the NPU to TransposeConv2D which are not supported in the VSI NPU Execution Provider. If your model uses TransposeConv2D heavily try to use TF Lite with VXDelegate instead.
- Running SqueezeNet with the NNAPI execution provider produces incorrect results.
- DeepViewRT:
  - Features and improvements:
    - Minor version update from 2.4.30 to 2.4.36.
    - C API for NPU support is available.
    - Performance optimization for DeepViewRT CPU.
    - Bug fix for shuffle layer.
  - Issues and limitations:
    - nn tensor load file ex is one convenience function and not well optimized.

# 10 List of Used Variables

The following table provides the summary of used variables described in this document for the particular inference engine. Use the export command to apply these variables.

Variable name	Description
CNN_PERF	0: Disable (default) 1: Prints the execution time for each operation (requires VIV_VX_ DEBUG_LEVEL=1). If VIV_VX_PROFILE=1 is set, the default value is 1.
NN_EXT_SHOW_PERF	0: Disable (default) 1: Shows more profiling details (requires VIV_VX_DEBUG_LEVEL=1)
PATH_ASSETS	Sets the export path for user assets.
USE_GPU_INFERENCE	Selection between the 3D GPU (1) and the NPU (otherwise).
VIV_VX_CACHE_BINARY_GRAPH_DIR	Specifies the path of the cached NBG. Default is the current work directory.
VIV_VX_DEBUG_LEVEL	0: Disable (default) 1: Prints the debug information of driver on the console. Generally, this environment variable is used together with other environment variables to print logs.
VIV_VX_ENABLE_CACHE_GRAPH_BINARY	0: Disable (default) 1: Enables graph cache mode. The network loads the NBG file to run if the cached NBG file exists. Otherwise, it generates an NBG file. It can save the time for the verification stage.

Table 8. System variables summary

Table 8. System variables summarycontinued			
Variable name	Description		
VIV_MEMORY_PROFILE	0: Disable (default) 1: Prints the memory footprint of the system (CPU) and GPU (VIP) (requires VIV_VX_DEBUG_LEVEL=1)		
VIV_VX_PROFILE	<ul> <li>0: Disable (default)</li> <li>1: Prints the DDR read and write bandwidth, AXI_SRAM read and write bandwidth, and the cycle count of VIP execution. The counter is per-node-process (requires VIV_VX_DEBUG_LEVEL=1).</li> <li>2: Prints the DDR read and write bandwidth, AXI_SRAM read and write bandwidth, and the cycle count of VIP execution. The counter is per-graph-process (requires VIV_VX_DEBUG_LEVEL=1).</li> </ul>		

# Table 8. System variables summary...continued

# 11 Neural Network API Reference

The neural-network operations and corresponding supported API functions are listed in the following table. See also <u>Section 2.2.3</u> for details about supported operators.

Op Category/Name	Android NNAPI 1.2	TensorFlow Lite 2.8.0	ONNX 1.10.0		
Activation					
elu	-	ELU	Elu		
floor	ANEURALNETWORKS_FLOOR	Floor	Floor		
leakyrelu	-	-	LeakyReL		
prelu	ANEURALNETWORKS_PRELU	PRELU	PreLu		
relu	ANEURALNETWORKS_RELU	RELU	ReLu		
relu1	ANEURALNETWORKS_RELU1	RELU1	-		
relu6	ANEURALNETWORKS_RELU6	RELU6	-		
Hard_swish	ANEURALNETWORKS_HARD_SWISH	HARD_SWISH	-		
rsqrt	ANEURALNETWORKS_RSQRT	RSQRT	-		
sigmoid	ANEURALNETWORKS_LOGISTIC	LOGISTIC	Sigmoid		
softmax	ANEURALNETWORKS_SOFTMAX	SOFTMAX	Softmax		
softrelu	-	-	-		
sqrt	ANEURALNETWORKS_SQRT	SQRT	Sqrt		
tanh	ANEURALNETWORKS_TANH	TANH	TanH		
bounded	-	-	-		
linear	-	-	-		
Dense Layers					
dense	-	-	-		
Element Wise					
abs	ANEURALNETWORKS_ABS	ABS	Abs		
add	ANEURALNETWORKS_ADD	ADD	Add		

Table 9. Neural-network operations and supported API functions

Op Category/Name	Android NNAPI 1.2	TensorFlow Lite 2.8.0	ONNX 1.10.0
clip_by_value	-	-	Clip
div	ANEURALNETWORKS_DIdV	DIV	Div
equal	ANEURALNETWORKS_EQUAL	EQUAL	Equal
ехр	ANEURALNETWORKS_EXP	EXP	Exp
log	ANEURALNETWORKS_LOG	LOG	Log
greater	ANEURALNETWORKS_GREATER	GREATER	Greater
greater_equal	ANEURALNETWORKS_GREATER_ EQUAL	GREATER_EQUAL	-
less	ANEURALNETWORKS_LESS	LESS	Less
less_equal	ANEURALNETWORKS_LESS_EQUAL	LESS_EQUAL	-
logical_and	ANEURALNETWORKS_LOGICAL_AND	LOGICAL_AND	And
logical_or	ANEURALNETWORKS_LOGICAL_OR	LOGICAL_OR	Or
minimum	ANEURALNETWORKS_MINIMUM	MINIMUM	Min
maximum	ANEURALNETWORKS_MAXIMUM	MAXIMUM	Мах
multiply	ANEURALNETWORKS_MUL	MUL	Mul
negative	ANEURALNETWORKS_NEG	NEG	Neg
not_equal	ANEURALNETWORKS_NOT_EQUAL	NOT_EQUAL	-
pow	ANEURALNETWORKS_POW	POW	POW
select	ANEURALNETWORKS_SELECT	SELECT	-
square	-	-	-
sub	ANEURALNETWORKS_SUB	SUB	Sub
where	-	-	Where
Image Processing			
resize_bilinear	ANEURALNETWORKS_RESIZE_ BILINEAR	RESIZE_BILINEAR	Unsample
resize_nearest_neighbor	ANEURALNETWORKS_RESIZE_ NEAREST_NEIGHBOR	RESIZE_ NEAREST_ NEIGHBOR	Resize
Matrix Multiplication			
fullconnect	ANEURALNETWORKS_FULLY_ CONNECTED	FULLY_ CONNECTED	-
matrix_mul	-	-	-
Normalization			-
batch_normalize	-	-	BatchNormalization
instance _normalize	-	-	InstanceNormalization
l2normalize	ANEURALNETWORKS_L2_ NORMALIZATION	L2_ NORMALIZATION	-

#### Table 9. Neural-network operations and supported API functions...continued

UG10166 User guide
		**	
Op Category/Name	Android NNAPI 1.2	TensorFlow Lite 2.8.0	ONNX 1.10.0
localresponsenormalization	ANEURALNETWORKS_LOCAL_ RESPONSE_NORMALIZATION	LOCAL_ RESPONSE_ NORMALIZATION	LRN
Reshape			
batch2space	ANEURALNETWORKS_BATCH_TO_ SPACE_ND	BATH_TO_SPACE_ ND	-
concat	ANEURALNETWORKS_ CONCATENATION	CONCATENATION	Concat
depth_to_space	ANEURALNETWORKS_DEPTH_TO_ SPACE	DEPTH_TO_SPACE	DepthToSpace
expanddims	ANEURALNETWORKS_EXPAND_DIMS	EXPAND_DIMS	-
flatten	ANEURALNETWORKS_RESHAPE	-	-
gather	ANEURALNETWORKS_GATHER	GATHER	Gather
pad	ANEURALNETWORKS_PAD	PAD	Pad
permute	ANEURALNETWORKS_TRANSPOSE	TRANSPOSE	Transpose
reducemean	ANEURALNETWORKS_MEAN	MEAN	ReduceMean
reducesum	ANEURALNETWORKS_SUM	REDUCE_SUM	ReduseSum
gathernd	-	-	GatherND
reducemax	ANEURALNETWORKS_REDUCE_MAX	REDUCE_MAX	ReduceMax
reducemin	ANEURALNETWORKS_REDUCE_MIN	REDUCE_MIN	ReduceMin
reduceproduct	-	-	-
reshape	ANEURALNETWORKS_RESHAPE	RESHAPE	Reshape
reverse	-	-	ReverseSequence
slice	ANEURALNETWORKS_SLICE	SLICE	Slice
space2batch	ANEURALNETWORKS_SPACE_TO_ BATCH_ND	SPACE_TO_ BATCH_ND	-
split	ANEURALNETWORKS_SPLIT	SPLIT	Split
squeeze	ANEURALNETWORKS_SQUEEZE	SQUEEZE	Squeeze
strided_slice	ANEURALNETWORKS_STRIDED_SLICE	STRIDED_SLICE	-
unstack	-	-	-
RNN			
gru	-	-	GRU
Istm	-	UNIDIRECTIONAL_ SEQUEENCE_ LSTM	-
Istmunit	ANEURALNETWORKS_LSTM	LSTM	LSTM
rnn	ANEURALNETWORKS_RNN	RNN	-
Sliding Window	·		

#### Table 9. Neural-network operations and supported API functions...continued

UG10166

User guide

Op Category/Name	Android NNAPI 1.2	TensorFlow Lite 2.8.0	ONNX 1.10.0
avg_pool	ANEURALNETWORKS_AVERAGE_POOL	AVERAGE_POOL_ 2D	AveragePool
convolution	ANEURALNETWORKS_CONV_2D	CONV_2D	Conv
deconvolution	ANEURALNETWORKS_TRANSPOSE_ CONV_2D	TRANSPOSE_ CONV	ConvTranspose
depthhwise_convolution	ANEURALNETWORKS_DEPTHWISE_ CONV_2D	DEPTHWISE_ CONV_2D	-
Log_softmax	ANEURALNETWORKS_LOG_SOFTMAX	LOG_SOFTMAX	Logsoftmax
I2pooling	ANEURALNETWORKS_L2_POOL	L2_POOL_2D	-
max_pool	ANEURALNETWORKS_MAX_POOL	MAX_POOL_2D	MaxPool
Others		·	
argmax	ANEURALNETWORKS_ARGMAX	ARGMAX	ArgMax
argmin	ANEURALNETWORKS_ARGMIN	ARGMIN	ArgMin
dequantize	ANEURALNETWORKS_DEQUANTIZE	DEQUANTIZE	DequantizeLinear
quantize	ANEURALNETWORKS_QUANTIZE	QUANTIZE	QuantizeLinear
roi_pool	ANEURALNETWORKS_ROI_ALIGN	-	-
shuffle_channel	ANEURALNETWORKS_CHANNEL_ SHUFFLE	-	-
tile	ANEURALNETWORKS_TILE	TILE	Tile
svdf	ANEURALNETWORKS_SVDF	SVDF	-
embedding_lookup	ANEURALNETWORKS_EMBEDDING_ LOOKUP	EMBEDDING_ LOOKUP	-
cast	ANEURALNETWORKS_CAST	CAST	Cast
ssd	-	-	-

#### Table 9. Neural-network operations and supported API functions...continued

## 12 OVXLIB Operation Support with GPU

This section provides a summary of the neural network OVXLIB operations supported by the NXP Graphics Processing Unit (GPU) IP with hardware support for OpenVX and OpenCL and a compatible Software stacks. OVXLIB operations are listed in the following table.

The following abbreviations are used for format types:

- asym-u8: asymmetric\_affine-uint8
- **asym-i8**: asymmetric\_affine-int8
- fp32: float32
- pc-sym-i8: perchannel\_symmetric\_int8
- **fp16**: float16
- **bool8**: bool8
- int16: int16
- int32: int32

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine		
	Input	Kernel	Output	OpenVX	OpenCL	
Basic Operations						
VSI_NN_OP_CONV2D	asym-u8	asym-u8	asym-u8	1	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8	<b>√</b>	$\checkmark$	
	fp32	fp32	fp32	1	$\checkmark$	
	fp16	fp16	fp16	1	$\checkmark$	
VSI_NN_OP_CONV1D	asym-u8	asym-u8	asym-u8	1	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8	1	$\checkmark$	
	fp32	fp32	fp32	1	$\checkmark$	
	fp16	fp16	fp16	1	$\checkmark$	
VSI_NN_OP_DEPTHWISE_	asym-u8	asym-u8	asym-u8	1		
CONV1D	asym-i8	asym-i8	asym-i8	1		
VSI_NN_OP_DECONVOLUTION1D	asym-u8	asym-u8	asym-u8	1	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8	1	$\checkmark$	
	fp32	fp32	fp32	<b>√</b>	$\checkmark$	
	fp16	fp16	fp16	1	$\checkmark$	
VSI_NN_OP_DECONVOLUTION	asym-u8	asym-u8	asym-u8	1	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8	1	$\checkmark$	
	fp32	fp32	fp32	1	$\checkmark$	
	fp16	fp16	fp16	<b>√</b>	$\checkmark$	
VSI_NN_OP_FCL	asym-u8	asym-u8	asym-u8	<b>√</b>	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8	1	✓	
	fp32	fp32	fp32	1	✓	
	fp16	fp16	fp16	1	✓	
VSI_NN_OP_GROUPED_CONV1D	asym-u8	asym-u8	asym-u8	1	✓	
	asym-i8	pc-sym-i8	asym-i8	✓	✓	
	fp32	fp32	fp32	1	✓	
	fp16	fp16	fp16	✓	✓	
VSI_NN_OP_GROUPED_CONV2D	asym-u8	asym-u8	asym-u8	✓	✓	
	asym-i8	pc-sym-i8	asym-i8	✓	✓	
	fp32	fp32	fp32	1	✓	
	fp16	fp16	fp16	<b>√</b>	✓	
Activation Operations			· ·	· ·		
VSI_NN_OP_ELU	asym-u8		asym-u8	<b>√</b>	$\checkmark$	
	asym-i8		asym-i8	1	$\checkmark$	
	fp32		fp32	<b>v</b>	✓	

### Table 10. OVXLIB operation support with GPU

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine		
	Input	Kernel	Output	OpenVX	OpenCL	
	fp16		fp16	<b>v</b>	✓	
VSI_NN_OP_HARD_SIGMOID	asym-u8		asym-u8	<b>v</b>	✓	
	asym-i8		asym-i8	1	$\checkmark$	
	fp32		fp32	<b>v</b>	✓	
	fp16		fp16	1	✓	
VSI_NN_OP_SWISH	asym-u8		asym-u8	1	$\checkmark$	
	asym-i8		asym-i8	1	$\checkmark$	
	fp32		fp32	<b>√</b>	$\checkmark$	
	fp16		fp16	<b>√</b>	$\checkmark$	
VSI_NN_OP_LEAKY_RELU	asym-u8		asym-u8	<b>v</b>	$\checkmark$	
	asym-i8		asym-i8	1	$\checkmark$	
	fp32		fp32	<b>v</b>	$\checkmark$	
	fp16		fp16	<b>v</b>	$\checkmark$	
VSI_NN_OP_PRELU	asym-u8		asym-u8	1	✓	
	asym-i8		asym-i8	1	✓	
	fp32		fp32	<b>v</b>	$\checkmark$	
	fp16		fp16	1	$\checkmark$	
VSI_NN_OP_RELU	asym-u8		asym-u8	<b>v</b>	$\checkmark$	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	✓	
	fp16		fp16	1	✓	
VSI_NN_OP_RELUN	asym-u8		asym-u8	1	✓	
	asym-i8		asym-i8	✓	√	
	fp32		fp32	✓	$\checkmark$	
	fp16		fp16	<b>√</b>	$\checkmark$	
VSI_NN_OP_RSQRT	asym-u8		asym-u8	<b>√</b>	$\checkmark$	
	asym-i8		asym-i8	<b>√</b>	$\checkmark$	
	fp32		fp32	✓	√	
	fp16		fp16	✓	√	
VSI_NN_OP_SIGMOID	asym-u8		asym-u8	✓	$\checkmark$	
	asym-i8		asym-i8	<b>√</b>	$\checkmark$	
	fp32		fp32	✓	✓	
	fp16		fp16	1	$\checkmark$	
VSI_NN_OP_SOFTRELU	asym-u8		asym-u8	1	$\checkmark$	
	asym-i8		asym-i8	✓	$\checkmark$	

#### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors	Tensors			Execution Engine		
	Input	Kernel	Output	OpenVX	OpenCL		
	fp32		fp32	√	1		
	fp16		fp16	✓	✓		
VSI_NN_OP_SQRT	asym-u8		asym-u8	✓	✓		
	asym-i8		asym-i8	1	<b>√</b>		
	fp32		fp32	√	✓		
	fp16		fp16	√	$\checkmark$		
VSI_NN_OP_TANH	asym-u8		asym-u8	√	√		
	asym-i8		asym-i8	√	1		
	fp32		fp32	√	1		
	fp16		fp16	√	1		
VSI_NN_OP_ABS	asym-u8		asym-u8	√	√		
	asym-i8		asym-i8	√	√		
	fp32		fp32	√	✓		
	fp16		fp16	✓	1		
/SI_NN_OP_CLIP	asym-u8		asym-u8	√	1		
	asym-i8		asym-i8	√	1		
	fp32		fp32	√	1		
	fp16		fp16	√	1		
VSI_NN_OP_EXP	asym-u8		asym-u8	✓	1		
	asym-i8		asym-i8	✓	1		
	fp32		fp32	✓	1		
	fp16		fp16	✓	1		
VSI_NN_OP_LOG	asym-u8		asym-u8	1	1		
	asym-i8		asym-i8	√	1		
	fp32		fp32	✓	1		
	fp16		fp16	✓	1		
/SI_NN_OP_NEG	asym-u8		asym-u8	√	1		
	asym-i8		asym-i8	√	1		
	fp32		fp32	√	1		
	fp16		fp16	√	√		
/SI_NN_OP_MISH	asym-u8		asym-u8	✓	$\checkmark$		
	asym-i8		asym-i8	√	$\checkmark$		
	fp32		fp32	✓	$\checkmark$		
	fp16		fp16	✓	✓		
VSI_NN_OP_LINEAR	asym-u8		asym-u8	1	1		

#### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine	
	Input	Kernel	Output	OpenVX	OpenCL
	asym-i8		asym-i8	√	1
	fp32		fp32	✓	1
	fp16		fp16	✓	1
VSI_NN_OP_ERF	asym-u8		asym-u8	✓	1
	asym-i8		asym-i8	✓	✓
	fp32		fp32	✓	1
	fp16		fp16	√	1
VSI_NN_OP_SOFTMAX	asym-u8		asym-u8	✓	1
	asym-i8		asym-i8	✓	1
	fp32		fp32	√	1
	fp16		fp16	√	1
VSI_NN_OP_LOG_SOFTMAX	asym-u8		asym-u8	√	1
	asym-i8		asym-i8	√	1
	fp32		fp32	√	1
	fp16		fp16	√	1
VSI_NN_OP_SQUARE	asym-u8		asym-u8	√	1
	asym-i8		asym-i8	√	1
	fp32		fp32	✓	1
	fp16		fp16	✓	1
VSI_NN_OP_SIN	asym-u8		asym-u8	√	1
	asym-i8		asym-i8	√	1
	fp32		fp32	√	1
	fp16		fp16	√	1
Elementwise Operations		I	I		I.
VSI_NN_OP_ADD	asym-u8		asym-u8	$\checkmark$	1
	asym-i8		asym-i8	√	1
	fp32		fp32	√	1
	fp16		fp16	√	1
VSI_NN_OP_SUBTRACT	asym-u8		asym-u8	√	1
	asym-i8		asym-i8	√	1
	fp32		fp32	√	$\checkmark$
	fp16		fp16	√	$\checkmark$
VSI_NN_OP_MULTIPLY	asym-u8		asym-u8	√	$\checkmark$
	asym-i8		asym-i8	√	$\checkmark$
	fp32		fp32	√	1

#### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine	
	Input	Kernel	Output	OpenVX	OpenCL
	fp16		fp16	<b>v</b>	<b>√</b>
VSI_NN_OP_DIVIDE	asym-u8		asym-u8	<b>√</b>	1
	asym-i8		asym-i8	1	<b>v</b>
	fp32		fp32	1	<b>√</b>
	fp16		fp16	1	<b>√</b>
VSI_NN_OP_MAXIMUN	asym-u8		asym-u8	<b>v</b>	<b>√</b>
	asym-i8		asym-i8	<b>√</b>	1
	fp32		fp32	<b>√</b>	1
	fp16		fp16	<b>√</b>	1
VSI_NN_OP_MINIMUM	asym-u8		asym-u8	<b>√</b>	✓
	asym-i8		asym-i8	<b>√</b>	1
	fp32		fp32	<b>√</b>	1
	fp16		fp16	<b>√</b>	1
VSI_NN_OP_POW	asym-u8		asym-u8	1	1
	asym-i8		asym-i8	<b>√</b>	✓
	fp32		fp32	<b>√</b>	1
	fp16		fp16	<b>√</b>	1
VSI_NN_OP_FLOORDIV	asym-u8		asym-u8	<b>√</b>	1
	asym-i8		asym-i8	<b>√</b>	1
	fp32		fp32	1	1
	fp16		fp16	1	1
VSI_NN_OP_MATRIXMUL	asym-u8		asym-u8	1	1
	asym-i8		asym-i8	1	1
	fp32		fp32	1	1
	fp16		fp16	1	1
VSI_NN_OP_RELATIONAL_OPS	asym-u8		bool8	1	1
	asym-i8		bool8	1	1
	fp32		bool8	<b>√</b>	1
	fp16		bool8	1	1
	bool8		bool8	<b>√</b>	1
VSI_NN_OP_LOGICAL_OPS	bool8		bool8	✓	1
VSI_NN_OP_LOGICAL_NOT	bool8		bool8	<b>√</b>	1
VSI_NN_OP_SELECT	asym-u8		asym-u8	✓	$\checkmark$
	asym-i8		asym-i8	✓	$\checkmark$
	fp32		fp32	1	1

#### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine		
	Input	Kernel	Output	OpenVX	OpenCL	
	fp16		fp16	1	1	
	bool8		bool8	✓	√	
VSI_NN_OP_ADDN	asym-u8		asym-u8	1	√	
	asym-i8		asym-i8	1	√	
	fp32		fp32	1	✓	
	fp16		fp16	<b>√</b>	√	
Normalization Operations						
VSI_NN_OP_BATCH_NORM	asym-u8		asym-u8	√	$\checkmark$	
	asym-i8		asym-i8	√	√	
	fp32		fp32	√	√	
	fp16		fp16	√	√	
VSI_NN_OP_LRN	asym-u8		asym-u8	√	√	
	asym-i8		asym-i8	√	√	
	fp32		fp32	1	$\checkmark$	
	fp16		fp16	√	$\checkmark$	
VSI_NN_OP_LRN2	asym-u8		asym-u8	√	$\checkmark$	
	asym-i8		asym-i8	1	$\checkmark$	
	fp32		fp32	1	$\checkmark$	
	fp16		fp16	√	$\checkmark$	
VSI_NN_OP_L2_NORMALIZE	asym-u8		asym-u8	1	$\checkmark$	
	asym-i8		asym-i8	1	$\checkmark$	
	fp32		fp32	1	$\checkmark$	
	fp16		fp16	1	$\checkmark$	
VSI_NN_OP_L2NORMALZESCALE	asym-u8		asym-u8	1	$\checkmark$	
	asym-i8		asym-i8	1	$\checkmark$	
	fp32		fp32	1	$\checkmark$	
	fp16		fp16	1	$\checkmark$	
VSI_NN_OP_LAYER_NORM	asym-u8		asym-u8	✓	$\checkmark$	
	asym-i8		asym-i8	1	$\checkmark$	
	fp32		fp32	1	$\checkmark$	
	fp16		fp16	✓	√	
VSI_NN_OP_INSTANCE_NORM	asym-u8		asym-u8	1	$\checkmark$	
	asym-i8		asym-i8	1	$\checkmark$	
	fp32		fp32	1	$\checkmark$	
	fp16		fp16	$\checkmark$	$\checkmark$	

### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine		
	Input	Kernel	Output	OpenVX	OpenCL	
VSI_NN_OP_GROUP_NORM	asym-u8		asym-u8	√	$\checkmark$	
	asym-i8		asym-i8	√	✓	
	fp32		fp32	√	$\checkmark$	
	fp16		fp16	√	$\checkmark$	
VSI_NN_OP_BATCHNORM_SINGLE	asym-u8		asym-u8	√	$\checkmark$	
	asym-i8		asym-i8	✓	$\checkmark$	
	fp32		fp32	✓	$\checkmark$	
	fp16		fp16	✓	✓	
VSI_NN_OP_MOMENTS	asym-u8		asym-u8	✓	✓	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	✓	
	fp16		fp16	✓	$\checkmark$	
Reshape Operations		·	·			
VSI_NN_OP_EXPAND_BROADCAST	asym-u8		asym-u8	✓	$\checkmark$	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	$\checkmark$	
	fp16		fp16	✓	✓	
VSI_NN_OP_SLICE	asym-u8		asym-u8	✓	✓	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	✓	
	fp16		fp16	✓	√	
VSI_NN_OP_SPLIT	asym-u8		asym-u8	✓	√	
	asym-i8		asym-i8	✓	√	
	fp32		fp32	✓	$\checkmark$	
	fp16		fp16	✓	$\checkmark$	
VSI_NN_OP_CONCAT	asym-u8		asym-u8	✓	$\checkmark$	
	asym-i8		asym-i8	✓	$\checkmark$	
	fp32		fp32	✓	√	
	fp16		fp16	✓	√	
VSI_NN_OP_STACK	asym-u8		asym-u8	✓	✓	
	asym-i8		asym-i8	✓	$\checkmark$	
	fp32		fp32	✓	$\checkmark$	
	fp16		fp16	1	$\checkmark$	
VSI_NN_OP_UNSTACK	asym-u8		asym-u8	1	✓	
	asym-i8		asym-i8	$\checkmark$	$\checkmark$	

#### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine		
	Input	Kernel	Output	OpenVX	OpenCL	
	fp32		fp32	✓	✓	
	fp16		fp16	1	✓	
VSI_NN_OP_RESHAPE	asym-u8		asym-u8	1	✓	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	✓	
	fp16		fp16	✓	✓	
VSI_NN_OP_SQUEEZE	asym-u8		asym-u8	✓	✓	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	✓	
	fp16		fp16	✓	✓	
VSI_NN_OP_PERMUTE	asym-u8		asym-u8	✓	✓	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	✓	
	fp16		fp16	✓	✓	
VSI_NN_OP_REORG	asym-u8		asym-u8	✓	✓	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	✓	
	fp16		fp16	✓	✓	
VSI_NN_OP_SPACE2DEPTH	asym-u8		asym-u8	√	$\checkmark$	
	asym-i8		asym-i8	√	$\checkmark$	
	fp32		fp32	√	$\checkmark$	
	fp16		fp16	√	$\checkmark$	
VSI_NN_OP_DEPTH2SPACE	asym-u8		asym-u8	√	$\checkmark$	
	asym-i8		asym-i8	√	$\checkmark$	
	fp32		fp32	✓	✓	
	fp16		fp16	✓	✓	
VSI_NN_OP_BATCH2SPACE	asym-u8		asym-u8	√	$\checkmark$	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	✓	
	fp16		fp16	√	$\checkmark$	
VSI_NN_OP_SPACE2BATCH	asym-u8		asym-u8	√	$\checkmark$	
	asym-i8		asym-i8	√	$\checkmark$	
	fp32		fp32	√	$\checkmark$	
	fp16		fp16	√	√	
VSI_NN_OP_PAD	asym-u8		asym-u8	✓	$\checkmark$	

#### Table 10. OVXLIB operation support with GPU...continued

UG10166

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine		
	Input	Kernel	Output	OpenVX	OpenCL	
	asym-i8		asym-i8	✓	✓	
	fp32		fp32	✓	✓	
	fp16		fp16	1	✓	
VSI_NN_OP_REVERSE	asym-u8		asym-u8	1	✓	
	asym-i8		asym-i8	1	✓	
	fp32		fp32	1	✓	
	fp16		fp16	1	✓	
VSI_NN_OP_STRIDED_SLICE	asym-u8		asym-u8	1	✓	
	asym-i8		asym-i8	1	✓	
	fp32		fp32	1	√	
	fp16		fp16	1	✓	
VSI_NN_OP_CROP	asym-u8		asym-u8	1	✓	
	asym-i8		asym-i8	1	✓	
	fp32		fp32	✓	$\checkmark$	
	fp16		fp16	√	$\checkmark$	
VSI_NN_OP_REDUCE	asym-u8		asym-u8	√	√	
	asym-i8		asym-i8	$\checkmark$	V	
	fp32		fp32	$\checkmark$	$\checkmark$	
	fp16		fp16	$\checkmark$	$\checkmark$	
VSI_NN_OP_ARGMX	asym-u8		asym-u8/int16/ int32	1	✓	
	asym-i8		asym-u8/int16/ int32	1	√	
	fp32		int32	$\checkmark$	$\checkmark$	
	fp16		asym-u8/int16/ int32	1	√	
VSI_NN_OP_ARGMIN	asym-u8		asym-u8/int16/ int32	1	√	
	asym-i8		asym-u8/int16/ int32	1	√	
	fp32		int32	√	$\checkmark$	
	fp16		asym-u8/int16/ int32	✓	✓	
VSI_NN_OP_SHUFFLECHANNEL	asym-u8		asym-u8	$\checkmark$	<b>√</b>	
	asym-i8		asym-i8	√	<b>√</b>	
	fp32		fp32	$\checkmark$	V	
	fp16		fp16	1	1	

#### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine		
	Input	Kernel	Output	OpenVX	OpenCL	
RNN Operations						
VSI_NN_OP_LSTMUNIT_OVXLIB	asym-u8	asym-u8	asym-u8	✓	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8	✓	✓	
	fp32	fp32	fp32	✓	✓	
	fp16	fp16	fp16	✓	✓	
VSI_NN_OP_LSTM_OVXLIB	asym-u8	asym-u8	asym-u8	✓	✓	
	asym-i8	pc-sym-i8	asym-i8	✓	✓	
	fp32	fp32	fp32	✓	$\checkmark$	
	fp16	fp16	fp16	✓	✓	
VSI_NN_OP_GRUCELL_OVXLIB	asym-u8	asym-u8	asym-u8	√	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8	✓	√	
	fp32	fp32	fp32	✓	√	
	fp16	fp16	fp16	√	√	
VSI_NN_OP_GRU_OVXLIB	asym-u8	asym-u8	asym-u8	✓	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8	√	$\checkmark$	
	fp32	fp32	fp32	√	$\checkmark$	
	fp16	fp16	fp16	√	$\checkmark$	
VSI_NN_OP_SVDF	asym-u8	asym-u8	asym-u8	√	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8	✓	$\checkmark$	
	fp32	fp32	fp32	✓	✓	
	fp16	fp16	fp16	✓	$\checkmark$	
Pooling Operations						
VSI_NN_OP_ROI_POOL	asym-u8		asym-u8	✓	$\checkmark$	
	asym-i8		asym-i8	✓	$\checkmark$	
	fp32		fp32	✓	$\checkmark$	
	fp16		fp16	✓	$\checkmark$	
VSI_NN_OP_POOLWITHARGMAX	asym-u8		asym-u8	$\checkmark$	$\checkmark$	
	asym-i8		asym-i8	✓	$\checkmark$	
	fp32		fp32	✓	$\checkmark$	
	fp16		fp16	✓	$\checkmark$	
VSI_NN_OP_UPSAMPLE	asym-u8		asym-u8	$\checkmark$	$\checkmark$	
	asym-i8		asym-i8	✓	$\checkmark$	
	fp32		fp32	✓	✓	
	fp16		fp16	✓	✓	
Miscellaneous Operations						

#### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution Engine		
	Input	Kernel	Output	OpenVX	OpenCL	
VSI_NN_OP_PROPOSAL	asym-u8		asym-u8	<b>√</b>		
	asym-i8		asym-i8	<b>√</b>		
	fp32		fp32	√		
	fp16		fp16	<b>√</b>		
VSI_NN_OP_VARIABLE	asym-u8		asym-u8	1	1	
	asym-i8		asym-i8	1	1	
	fp32		fp32	1	1	
	fp16		fp16	1	1	
VSI_NN_OP_DROPOUT	asym-u8		asym-u8	<b>√</b>	1	
	asym-i8		asym-i8	1	1	
	fp32		fp32	1	1	
	fp16		fp16	1	1	
VSI_NN_OP_RESIZE	asym-u8		asym-u8	1	1	
	asym-i8		asym-i8	1	1	
	fp32		fp32	1	1	
	fp16		fp16	1	1	
VSI_NN_OP_INTERP	asym-u8		asym-u8	1	1	
	asym-i8		asym-i8	1	1	
	fp32		fp32	1	1	
	fp16		fp16	1	1	
VSI_NN_OP_DATACONVERT	asym-u8		asym-u8	1	1	
	asym-i8		asym-i8	1	1	
	fp32		fp32	1	1	
	fp16		fp16	<b>√</b>	1	
VSI_NN_OP_A_TIMES_B_PLUS_C	asym-u8		asym-u8	1	1	
	asym-i8		asym-i8	1	1	
	fp32		fp32	✓	1	
	fp16		fp16	1	1	
VSI_NN_OP_FLOOR	asym-u8		asym-u8	1	1	
	asym-i8		asym-i8	1	1	
	fp32		fp32	1	1	
	fp16		fp16	√	1	
VSI_NN_OP_EMBEDDING_	asym-u8		asym-u8	✓	1	
LOOKUP	asym-i8		asym-i8	✓	1	
	fp32		fp32	1	1	
	L					

#### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution	Engine
	Input	Kernel	Output	OpenVX	OpenCL
	fp16		fp16	<b>v</b>	✓
VSI_NN_OP_GATHER	asym-u8		asym-u8	<b>v</b>	✓
	asym-i8		asym-i8	<b>v</b>	✓
	fp32		fp32	<b>v</b>	√
	fp16		fp16	<b>v</b>	√
VSI_NN_OP_GATHER_ND	asym-u8		asym-u8	<b>v</b>	√
	asym-i8		asym-i8	<b>v</b>	√
	fp32		fp32	<b>√</b>	$\checkmark$
	fp16		fp16	<b>√</b>	$\checkmark$
VSI_NN_OP_SCATTER_ND	asym-u8	Kernel         Output         Open           ill         fp16         ✓           iull         asym-u8         ✓           iull         asym-u8         ✓           iull         asym-u8         ✓           iull         asym-u8         ✓           iull         fp16         ✓           iull         fp12         ✓           iull         asym-u8         ✓           iull         fp16         ✓           iull         asym-u8         ✓           iull         asym-u8         ✓           iull         fp16         ✓           iull         asym-u8         ✓           iull         asym-u8         ✓           iull         asym-u8         ✓           iull         fp16         ✓           iull         fp16         ✓           iull         asym-u8         ✓           iull         fp16         ✓           iull         asym-u8         ✓           iull         asym-u8         ✓           iull         asym-u8         ✓           iull         asym-u8         ✓           iul	<b>√</b>	$\checkmark$	
	asym-i8		asym-i8	<b>v</b>	$\checkmark$
	fp32		fp32	<b>v</b>	$\checkmark$
	fp16		fp16	<b>v</b>	$\checkmark$
VSI_NN_OP_TILE	asym-u8		asym-u8	1	√
	asym-i8		asym-i8	1	$\checkmark$
	fp32		fp32	1	$\checkmark$
	fp16		fp16	1	$\checkmark$
VSI_NN_OP_RELU_KERAS	asym-u8		asym-u8	1	$\checkmark$
	asym-i8		asym-i8	✓	✓
	fp32		fp32	✓	√
	fp16		fp16	✓	$\checkmark$
VSI_NN_OP_ELTWISEMAX	asym-u8		asym-u8	✓	$\checkmark$
	asym-i8		asym-i8	<b>√</b>	$\checkmark$
	fp32		fp32	<b>√</b>	$\checkmark$
	fp16		fp16	$\checkmark$	$\checkmark$
VSI_NN_OP_INSTANCE_NORM	asym-u8		asym-u8	<b>√</b>	$\checkmark$
	asym-i8		asym-i8	$\checkmark$	$\checkmark$
	fp32		fp32	$\checkmark$	$\checkmark$
	fp16		fp16	<b>√</b>	$\checkmark$
VSI_NN_OP_FCL2	asym-u8		asym-u8	<b>√</b>	$\checkmark$
	asym-i8		asym-i8	$\checkmark$	$\checkmark$
	fp32		fp32	$\checkmark$	$\checkmark$
	fp16		fp16	<b>√</b>	✓
VSI_NN_OP_POOL	asym-u8		asym-u8	<b>√</b>	✓
	asym-i8		asym-i8	✓	$\checkmark$

#### Table 10. OVXLIB operation support with GPU...continued

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors			Execution	Engine
	Input	Kernel	Output	OpenVX	OpenCL
	fp32		fp32	<b>v</b>	√
	fp16		fp16	1	✓
VSI_NN_OP_SIGNAL_FRAME	asym-u8		asym-u8	<b>√</b>	
	asym-i8		asym-i8	<b>√</b>	
	fp32		fp32	<b>√</b>	
	fp16		fp16	<b>√</b>	
VSI_NN_OP_CONCATSHIFT	asym-u8		asym-u8	<b>√</b>	√
	asym-i8		asym-i8	<b>√</b>	$\checkmark$
	fp32		fp32	1	$\checkmark$
	fp16		fp16	<b>√</b>	$\checkmark$
VSI_NN_OP_UPSAMPLESCALE	asym-u8		asym-u8	<b>√</b>	
	asym-i8		asym-i8	<b>√</b>	
	fp16		fp16	1	
VSI_NN_OP_ROUND	asym-u8		asym-u8	<b>√</b>	√
	asym-i8		asym-i8	<b>√</b>	√
	fp32		fp32	<b>√</b>	√
	fp16		fp16	<b>√</b>	√
VSI_NN_OP_CEIL	asym-u8		asym-u8	<b>v</b>	✓
	asym-i8		asym-i8	<b>√</b>	$\checkmark$
	fp32		fp32	1	$\checkmark$
	fp16		fp16	1	$\checkmark$
VSI_NN_OP_SEQUENCE_MASK	asym-u8		asym-u8	<b>√</b>	√
	asym-i8		asym-i8	<b>√</b>	√
	fp32		fp32	<b>√</b>	√
	fp16		fp16	✓	$\checkmark$
VSI_NN_OP_REPEAT	asym-u8		asym-u8	1	$\checkmark$
	asym-i8		asym-i8	<b>√</b>	√
	fp32		fp32	<b>√</b>	√
	fp16		fp16	<b>v</b>	✓
VSI_NN_OP_ONE_HOT	asym-u8		asym-u8	<b>v</b>	✓
	asym-i8		asym-i8	1	$\checkmark$
	fp32		fp32	1	√
	fp16		fp16	1	✓
VSI_NN_OP_CAST	asym-u8		asym-u8	<b>√</b>	✓
	asym-i8		asym-i8	1	√

#### Table 10. OVXLIB operation support with GPU...continued

OVXLIB Operations	Tensors	Execution Engine				
	Input	Kernel	Output	OpenVX	OpenCL	
	fp32		fp32	<b>√</b>	✓	
	fp16		fp16	✓	✓	

#### Table 10. OVXLIB operation support with GPU...continued

## 13 OVXLIB Operation Support with NPU

This section provides a summary of the neural network OVXLIB operations supported by the NXP Neural Processor Unit (NPU) IP and a compatible Software stacks. OVXLIB operations are listed in the following table.

The following abbreviations are used for format types:

- **asym-u8**: asymmetric\_affine-uint8
- asym-i8: asymmetric\_affine-int8
- fp32: float32
- pc-sym-i8: perchannel\_symmetric-int8
- fp16: float16
- **bool8**: bool8
- int16: int16
- int32: int32

The following abbreviations are used to reference key Execution Engines (NPU) in the hardware:

- NN: Neural-Network Engine
- PPU: Parallel Processing Unit
- TP: Tensor Processor

#### Table 11. OVXLIB operation support with NPU

OVXLIB Operations	Tensors			Execut	ion Engine	e (NPU)
	Input	Kernel	Output	NN	TP	PPU
Basic Operations			·		·	
VSI_NN_OP_CONV2D	asym-u8	asym-u8	asym-u8	1		
	asym-i8	pc-sym-i8	asym-i8	1		$\checkmark$
	fp32	fp32	fp32			$\checkmark$
	fp16	fp16	fp16			$\checkmark$
VSI_NN_OP_CONV1D	asym-u8	asym-u8	asym-u8	1		
	asym-i8	pc-sym-i8	asym-i8	1		$\checkmark$
	fp32	fp32	fp32			$\checkmark$
	fp16	fp16	fp16			$\checkmark$
VSI_NN_OP_CONV3D	asym-u8	asym-u8	asym-u8	1		
	asym-i8	pc-sym-i8	asym-i8	1		$\checkmark$
	fp32	fp32	fp32			$\checkmark$
	fp16	fp16	fp16			$\checkmark$
VSI_NN_OP_DEPTHWISE_CONV1D	asym-u8	asym-u8	asym-u8			$\checkmark$

OVXLIB Operations	Tensors			Execut	tion Engine	e (NPU)
	Input	Kernel	Output	NN	TP	PPU
	asym-i8	asym-i8	asym-i8			√
VSI_NN_OP_DECONVOLUTION	asym-u8	asym-u8	asym-u8	1		
	asym-i8	pc-sym-i8	asym-i8	1		$\checkmark$
	fp32	fp32	fp32			$\checkmark$
	fp16	fp16	fp16			$\checkmark$
VSI_NN_OP_DECONVOLUTION1D	asym-u8	asym-u8	asym-u8	1		
	asym-i8	pc-sym-i8	asym-i8	<b>√</b>		√
	fp32	fp32	fp32			$\checkmark$
	fp16	fp16	fp16			$\checkmark$
VSI_NN_OP_FCL	asym-u8	asym-u8	asym-u8		√	
	asym-i8	pc-sym-i8	asym-i8		√	$\checkmark$
	fp32	fp32	fp32			$\checkmark$
	fp16	fp16	fp16		√	
VSI_NN_OP_GROUPED_CONV1D	asym-u8	asym-u8	asym-u8	1		
	asym-i8	pc-sym-i8	asym-i8	1		$\checkmark$
	fp32	fp32	fp32			$\checkmark$
	fp16	fp16	fp16			$\checkmark$
VSI_NN_OP_GROUPED_CONV2D	asym-u8	asym-u8	asym-u8			
	asym-i8	pc-sym-i8	asym-i8			$\checkmark$
	fp32	fp32	fp32			$\checkmark$
	fp16	fp16	fp16			$\checkmark$
Activation Operations		I	1			
VSI_NN_OP_ELU	asym-u8		asym-u8			$\checkmark$
	asym-i8		asym-i8			$\checkmark$
	fp32		fp32			$\checkmark$
	fp16		fp16			$\checkmark$
VSI_NN_OP_HARD_SIGMOID	asym-u8		asym-u8			$\checkmark$
	asym-i8		asym-i8			$\checkmark$
	fp32		fp32			$\checkmark$
	fp16		fp16			$\checkmark$
VSI_NN_OP_SWISH	asym-u8		asym-u8		✓	
	asym-i8		asym-i8		√	
	fp32		fp32			1
	fp16		fp16		$\checkmark$	
VSI_NN_OP_LEAKY_RELU	asym-u8		asym-u8		1	

#### Table 11. OVXLIB operation support with NPU...continued

OVXLIB Operations	Tensors			Execut	tion Engine	Engine (NPU)		
	Input	Kernel	Output	NN	ТР	PPU		
	asym-i8		asym-i8		√			
	fp32		fp32			1		
	fp16		fp16		√			
VSI_NN_OP_PRELU	asym-u8		asym-u8		√			
	asym-i8		asym-i8		√			
	fp32		fp32			√		
	fp16		fp16		√			
VSI_NN_OP_RELU	asym-u8		asym-u8		$\checkmark$			
	asym-i8		asym-i8		$\checkmark$			
	fp32		fp32			$\checkmark$		
	fp16		fp16		√			
VSI_NN_OP_RELUN	asym-u8		asym-u8		√			
	asym-i8		asym-i8		√			
	fp32		fp32			$\checkmark$		
	fp16		fp16		√			
VSI_NN_OP_RSQRT	asym-u8		asym-u8			$\checkmark$		
	asym-i8		asym-i8			$\checkmark$		
	fp32		fp32			1		
	fp16		fp16			√		
VSI_NN_OP_SIGMOID	asym-u8		asym-u8		$\checkmark$			
	asym-i8		asym-i8		$\checkmark$			
	fp32		fp32			$\checkmark$		
	fp16		fp16		√			
VSI_NN_OP_SOFTRELU	asym-u8		asym-u8			√		
	asym-i8		asym-i8			$\checkmark$		
	fp32		fp32			1		
	fp16		fp16			1		
VSI_NN_OP_SQRT	asym-u8		asym-u8			1		
	asym-i8		asym-i8			1		
	fp32		fp32			1		
	fp16		fp16			✓		
VSI_NN_OP_TANH	asym-u8		asym-u8		✓			
	asym-i8		asym-i8		$\checkmark$			
	fp32		fp32			V		
	fp16		fp16		$\checkmark$			

#### Table 11. OVXLIB operation support with NPU...continued

OVXLIB Operations	Tensors			Execut	tion Engine	on Engine (NPU)		
	Input	Kernel	Output	NN	ТР	PPU		
VSI_NN_OP_ABS	asym-u8		asym-u8		√			
	asym-i8		asym-i8		√			
	fp32		fp32			√		
	fp16		fp16		√			
VSI_NN_OP_CLIP	asym-u8		asym-u8			√		
	asym-i8		asym-i8			✓		
	fp32		fp32			✓		
	fp16		fp16			✓		
VSI_NN_OP_EXP	asym-u8		asym-u8			√		
	asym-i8		asym-i8			√		
	fp32		fp32			✓		
	fp16		fp16			✓		
VSI_NN_OP_LOG	asym-u8		asym-u8			√		
	asym-i8		asym-i8			✓		
	fp32		fp32			✓		
	fp16		fp16			✓		
VSI_NN_OP_NEG	asym-u8		asym-u8			✓		
	asym-i8		asym-i8			✓		
	fp32		fp32			✓		
	fp16		fp16			✓		
VSI_NN_OP_MISH	asym-u8		asym-u8			√		
	asym-i8		asym-i8			✓		
	fp32		fp32			✓		
	fp16		fp16			✓		
VSI_NN_OP_SOFTMAX	asym-u8		asym-u8			✓		
	asym-i8		asym-i8			✓		
	fp32		fp32			✓		
	fp16		fp16			✓		
VSI_NN_OP_LOG_SOFTMAX	asym-u8		asym-u8			✓		
	asym-i8		asym-i8			✓		
	fp32		fp32			✓		
	fp16		fp16			√		
VSI_NN_OP_SQUARE	asym-u8		asym-u8			✓		
	asym-i8		asym-i8			✓		
	fp32		fp32			✓		

#### Table 11. OVXLIB operation support with NPU...continued

OVXLIB Operations	Tensors			Execut	tion Engine	ngine (NPU)		
	Input	Kernel	Output	NN	TP	PPU		
	fp16		fp16			1		
VSI_NN_OP_SIN	asym-u8		asym-u8			<b>√</b>		
	asym-i8		asym-i8			1		
	fp32		fp32			1		
	fp16		fp16			1		
VSI_NN_OP_LINEAR	asym-u8		asym-u8			1		
	asym-i8		asym-i8			1		
	fp32		fp32			√		
	fp16		fp16			1		
VSI_NN_OP_ERF	asym-u8		asym-u8		√	√		
	asym-i8		asym-i8		√	1		
	fp32		fp32			1		
	fp16		fp16		1	1		
Elementwise Operations		I			I			
VSI_NN_OP_ADD	asym-u8		asym-u8	$\checkmark$				
	asym-i8		asym-i8	√				
	fp32		fp32			1		
	fp16		fp16			1		
VSI_NN_OP_SUBTRACT	asym-u8		asym-u8	√				
	asym-i8		asym-i8	$\checkmark$				
	fp32		fp32			$\checkmark$		
	fp16		fp16			1		
VSI_NN_OP_MULTIPLY	asym-u8		asym-u8			1		
	asym-i8		asym-i8			1		
	fp32		fp32			1		
	fp16		fp16			✓		
VSI_NN_OP_DIVIDE	asym-u8		asym-u8			1		
	asym-i8		asym-i8			1		
	fp32		fp32			1		
	fp16		fp16			1		
VSI_NN_OP_MAXIMUN	asym-u8		asym-u8			1		
	asym-i8		asym-i8			✓		
	fp32		fp32			<b>√</b>		
	fp16		fp16			<b>√</b>		
VSI_NN_OP_MINIMUM	asym-u8		asym-u8			1		

#### Table 11. OVXLIB operation support with NPU...continued

UG10166

User guide

OVXLIB Operations	Tensors			Execut	ion Engine	ngine (NPU)		
	Input	Kernel	Output	NN	ТР	PPU		
	asym-i8		asym-i8			√		
	fp32		fp32			√		
	fp16		fp16			√		
VSI_NN_OP_POW	asym-u8		asym-u8			√		
	asym-i8		asym-i8			√		
	fp32		fp32			✓		
	fp16		fp16			√		
VSI_NN_OP_FLOORDIV	asym-u8		asym-u8			√		
	asym-i8		asym-i8			√		
	fp32		fp32			$\checkmark$		
	fp16		fp16			√		
VSI_NN_OP_MATRIXMUL	asym-u8		asym-u8			✓		
	asym-i8		asym-i8			√		
	fp32		fp32			√		
	fp16		fp16			$\checkmark$		
VSI_NN_OP_RELATIONAL_OPS	asym-u8		bool8			√		
	asym-i8		bool8			1		
	fp32		bool8			√		
	fp16		bool8			$\checkmark$		
	bool8		bool8			√		
VSI_NN_OP_LOGICAL_OPS	bool8		bool8			$\checkmark$		
VSI_NN_OP_LOGICAL_NOT	bool8		bool8			√		
VSI_NN_OP_SELECT	asym-u8		asym-u8			√		
	asym-i8		asym-i8			√		
	fp32		fp32			$\checkmark$		
	fp16		fp16			$\checkmark$		
	bool8		bool8			$\checkmark$		
VSI_NN_OP_ADDN	asym-u8		asym-u8			$\checkmark$		
	asym-i8		asym-i8			$\checkmark$		
	fp32		fp32			$\checkmark$		
	fp16		fp16			✓		
Normalization Operations		·						
VSI_NN_OP_BATCH_NORM	asym-u8		asym-u8			√		
	asym-i8		asym-i8			1		
	fp32		fp32			$\checkmark$		

#### Table 11. OVXLIB operation support with NPU...continued

OVXLIB Operations	Tensors			Execut	ion Engine	Engine (NPU)		
	Input	Kernel	Output	NN	TP	PPU		
	fp16		fp16			1		
VSI_NN_OP_LRN	asym-u8		asym-u8		<b>v</b>			
	asym-i8		asym-i8		<b>v</b>			
	fp32		fp32			1		
	fp16		fp16		<b>√</b>			
VSI_NN_OP_LRN2	asym-u8		asym-u8		<b>√</b>			
	asym-i8		asym-i8		<b>√</b>			
	fp32		fp32			$\checkmark$		
	fp16		fp16		1			
VSI_NN_OP_L2_NORMALIZE	asym-u8		asym-u8			$\checkmark$		
	asym-i8		asym-i8			$\checkmark$		
	fp32		fp32			$\checkmark$		
	fp16		fp16			$\checkmark$		
VSI_NN_OP_L2NORMALZESCALE	asym-u8		asym-u8			$\checkmark$		
	asym-i8		asym-i8			$\checkmark$		
	fp32		fp32			$\checkmark$		
	fp16		fp16			$\checkmark$		
VSI_NN_OP_LAYER_NORM	asym-u8		asym-u8			√		
	asym-i8		asym-i8			$\checkmark$		
	fp32		fp32			$\checkmark$		
	fp16		fp16			$\checkmark$		
VSI_NN_OP_INSTANCE_NORM	asym-u8		asym-u8			1		
	asym-i8		asym-i8			1		
	fp32		fp32			1		
	fp16		fp16			1		
VSI_NN_OP_BATCHNORM_SINGLE	asym-u8		asym-u8			✓		
	asym-i8		asym-i8			$\checkmark$		
	fp32		fp32			✓		
	fp16		fp16			✓		
VSI_NN_OP_MOMENTS	asym-u8		asym-u8			$\checkmark$		
	asym-i8		asym-i8			$\checkmark$		
	fp32		fp32			$\checkmark$		
	fp16		fp16			✓		
VSI_NN_OP_GROUP_NORM	asym-u8		asym-u8			<b>√</b>		
	asym-i8		asym-i8			$\checkmark$		

#### Table 11. OVXLIB operation support with NPU...continued

OVXLIB Operations	Tensors			Execut	Execution Engine (NPU		
	Input	Kernel	Output	NN	TP	PPU	
	fp32		fp32			√	
	fp16		fp16			√	
Reshape Operations		1	I				
VSI_NN_OP_EXPAND_BROADCAST	asym-u8		asym-u8			$\checkmark$	
	asym-i8		asym-i8			√	
	fp32		fp32			✓	
	fp16		fp16			√	
VSI_NN_OP_SLICE	asym-u8		asym-u8		1		
	asym-i8		asym-i8		<b>√</b>		
	fp32		fp32			$\checkmark$	
	fp16		fp16		1		
VSI_NN_OP_SPLIT	asym-u8		asym-u8		1		
	asym-i8		asym-i8		1		
	fp32		fp32			$\checkmark$	
	fp16		fp16		<b>√</b>		
VSI_NN_OP_CONCAT	asym-u8		asym-u8		1		
	asym-i8		asym-i8		1		
	fp32		fp32			1	
	fp16		fp16		1		
VSI_NN_OP_STACK	asym-u8		asym-u8		1		
	asym-i8		asym-i8		1		
	fp32		fp32			$\checkmark$	
	fp16		fp16		1		
VSI_NN_OP_UNSTACK	asym-u8		asym-u8		1		
	asym-i8		asym-i8		<b>√</b>		
	fp32		fp32			$\checkmark$	
	fp16		fp16		<b>√</b>		
VSI_NN_OP_RESHAPE	asym-u8		asym-u8		<b>√</b>		
	asym-i8		asym-i8		<b>√</b>		
	fp32		fp32			$\checkmark$	
	fp16		fp16		✓		
VSI_NN_OP_SQUEEZE	asym-u8		asym-u8		✓		
	asym-i8		asym-i8		1		
	fp32		fp32			1	
	fp16		fp16		1		

#### Table 11. OVXLIB operation support with NPU...continued

OVXLIB Operations	Tensors				Execution Engine (NPU)		
	Input	Kernel	Output	NN	ТР	PPU	
VSI_NN_OP_PERMUTE	asym-u8		asym-u8		√		
	asym-i8		asym-i8		✓		
	fp32		fp32			✓	
	fp16		fp16		✓		
VSI_NN_OP_REORG	asym-u8		asym-u8		✓		
	asym-i8		asym-i8		✓		
	fp32		fp32			√	
	fp16		fp16		✓		
VSI_NN_OP_SPACE2DEPTH	asym-u8		asym-u8		✓		
	asym-i8		asym-i8		√		
	fp32		fp32			✓	
	fp16		fp16		✓		
VSI_NN_OP_DEPTH2SPACE	asym-u8		asym-u8		√		
	asym-i8		asym-i8		√		
	fp32		fp32			✓	
	fp16		fp16		✓		
	bool8		bool8				
VSI_NN_OP_BATCH2SPACE	asym-u8		asym-u8		√		
	asym-i8		asym-i8		✓		
	fp32		fp32			✓	
	fp16		fp16		√		
VSI_NN_OP_SPACE2BATCH	asym-u8		asym-u8		✓		
	asym-i8		asym-i8		√		
	fp32		fp32			✓	
	fp16		fp16		√		
VSI_NN_OP_PAD	asym-u8		asym-u8		✓		
	asym-i8		asym-i8		✓		
	fp32		fp32			√	
	fp16		fp16		✓		
VSI_NN_OP_REVERSE	asym-u8		asym-u8		✓		
	asym-i8		asym-i8		$\checkmark$		
	fp32		fp32			$\checkmark$	
	fp16		fp16		✓		
VSI_NN_OP_STRIDED_SLICE	asym-u8		asym-u8		✓		
	asym-i8		asym-i8		✓		

Table 11. OVXLIB operation support with NPU...continued

OVXLIB Operations	Tensors				Execution Engine (NPU)		
	Input	Kernel	Output	NN	TP	PPU	
	fp32		fp32			✓	
	fp16		fp16		1		
VSI_NN_OP_CROP	asym-u8		asym-u8		1		
	asym-i8		asym-i8		1		
	fp32		fp32			√	
	fp16		fp16		1		
VSI_NN_OP_REDUCE	asym-u8		asym-u8			✓	
	asym-i8		asym-i8			√	
	fp32		fp32			√	
	fp16		fp16			√	
VSI_NN_OP_ARGMAX	asym-u8		asym-u8/int16/ int32			1	
	asym-i8		asym-u8/int16/ int32			1	
	fp32		int32			✓	
	fp16		asym-u8/int16/ int32			1	
VSI_NN_OP_ARGMIN	asym-u8		asym-u8/int16/ int32			1	
	asym-i8		asym-u8/int16/ int32			1	
	fp32		int32			$\checkmark$	
	fp16		asym-u8/int16/ int32			1	
VSI_NN_OP_SHUFFLECHANNEL	asym-u8		asym-u8		1		
	asym-i8		asym-i8		1		
	fp32		fp32			✓	
	fp16		fp16		1		
RNN Operations			·				
VSI_NN_OP_LSTMUNIT_OVXLIB	asym-u8	asym-u8	asym-u8		✓	✓	
	asym-i8	pc-sym-i8	asym-i8		✓	✓	
	fp32	fp32	fp32			✓	
	fp16	fp16	fp16		$\checkmark$	$\checkmark$	
VSI_NN_OP_LSTM_OVXLIB	asym-u8	asym-u8	asym-u8		$\checkmark$	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8		1	<b>√</b>	
	fp32	fp32	fp32			<b>√</b>	
	fp16	fp16	fp16		$\checkmark$	$\checkmark$	

#### Table 11. OVXLIB operation support with NPU...continued

OVXLIB Operations	Tensors				Execution Engine (NPU)		
	Input	Kernel	Output	NN	ТР	PPU	
VSI_NN_OP_GRUCELL_OVXLIB	asym-u8	asym-u8	asym-u8		√	1	
	asym-i8	pc-sym-i8	asym-i8		√	√	
	fp32	fp32	fp32			$\checkmark$	
	fp16	fp16	fp16		✓	√	
VSI_NN_OP_GRU_OVXLIB	asym-u8	asym-u8	asym-u8		✓	√	
	asym-i8	pc-sym-i8	asym-i8		$\checkmark$	$\checkmark$	
	fp32	fp32	fp32			$\checkmark$	
	fp16	fp16	fp16		$\checkmark$	$\checkmark$	
VSI_NN_OP_SVDF	asym-u8	asym-u8	asym-u8		√	$\checkmark$	
	asym-i8	pc-sym-i8	asym-i8		$\checkmark$	$\checkmark$	
	fp32	fp32	fp32			$\checkmark$	
	fp16	fp16	fp16		$\checkmark$	1	
Pooling Operations							
VSI_NN_OP_ROI_POOL	asym-u8		asym-u8		$\checkmark$	$\checkmark$	
	asym-i8		asym-i8		$\checkmark$	$\checkmark$	
	fp32		fp32			1	
	fp16		fp16		$\checkmark$	$\checkmark$	
VSI_NN_OP_POOLWITHARGMAX	asym-u8		asym-u8			$\checkmark$	
	asym-i8		asym-i8			$\checkmark$	
	fp32		fp32			1	
	fp16		fp16			$\checkmark$	
VSI_NN_OP_UPSAMPLE	asym-u8		asym-u8			$\checkmark$	
	asym-i8		asym-i8			$\checkmark$	
	fp32		fp32			<b>√</b>	
	fp16		fp16			$\checkmark$	
Miscellaneous Operations							
VSI_NN_OP_PROPOSAL	asym-u8		asym-u8			$\checkmark$	
	asym-i8		asym-i8			<b>√</b>	
	fp32		fp32			$\checkmark$	
	fp16		fp16			$\checkmark$	
VSI_NN_OP_VARIABLE	asym-u8		asym-u8		√		
	asym-i8		asym-i8		$\checkmark$		
	fp32		fp32			$\checkmark$	
	fp16		fp16		$\checkmark$		
VSI_NN_OP_DROPOUT	asym-u8		asym-u8			<b>√</b>	

#### Table 11. OVXLIB operation support with NPU...continued

UG10166

User guide

OVXLIB Operations	Tensors				Execution Engine (NPU)		
	Input	Kernel	Output	NN	TP	PPU	
	asym-i8		asym-i8			<b>√</b>	
	fp32		fp32			1	
	fp16		fp16			1	
VSI_NN_OP_RESIZE	asym-u8		asym-u8			1	
	asym-i8		asym-i8			1	
	fp32		fp32			<b>√</b>	
	fp16		fp16			1	
VSI_NN_OP_INTERP	asym-u8		asym-u8			√	
	asym-i8		asym-i8			√	
	fp32		fp32			√	
	fp16		fp16			1	
VSI_NN_OP_DATACONVERT	asym-u8		asym-u8		1		
	asym-i8		asym-i8		1		
	fp32		fp32			1	
	fp16		fp16		1		
VSI_NN_OP_A_TIMES_B_PLUS_C	asym-u8		asym-u8			√	
	asym-i8		asym-i8			1	
	fp32		fp32			$\checkmark$	
	fp16		fp16			1	
VSI_NN_OP_FLOOR	asym-u8		asym-u8			1	
	asym-i8		asym-i8			1	
	fp32		fp32			1	
	fp16		fp16			1	
VSI_NN_OP_EMBEDDING_	asym-u8		asym-u8			1	
LOOKUP	asym-i8		asym-i8			1	
	fp32		fp32			1	
	fp16		fp16			$\checkmark$	
VSI_NN_OP_GATHER	asym-u8		asym-u8			✓	
	asym-i8		asym-i8			✓	
	fp32		fp32			1	
	fp16		fp16			1	
VSI_NN_OP_GATHER_ND	asym-u8		asym-u8			$\checkmark$	
	asym-i8		asym-i8			<b>√</b>	
	fp32		fp32			<b>√</b>	
	fp16		fp16			$\checkmark$	

#### Table 11. OVXLIB operation support with NPU...continued

UG10166

## i.MX Machine Learning User's Guide

OVXLIB Operations	Tensors				Execution Engine (NPU)		
	Input	Kernel	Output	NN	ТР	PPU	
VSI_NN_OP_SCATTER_ND	asym-u8		asym-u8			<b>√</b>	
	asym-i8		asym-i8			<b>√</b>	
	fp32		fp32			<b>√</b>	
	fp16		fp16			<b>√</b>	
VSI_NN_OP_TILE	asym-u8		asym-u8			<b>√</b>	
	asym-i8		asym-i8			<b>√</b>	
	fp32		fp32			1	
	fp16		fp16			1	
VSI_NN_OP_RELU_KERAS	asym-u8		asym-u8		√		
	asym-i8		asym-i8		√		
	fp32		fp32			1	
	fp16		fp16		√		
VSI_NN_OP_ELTWISEMAX	asym-u8		asym-u8			1	
	asym-i8		asym-i8			1	
	fp32		fp32			1	
	fp16		fp16			1	
VSI_NN_OP_INSTANCE_NORM	asym-u8		asym-u8			1	
	asym-i8		asym-i8			1	
	fp32		fp32			1	
	fp16		fp16			1	
VSI_NN_OP_FCL2	asym-u8		asym-u8		$\checkmark$		
	asym-i8		asym-i8		√		
	fp32		fp32			1	
	fp16		fp16		√		
VSI_NN_OP_POOL	asym-u8		asym-u8	√	√		
	asym-i8		asym-i8	√	✓		
	fp32		fp32			✓	
	fp16		fp16		√		
VSI_NN_OP_SIGNAL_FRAME	asym-u8		asym-u8			1	
	asym-i8		asym-i8			1	
	fp32		fp32			✓	
	fp16		fp16			✓	
VSI_NN_OP_CONCATSHIFT	asym-u8		asym-u8			<b>√</b>	
	asym-i8		asym-i8			✓	
	fp32		fp32			1	

#### Table 11. OVXLIB operation support with NPU...continued

OVXLIB Operations	Tensors			Execut	Execution Engine (NPU)		
	Input	Kernel	Output	NN	ТР	PPU	
	fp16		fp16			✓	
VSI_NN_OP_UPSAMPLESCALE	asym-u8		asym-u8			√	
	asym-i8		asym-i8			$\checkmark$	
	fp16		fp16			$\checkmark$	
VSI_NN_OP_ROUND	asym-u8		asym-u8			$\checkmark$	
	asym-i8		asym-i8			$\checkmark$	
	fp32		fp32			$\checkmark$	
	fp16		fp16			✓	
VSI_NN_OP_CEIL	asym-u8		asym-u8			✓	
	asym-i8		asym-i8			✓	
	fp32		fp32			✓	
	fp16		fp16			✓	
VSI_NN_OP_SEQUENCE_MASK	asym-u8		asym-u8			✓	
	asym-i8		asym-i8			✓	
	fp32		fp32			√	
	fp16		fp16			√	
VSI_NN_OP_REPEAT	asym-u8		asym-u8			√	
	asym-i8		asym-i8			✓	
	fp32		fp32			✓	
	fp16		fp16			√	
VSI_NN_OP_ONE_HOT	asym-u8		asym-u8			√	
	asym-i8		asym-i8			√	
	fp32		fp32			√	
	fp16		fp16			$\checkmark$	
VSI_NN_OP_CAST	asym-u8		asym-u8			√	
	asym-i8		asym-i8			√	
	fp32		fp32			<b>√</b>	
	fp16		fp16			$\checkmark$	

#### Table 11. OVXLIB operation support with NPU...continued

## 14 Note About the Source Code in the Document

Example code shown in this document has the following copyright and BSD-3-Clause license:

Copyright 2025 NXP Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3. Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## **15 Revision History**

This table provides the revision history.

Document ID	Release date	Description
UG10166 v.LF6.12.3_1.0.0	31 March 2025	Upgraded to the 6.12.3 kernel.
UG10166 v.LF6.6.52_2.2.0	16 December 2024	Upgraded to the 6.6.52 kernel.
UG10166 v.LF6.6.36_2.1.0	30 September 2024	Upgraded to the 6.6.36 kernel.
IMXMLUG_6.6.23_2.0.0	28 June 2024	Upgraded to the 6.6.23 kernel, U-Boot v2024.04, TF-A v2.10, OP-TEE 4.2.0, Yocto 5.0 Scarthgap, and added the i.MX 91 as Alpha quality, i.MX 95 as Beta quality.
IMXMLUG v.LF6.6.3_1.0.0	29 March 2024	Upgraded to the 6.6.3 kernel, removed the i.MX 91P, and added the i.MX 95 as Alpha Quality.
IMXMLUG v.LF6.1.55_2.2.0	12/2023	Upgraded to the 6.1.55 kernel.
IMXMLUG v.LF6.1.36_2.1.0	09/2023	Upgraded to the 6.1.36 kernel.
IMXMLUG v.LF6.1.22_2.0.0	06/2023	Upgraded to the 6.1.22 kernel.
IMXMLUG v.LF6.1.1_1.0.0	03/2023	Upgraded to the 6.1.1 kernel.
IMXMLUG v.LF5.15.71_2.2.0	12/2022	Upgraded to the 5.15.71 kernel.
IMXMLUG v.LF5.15.52_2.1.0	09/2022	Upgraded to the 5.15.52 kernel, and added the i.MX 93.
IMXMLUG v.LF5.15.32_2.0.0	06/2022	Upgraded to the 5.15.32 kernel, U-Boot 2022.04, and Kirkstone Yocto.
IMXMLUG v.LF5.15.5_1.0.0	03/2022	Upgraded to the 5.15.5 kernel, Honister Yocto, and Qt6.
IMXMLUG v.LF5.10.72_2.2.0	12/2021	Upgraded the kernel to 5.10.72 and updated the BSP.
IMXMLUG v.LF5.10.52_2.1.0	09/2021	Updated for i.MX 8ULP Alpha and the kernel upgraded to 5.10.52.
IMXMLUG v.LF5.10.35_2.0.0	06/2021	Upgraded to Yocto Project Hardknott and the kernel upgraded to 5.10.35.
IMXMLUG v.L5.4.70_2.3.2	04/2021	Patch release.

Revision history

User guide

### Revision history...continued

Document ID	Release date	Description
IMXMLUG v.LF5.10.9_1.0.0	03/2021	Kernel upgrade to 5.10.9 and Machine Learning upgrades.
IMXMLUG v.L5.4.70_2.3.0	01/2021	i.MX 5.4 consolidated GA for release i.MX boards including i. MX 8M Plus and i.MX 8DXL.
IMXMLUG v.L5.4.47_2.2.0	09/2020	Initial release.

### i.MX Machine Learning User's Guide

## Legal information

### Definitions

**Draft** — A draft status on a document indicates that the content is still under internal review and subject to formal approval, which may result in modifications or additions. NXP Semiconductors does not give any representations or warranties as to the accuracy or completeness of information included in a draft version of a document and shall have no liability for the consequences of use of such information.

## Disclaimers

Limited warranty and liability — Information in this document is believed to be accurate and reliable. However, NXP Semiconductors does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information and shall have no liability for the consequences of use of such information. NXP Semiconductors takes no responsibility for the content in this document if provided by an information source outside of NXP Semiconductors.

In no event shall NXP Semiconductors be liable for any indirect, incidental, punitive, special or consequential damages (including - without limitation lost profits, lost savings, business interruption, costs related to the removal or replacement of any products or rework charges) whether or not such damages are based on tort (including negligence), warranty, breach of contract or any other legal theory.

Notwithstanding any damages that customer might incur for any reason whatsoever, NXP Semiconductors' aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms and conditions of commercial sale of NXP Semiconductors.

**Right to make changes** — NXP Semiconductors reserves the right to make changes to information published in this document, including without limitation specifications and product descriptions, at any time and without notice. This document supersedes and replaces all information supplied prior to the publication hereof.

Suitability for use — NXP Semiconductors products are not designed, authorized or warranted to be suitable for use in life support, life-critical or safety-critical systems or equipment, nor in applications where failure or malfunction of an NXP Semiconductors product can reasonably be expected to result in personal injury, death or severe property or environmental damage. NXP Semiconductors and its suppliers accept no liability for inclusion and/or use of NXP Semiconductors products in such equipment or applications and therefore such inclusion and/or use is at the customer's own risk.

**Applications** — Applications that are described herein for any of these products are for illustrative purposes only. NXP Semiconductors makes no representation or warranty that such applications will be suitable for the specified use without further testing or modification.

Customers are responsible for the design and operation of their applications and products using NXP Semiconductors products, and NXP Semiconductors accepts no liability for any assistance with applications or customer product design. It is customer's sole responsibility to determine whether the NXP Semiconductors product is suitable and fit for the customer's applications and products planned, as well as for the planned application and use of customer's third party customer(s). Customers should provide appropriate design and operating safeguards to minimize the risks associated with their applications and products.

NXP Semiconductors does not accept any liability related to any default, damage, costs or problem which is based on any weakness or default in the customer's applications or products, or the application or use by customer's third party customer(s). Customer is responsible for doing all necessary testing for the customer's applications and products using NXP Semiconductors products in order to avoid a default of the applications and the products or of the application or use by customer's third party customer(s). NXP does not accept any liability in this respect.

Terms and conditions of commercial sale — NXP Semiconductors products are sold subject to the general terms and conditions of commercial sale, as published at https://www.nxp.com/profile/terms, unless otherwise agreed in a valid written individual agreement. In case an individual agreement is concluded only the terms and conditions of the respective agreement shall apply. NXP Semiconductors hereby expressly objects to applying the customer's general terms and conditions with regard to the purchase of NXP Semiconductors products by customer.

**Export control** — This document as well as the item(s) described herein may be subject to export control regulations. Export might require a prior authorization from competent authorities.

Suitability for use in non-automotive qualified products — Unless this document expressly states that this specific NXP Semiconductors product is automotive qualified, the product is not suitable for automotive use. It is neither qualified nor tested in accordance with automotive testing or application requirements. NXP Semiconductors accepts no liability for inclusion and/or use of non-automotive qualified products in automotive equipment or applications.

In the event that customer uses the product for design-in and use in automotive applications to automotive specifications and standards, customer (a) shall use the product without NXP Semiconductors' warranty of the product for such automotive applications, use and specifications, and (b) whenever customer uses the product for automotive applications beyond NXP Semiconductors' specifications such use shall be solely at customer's own risk, and (c) customer fully indemnifies NXP Semiconductors for any liability, damages or failed product claims resulting from customer design and use of the product for automotive applications beyond NXP Semiconductors' standard warranty and NXP Semiconductors' product specifications.

**HTML publications** — An HTML version, if available, of this document is provided as a courtesy. Definitive information is contained in the applicable document in PDF format. If there is a discrepancy between the HTML document and the PDF document, the PDF document has priority.

**Translations** — A non-English (translated) version of a document, including the legal information in that document, is for reference only. The English version shall prevail in case of any discrepancy between the translated and English versions.

Security — Customer understands that all NXP products may be subject to unidentified vulnerabilities or may support established security standards or specifications with known limitations. Customer is responsible for the design and operation of its applications and products throughout their lifecycles to reduce the effect of these vulnerabilities on customer's applications and products. Customer's responsibility also extends to other open and/or proprietary technologies supported by NXP products for use in customer's applications. NXP accepts no liability for any vulnerability. Customer should regularly check security updates from NXP and follow up appropriately. Customer shall select products with security features that best meet rules, regulations, and standards of the intended application and make the ultimate design decisions regarding its products and is solely responsible for compliance with all legal, regulatory, and security related requirements concerning its products, regardless of any information or support that may be provided by NXP.

NXP has a Product Security Incident Response Team (PSIRT) (reachable at <u>PSIRT@nxp.com</u>) that manages the investigation, reporting, and solution release to security vulnerabilities of NXP products.

 $\ensuremath{\mathsf{NXP}}\xspace$  B.V. — NXP B.V. is not an operating company and it does not distribute or sell products.

## Trademarks

Notice: All referenced brands, product names, service names, and trademarks are the property of their respective owners. **NXP** — wordmark and logo are trademarks of NXP B.V.

#### i.MX Machine Learning User's Guide

Amazon Web Services, AWS, the Powered by AWS logo, and FreeRTOS — are trademarks of Amazon.com, Inc. or its affiliates.

AMBA, Arm, Arm7, Arm7TDMI, Arm9, Arm11, Artisan, big.LITTLE, Cordio, CoreLink, CoreSight, Cortex, DesignStart, DynamIQ, Jazelle, Keil, Mali, Mbed, Mbed Enabled, NEON, POP, RealView, SecurCore, Socrates, Thumb, TrustZone, ULINK, ULINK2, ULINK-ME, ULINK-PLUS, ULINKpro, µVision, Versatile — are trademarks and/or registered trademarks of Arm Limited (or its subsidiaries or affiliates) in the US and/or elsewhere. The related technology may be protected by any or all of patents, copyrights, designs and trade secrets. All rights reserved. **Cadence** — the Cadence logo, and the other Cadence marks found at <u>www.cadence.com/go/trademarks</u> are trademarks or registered trademarks of Cadence Design Systems, Inc. All rights reserved worldwide.

**eIQ** — is a trademark of NXP B.V. **IAR** — is a trademark of IAR Systems AB.

**PyTorch, the PyTorch logo and any related marks** — are trademarks of The Linux Foundation.

TensorFlow, the TensorFlow logo and any related marks — are trademarks of Google Inc.

UG10166

## i.MX Machine Learning User's Guide

## Contents

1	Software Stack Introduction	.2
2	TensorFlow Lite	. 2
2.1	TensorFlow Lite software stack	. 3
2.2	Inference backends and delegates	.4
2.2.1	Built-in kernels	.5
2.2.2	XNNPACK Delegate	.5
223	VX Delegate	5
2.2.0	Ethos-II Delegate	5
2.2.4	Neutron Delegate	5
2.2.0	GPU Delegate	6
2.2.0	Delivery package	.0
2.0	Build details	.0
2. <del>4</del> 2.5	Application development	.0
2.5	Croate CMake project which uses	
2.3.1	TapaarElow Lita	7
252	Light Vesto SDK procempiled librarias	. /
2.5.2	Using Yoclo SDK precompiled libraries	. 8
2.0		0
0.0.4	TensorFlow Lite Runtime	.9
2.6.1	TensorFlow and TensorFlow Lite Operator	•
	Set	. 9
2.6.2	Building the TensorFlow Lite Library with	-
	the Flex Delegate for I.MX Linux platforms	.9
2.6.2.1	Checking out the TensorFlow repository	10
2.6.2.2	Setting up Docker VM	10
2.6.2.3	Building the TensorFlow Lite with Flex	
	Delegate	10
2.6.3	Reducing the size of the Flex Delegate	
	library	11
2.6.4	Flex Delegate deployment on NXP i.MX	
	Linux platform	12
2.6.5	Using hardware accelerators	13
2.6.6	Flex Delegate limitations	13
2.7	Running image classification example	14
2.7.1	Running the example on the i.MX 8	
	platform hardware accelerator	14
2.7.2	Running the example on the i.MX 93	
	platform with Ethos-U	15
2.7.3	Running the example on the i.MX 9	
	platform with Neutron-S	15
2.7.4	Running the Python example	16
2.7.5	Running the example on the i.MX 95	
	platform using GPU	16
2.8	Running benchmark applications	17
2.9	Post training quantization using TensorFlow	••
2.0	Lite converter	19
2 10	TensorFlow Lite for Microcontrollers on	
2.10	Xtensa HiFi4 core	20
3	ONNX Runtime	20 21
31	ONNX Runtime software stack	22
3.2	ONNIX model test	22
J.∠ 3 3	ONNIX nerformance test	~~ 23
J.J	DyTorch	22 22
<b>→</b> // 1	Installing DyTorch	24
4.1 12	Pupping image classification example	∠ <del>4</del> ว∕I
4.Z	TVM	∠4 ว⊿
5 F 4		<b>24</b>
5.T		25

5.2	Getting started	. 25
5.2.1	Running example with RPC verification	25
5.2.2	Running example individually on device	26
5.3	How to build TVM stack on host	26
5.4	Supported models	27
6	NN Execution on Hardware Accelerators	. 28
6.1	Hardware acceleration on i.MX 8 Series	28
6.1.1	Hardware accelerator description	. 28
6.1.2	Profiling on hardware accelerators	. 29
6.1.3	Hardware accelerators warmup time	. 29
6.1.4	Switching between GPU and NPU	. 30
6.1.5	Per-tensor vs. per-channel quantization	. 30
6.2	Hardware acceleration with Ethos-U on	
	i.MX 93 platform	. 30
6.2.1	Ethos-U subsystem overview	. 31
6.2.2	Ethos-U software architecture	. 31
6.2.3	Getting started	. 32
6.2.4	Vela tool	33
6.2.4.1	Installing the Vela tool	. 34
6.2.4.2	Compiling the TFLite model	34
6.2.5	Inference with Ethos-U inference API	36
6.2.5.1	Ethos-U driver library	36
6.2.5.2	Ethos-U Linux kernel driver interface	37
6.2.5.3	Device and Buffer class	38
6.2.5.4	Network class	. 38
6.2.5.5	Inference class	. 39
6.2.5.6	How to use the inference API	40
6.2.5.7	Interpreter class	41
6.2.5.8	Interpreter Python wrapper	42
6.2.6	Inference with TensorFlow Lite	. 42
6.2.6.1	Ethos-U Delegate	. 42
6.2.6.2	Delivery package	. 42
6.2.6.3	Running image classification example	42
6.2.6.4	Hardware accelerators warmup time	. 43
6.2.6.5	Ethos-U performance enhancement with	
	memory zero-copy	43
6.2.7	Building and deploying the Ethos-U	
	firmware	. 43
6.2.7.1	Getting the source	43
6.2.7.2	Ethos-U example applications	. 44
6.2.7.3	Deploy procedure	. 45
6.2.7.4	Using the Ethos-U on Cortex-M	46
6.2.8	Memory hierarchy for Cortex-M	47
6.2.9	Supported ML operators and constraints	48
6.2.10	Profiling on hardware accelerators	. 48
6.3	NPU transition guide from i.MX 8M Plus to	51
631	Tensorflow Lite difference between i MX 8M	. 01
0.0.1	Plus and i MX 93 NPU acceleration	51
6.3.2	NPU supported operator list	.51
6.4	Hardware acceleration with eIQ Neutron	5.
	NPU on i.MX 9 series platform	51
6.4.1	Neutron-S NPU overview	. 52
6.4.2	Neutron-S software architecture	. 52
6.4.3	NPU performance tuning	. 53
7	Vision Pipeline with NNStreamer	. 53

UG10166

### i.MX Machine Learning User's Guide

7.1	Object detection pipeline example	56
7.2	NXP NNStreamer pipeline examples	57
7.3	Pipeline profiling	57
7.3.1	Enable profiling with NNShark	57
7.3.2	Adding power measurement to NNShark	59
7.3.3	Known issues and limitations	59
8	elQ Demos	59
8.1	TensorFlow Lite Demos for i.MX 93	59
8.1.1	Image classification demo	59
8.1.2	SSD object detection demo	60
8.1.3	Hand gesture detection demo	61
8.1.4	Face recognition demo	62
9	Release Notes	63
9.1	Known issues and limitations	63
9.2	Release notes for LF6.6.52_2.2.0	63
9.3	Release notes for LF6.6.36_2.1.0	64
9.4	Release notes for LF6.6.23_2.0.0	64
9.5	Release notes for LF6.6.3_1.0.0	65
9.6	Release notes for LF6.1.55_2.2.0	65
9.7	Release notes for LF6.1.36_2.1.0	66
9.8	Release notes for LF6.1.22_2.0.0	66
9.9	Release notes for LF6.1.1_1.0.0	66
9.10	Release notes for LF5.15.71_2.2.0	67
9.11	Release notes for LF5.15.52_2.1.0	67
9.12	Release notes for LF5.15.32_2.0.0	68
9.13	Release notes for LF5.15.5-1.0.0	68
9.14	Release notes for LF5.10.72-2.2.0	69
10	List of Used Variables	70
11	Neural Network API Reference	71
12	OVXLIB Operation Support with GPU	74
13	OVXLIB Operation Support with NPU	88
14	Note About the Source Code in the	
	Document	101
15	Revision History	102
	Legal information	104

Please be aware that important notices concerning this document and the product(s) described herein, have been included in section 'Legal information'.

© 2025 NXP B.V.

All rights reserved.

For more information, please visit: https://www.nxp.com

Document feedback Date of release: 31 March 2025 Document identifier: UG10166