



ESPRESSIF

ESP-DL 用户指南

发行版本 *latest*

乐鑫信息科技

2026 年 05 月 09 日

1	Introduction	3
1.1	ESP-DL 简介	3
1.2	ESP-DL 项目组织	4
2	入门指南	7
2.1	硬件要求	7
2.2	软件要求	8
2.3	快速开始	9
2.4	模型量化	9
2.5	模型部署	10
3	Tutorials	11
3.1	如何量化模型	11
3.2	AutoQuant	14
3.3	如何加载、测试和性能分析模型	18
3.4	如何进行模型推理	23
3.5	如何创建新模块（算子）	26
3.6	使用 AI Agent 自动实现算子	28
3.7	如何部署 MobileNetV2	36
3.8	如何部署 YOLO11n	49
3.9	如何部署 YOLO11n-pose	62
3.10	如何部署流式模型	72
3.11	使用 TQT 量化模型	76
4	API Reference	89
4.1	Tensor API Reference	89
4.2	Module API Reference	99
4.3	Model API Reference	103
4.4	Fbs API Reference	118



入门指南



使用教程



API Reference

1.1 ESP-DL 简介

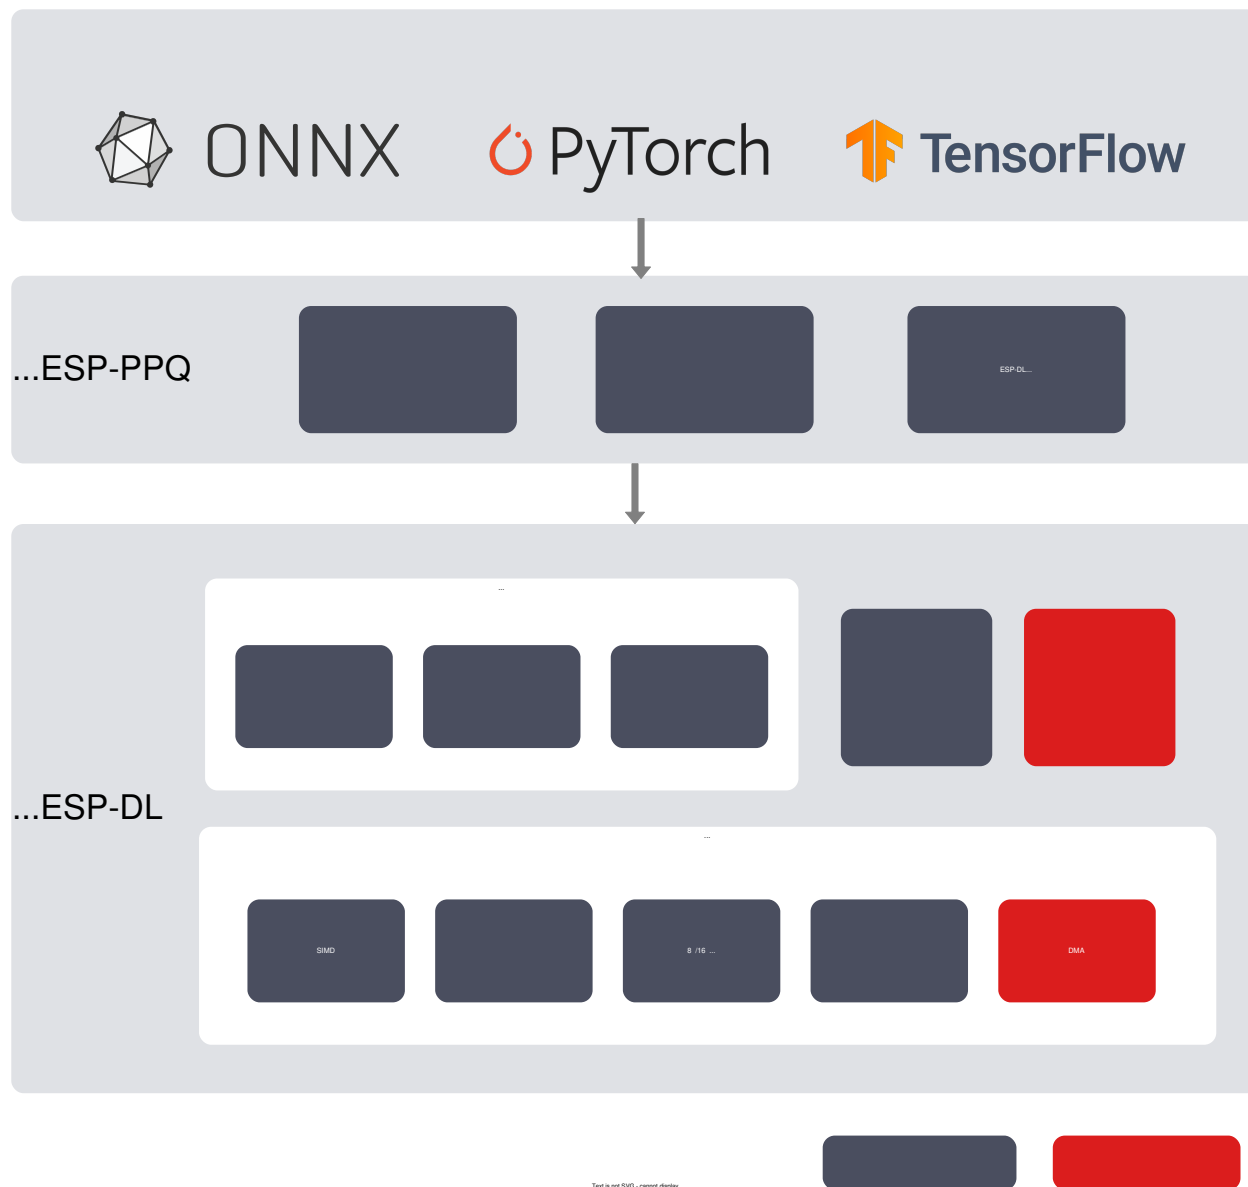
ESP-DL 是一个专为 ESP 系列芯片设计的轻量级且高效的神经网络推理框架。通过 ESP-DL，您可以轻松快速地使用乐鑫的系统级芯片 (SoC) 开发 AI 应用。

1.1.1 概述

ESP-DL 提供了加载、调试和运行 AI 模型的 API。该框架易于使用，并且可以与其他乐鑫 SDK 无缝集成。ESP-PPQ 作为 ESP-DL 的量化工具，能够量化来自 ONNX、Pytorch 和 TensorFlow 的模型，并将其导出为 ESP-DL 标准模型格式。

- **ESP-DL 标准模型格式：**该格式类似于 ONNX，但使用 FlatBuffers 而不是 Protobuf，使其更轻量级并支持零拷贝反序列化，文件后缀为 '.espdl'。
- **高效算子实现：**ESP-DL 高效地实现了常见的 AI 算子，如 Conv、Pool、Gemm、Add 和 Mul 等。目前支持的算子 [operator_support_state.md](#)
- **静态内存规划器：**内存规划器根据用户指定的内部 RAM 大小，自动将不同层分配到最佳内存位置，确保高效的整体运行速度同时最小化内存使用。
- **双核调度：**自动双核调度允许计算密集型算子充分利用双核计算能力。目前，Conv2D 和 DepthwiseConv2D 支持双核调度。
- **8bit LUT Activation：**除了 Relu, PRelu(n>1) 之外的所有激活函数，ESP-DL 默认使用 8bit LUT(Look Up Table) 方式实现，以加速推理。

ESP-DL 系统框架图如下所示：



1.2 ESP-DL 项目组织

ESP-DL 的模块化设计使其开发、维护和扩展变得高效。项目的组织结构如下：

1.2.1 dl (深度学习)

核心深度学习模块和工具，分为子模块：

- **model** 加载、管理和分配深度学习模型的内存。包含 `dl_model_base` 和 `dl_memory_manager`。
- **module** 60+ 个神经网络算子接口（卷积、池化、激活等）。文件：`dl_module_base.hpp`, `dl_module_conv.hpp`, `dl_module_pool.hpp`, `dl_module_relu.hpp` 等。
- **base** 具体的算子实现，包括对芯片（`esp32`, `esp32s3`, `esp32p4`）的 ISA 特定汇编加速。包含算子实现文件如 `dl_base_conv2d.cpp/hpp`, `dl_base_avg_pool2d.cpp/hpp` 等，以及 `isa/` 子目录中的 ISA 特定代码。
- **math** 数学操作（矩阵函数）。文件：`dl_math.hpp` 和 `dl_math_matrix.hpp`。
- **tool** 辅助功能（实用工具）。文件：`dl_tool.hpp` 和 `dl_tool.cpp`。包含 `isa/` 子目录中的 ISA 特定工具。
- **tensor** 张量类和操作。文件：`dl_tensor_base.hpp`。

1.2.2 vision (计算机视觉)

计算机视觉模块，分为子模块：

- **classification** 图像分类（模型推理）。推理：`dl_cls_base`。后处理器：`imagenet_cls_postprocessor`, `hand_gesture_cls_postprocessor`, `dl_cls_postprocessor`。
- **recognition** 特征提取（模型推理）。特征数据库管理（注册、删除、查询）。预处理器：`dl_feat_image_preprocessor`。推理：`dl_feat_base`。后处理器：`dl_feat_postprocessor`。数据库：`dl_recognition_database`。
- **image** 图像处理（调整大小、裁剪、仿射变换）。颜色转换（像素、图像）。图像预处理器（调整大小、裁剪、颜色转换、规范化、量化的管道）。图像解码/编码（JPEG/BMP）。绘制工具（点、空心矩形）。
图像处理：`dl_image_process`。颜色转换：`dl_image_color`。图像预处理器：`dl_image_preprocessor`。图像解码/编码：`dl_image_jpeg`, `dl_image_bmp`。绘制工具：`dl_image_draw`。
- **detect** 目标检测（模型推理）。推理：`dl_detect_base`。后处理器：`dl_detect_yolo11_postprocessor`, `dl_detect_espdet_postprocessor`, `dl_detect_msr_postprocessor`, `dl_detect_mnp_postprocessor`, `dl_detect_pico_postprocessor`。姿态估计：`dl_pose_yolo11_postprocessor`。

1.2.3 audio (音频处理)

音频处理模块，分为子模块：

- **common** 通用音频工具。文件：`dl_audio_common.cpp/hpp`, `dl_audio_wav.cpp/hpp`。
- **speech_features** 语音特征提取。文件：`dl_speech_features.cpp/hpp` (base class), `dl_fbank.cpp/hpp` (Filter Bank), `dl_mfcc.cpp/hpp` (MFCC), `dl_spectrogram.cpp/hpp` (Spectrogram)。

1.2.4 fbs_loader (FlatBuffers 加载器)

处理 FlatBuffers 模型:

- **include** 头文件: `fbs_loader.hpp`, `fbs_model.hpp`。
- **src** 实现: `fbs_loader.cpp`。
- **lib/** 针对不同目标的预编译库: `esp32/`, `esp32s3/`, `esp32p4/`。
- **espidl.fbs** FlatBuffers 模式文件。
- **pack_espidl_models.py** 模型打包脚本。

1.2.5 其他文件

- **CMakeLists.txt** 项目构建配置。
- **idf_component.yml** 组件元数据 (名称、版本、依赖项)。
- **README.md** 项目文档和使用说明。
- **LICENSE** 许可条款。

2.1 硬件要求

- 一块 ESP32-S3 或 ESP32-P4 开发板。推荐使用：ESP32-S3-EYE 或 ESP32-P4-Function-EV-Board
- 一台 PC (Linux 系统)

备注:

- 部分开发板目前采用 Type C 接口。请确保使用正确的线缆连接开发板!
- ESP-DL 也支持 ESP32，但其算子实现采用 C 编写，因此 ESP32 运行速度会远慢于 ESP32-S3 或 ESP32-P4。如有需要，可在项目中自行添加编译配置文件，ESP-DL 的函数接口调用方式完全一致。需要注意的是：
 - 使用 **ESP-PPQ** 量化 **ESP32** 平台模型时，需将 target 设置为 `c`。
 - 使用 **ESP-DL** 部署 **ESP32** 平台模型时，项目编译 target 则设置为 `esp32`。

2.2 软件要求

2.2.1 ESP-IDF

ESP-DL 基于 ESP-IDF 运行。有关如何获取 ESP-IDF 的详细说明, 请参阅 [ESP-IDF 编程指南](#)。

备注: 请使用 [ESP-IDF](#) 的 `release/v5.3` 或更高版本。

2.2.2 ESP-PPQ

ESP-PPQ 是基于 ppq 的量化工具, 其代码已全部开源。ESP-PPQ 在 PPQ 的基础上添加了乐鑫定制的 quantizer 和 exporter, 方便用户根据不同的芯片选择和 ESP-DL 匹配的量化规则, 并导出为 ESP-DL 可以直接加载的标准模型文件。ESP-PPQ 兼容 PPQ 所有的 API 和量化脚本。更多细节请参考 [PPQ 文档](#) 和 [视频](#)。如果您想量化自己的模型, 可以使用如下方式安装 esp-ppq:

方式一: 使用 pip 安装包

```
pip install torch torchvision torchaudio --index-url https://download.pytorch.org/whl/
↳cpu
pip install esp-ppq
```

方式二: 使用 pip 安装源码, 以便保持与 master 分支同步

```
git clone https://github.com/espressif/esp-ppq.git
cd esp-ppq
pip install torch torchvision torchaudio --index-url https://download.pytorch.org/whl/
↳cpu
pip install -e .
```

方式三: 使用 uv 安装包

```
uv pip install "esp-ppq[cpu]" --torch-backend=cpu
# GPU
# uv pip install "esp-ppq[cpu]" --torch-backend=cu124
# AMD GPU
# uv pip install "esp-ppq[cpu]" --torch-backend=rocm6.2
# Intel XPU
# uv pip install "esp-ppq[cpu]" --torch-backend=xpu
```

方式四: 使用 uv 安装源码, 以便保持与 master 分支同步

```
git clone https://github.com/espressif/esp-ppq.git
cd esp-ppq
uv pip install torch torchvision torchaudio --index-url https://download.pytorch.org/
↳whl/cpu
uv pip install -e .
```

方式五: 在 docker 中使用 esp-ppq

```
docker build -t esp-ppq:your_tag https://github.com/espressif/esp-ppq.git
```

备注:

- 示例代码中安装的是 linux pytorch cpu 版本，请根据实际情况安装对应的 pytorch。
- 如果使用 uv 安装包，仅需要更改 `--torch-backend` 参数即可，其会忽略项目中配置的 pytorch URLs 索引。

2.3 快速开始

ESP-DL 提供了一些开箱即用的 示例

2.3.1 示例编译 & 烧录

```
idf.py set-target [Soc]
idf.py flash monitor
```

使用具体的芯片替换 [Soc]，目前支持 esp32s3 和 esp32p4。示例暂未添加 esp32 的模型和编译配置文件。

2.3.2 示例配置

```
idf.py menuconfig
```

一些示例包含可配置的选项，可以在使用 `idf.py set-target` 指定芯片之后使用 `idf.py menuconfig` 进行配置。

2.3.3 故障排除

查看 ESP-IDF 文档

请参阅 [ESP-IDF DOC](#)

擦除 FLASH 和清除示例

```
idf.py eras-flash -p [PORT]
```

删除 `build/`、`sdkconfig`、`dependencies.lock`、`managed_components/` 并重试。

2.4 模型量化

首先，请参考 ESP-DL 算子支持状态 [operator_support_state.md](#)，确保您的模型中的算子已经得到支持。

ESP-DL 必须使用专有格式 `.espd1` 进行模型部署，深度学习模型需要进行量化和格式转换之后才能使用。ESP-PPQ 提供了 `espd1_quantize_onnx` 和 `espd1_quantize_torch` 两种接口以支持 ONNX 模型和 PyTorch 模型导出为 `.espd1` 模型。其他深度学习框架，如 TensorFlow, PaddlePaddle 等都需要先将模型转换为 ONNX。因此请确保您的模型可以转换为 ONNX 模型。更多详细信息，请参阅：

- [如何量化模型](#)

- 如何量化 *MobileNetV2*
- 如何量化 *YOLO11n*
- 如何量化 *YOLO11n-pose*
- 如何量化流式模型

2.5 模型部署

ESP-DL 提供了一系列 API 来快速加载和运行模型。更多详细信息，请参阅：

- 如何加载和测试模型
- 如何进行模型推理
- 如何部署流式模型

3.1 如何量化模型

ESP-DL 必须使用专有格式 `.espd1` 进行模型部署。这是一种量化模型格式，支持 8bit 和 16bit。在本教程中，我们将以 `quantize_sin_model` 为例，介绍如何使用 ESP-PPQ 量化并导出 `.espd1` 模型，量化方法为 Post Training Quantization (PTQ)。

- 准备工作
- 预训练模型
- 量化并导出 `.espd1`
 - 添加测试输入/输出
 - 量化模型推理 & 精度评估
- 高级量化方法
 - 训练后量化 (PTQ)
 - 量化感知训练 (QAT)

3.1.1 准备工作

安装 *ESP_PPQ*

3.1.2 预训练模型

```
python sin_model.py
```

执行 `sin_model.py`。该脚本会训练一个简单的 Pytorch 模型用于拟合 $[0, 2\pi]$ 范围内的 `sin` 函数。训练结束会保存相应的 `.pth` 权重, 并导出 ONNX 模型。

备注: ESP-PPQ 提供了 `espdl_quantize_onnx` 和 `espdl_quantize_torch` 两种接口以支持 ONNX 模型和 PyTorch 模型。其他深度学习框架, 如 TensorFlow, PaddlePaddle 等都需要先将模型转换为 ONNX。

- TensorFlow 转 ONNX `tf2onnx`
- TFLite 转 ONNX `tflite2onnx`
- TFLite 转 TensorFlow `tflite2tensorflow`
- PaddlePaddle 转 ONNX `paddle2onnx`

3.1.3 量化并导出 `.espdl`

参考 `quantize_torch_model.py` 和 `quantize_onnx_model.py`, 了解如何使用 `espdl_quantize_onnx` 和 `espdl_quantize_torch` 接口量化并导出 `.espdl` 模型。

执行脚本后会导出三个文件, 分别是:

- `** .espdl`: ESPDL 模型二进制文件, 可直接部署于芯片端执行推理, 并支持通过 [Netron](#) 可视化查看模型结构。
- `** .info`: ESPDL 模型文本文件, 用于调试和确定 `.espdl` 模型是否被正确导出。包含了模型结构, 量化完的模型权重, 测试输入/输出等信息。
- `** .json`: 量化信息文件, 用于量化信息的保存和加载。

备注:

1. 不同平台的 `.espdl` 模型不能混用, 推理结果会有误差。
 - ESP32 采用 Per-Tensor 量化策略, ROUND 策略为 `ROUND_HALF_UP`。
 - 使用 **ESP-PPQ** 量化 **ESP32** 平台模型时, 需将 `target` 设置为 `c`, 因为在 ESP-DL 中, 其算子实现采用 C 语言编写。
 - 使用 **ESP-DL** 部署 **ESP32** 平台模型时, 项目编译 `target` 则设置为 `esp32`。
 - ESP32S3 采用 Per-Tensor 量化策略, ROUND 策略为 `ROUND_HALF_UP`。
 - ESP32P4 的 Conv, GEMM 算子采用 Per-Channel 量化策略, 其它算子采用 Per-Tensor 量化策略, ROUND 策略为 `ROUND_HALF_EVEN`。
2. 目前 ESP-DL 使用的量化策略是对称量化 + POWER OF TWO。

添加测试输入/输出

验证模型在板端的推理结果是否正确，首先需要记录 PC 端的一组测试输入/输出。开启 api 中的 `export_test_values` 选项，就能将一组测试输入/输出固化在 `.espd1` 模型中。`input_shape` 参数和 `inputs` 参数必须指定其中的一个，`input_shape` 参数使用随机的测试输入，`inputs` 则可以指定一个特定的测试输入。`.info` 文件中可以查看测试输入/输出的值。搜索 `test inputs value` 和 `test outputs value` 查看它们。

量化模型推理 & 精度评估

`espd1_quantize_onnx` 和 `espd1_quantize_torch` API 会返回 `BaseGraph`。使用 `BaseGraph` 构建相应的 `TorchExecutor` 就可以在 PC 端使用量化模型进行推理了。

```
executor = TorchExecutor(graph=quanted_graph, device=device)
output = executor(input)
```

量化模型推理得到的输出可以用来计算各种精度指标。由于 `esp-dl` 板端推理的结果是能和 `esp-ppq` 对齐的，可以直接用该指标评估量化完模型的性能。

备注:

1. 当前 `esp-dl` 仅支持 `batch_size` 为 1，不支持多 `batch` 或者动态 `batch`。
2. `.info` 文件中的测试输入/输出，以及量化完的模型权重都是 16 字节对齐的，也就是说如果不满 16 字节，会在后面填充 0。

3.1.4 高级量化方法

如果你的模型使用默认的 8bit 量化方法无法达到满意的结果，我们也提供了如下量化方法可以进一步减少量化模型的性能损失：

训练后量化 (PTQ)

- 混合精度量化
- 层间均衡量化
- 算子分裂量化

量化感知训练 (QAT)

- *YOLO11n* 量化感知训练
- *YOLO11n-pose* 量化感知训练

3.2 AutoQuant

- 准备工作
- 快速开始
- 自定义模型量化部署
- *Stage 1*: 搜索与精度评估
- *Stage 2*: 板上延时测试
- 输出目录结构
- 故障排查

本文档介绍如何使用 AutoQuant 工具链量化并部署深度学习模型。AutoQuant 面向乐鑫芯片 (ESP32-P4 和 ESP32-S3) 设计, 你只需要定义量化策略和参数搜索空间, AutoQuant 会自动完成从量化、精度评估到端侧延时测试的整个流程, 并给出 Top-K 候选模型供你选择。相比手动反复调参, AutoQuant 可以显著降低试错成本。

3.2.1 准备工作

1. 安装 *ESP_IDF*
2. 安装 *ESP_PPQ*

3.2.2 快速开始

AutoQuant 默认基于 MobileNetV2 + ImageNet 进行量化部署。建议先运行该示例, 熟悉 AutoQuant 的整体流程, 再替换为自定义模型。

3.2.3 自定义模型量化部署

接入自定义模型时, 需要修改以下三个文件: `- config.py` —— 运行参数和搜索空间 - `calib_data_loader.py` —— 校准数据 - `eval.py` —— 精度评估

备注: 不需要修改“`main.py`”, 只需要实现 `create_calib_data_loader()` 和 `evaluation()`, AutoQuant 会自动调用它们。

1. config.py

首先修改 `runtime_config`:

```
runtime_config = {
    "onnx_path": "./quantize_mobilenetv2/models/torch/mobilenet_v2_relu.onnx",
    "export_path": "temp/model.espd1",
    "input_shape": [3, 224, 224],
    "device": "cpu",          # "cpu" 或 "cuda"
    "batchsz": 32,
    "target": "esp32p4",     # "esp32p4" 或 "esp32s3"
    "test_app_path": "../..../test_apps/autoquant_latency",
    "port": "/dev/ttyUSB0",  # ESP32-S3 使用 "/dev/ttyACM0"
    "num_of_candidates": 5,  # 保留多少个 Top-K 模型
    "primary_metric": "top1", # 见下文 eval.py 章节
}
```

`calib_dataloader` 不在 `runtime_config` 里, 它由 `main.py` 在运行时构建并注入。

同一文件中还定义了 `strategy_space` 和 `param_space`。 `strategy_space` 中的每一项配置定义某个量化模块 (如混合精度、权重均衡、TQT) 是始终启用、始终禁用, 还是在搜索时作为可选项自动决策。 `param_space` 则给出模块启用时 `AutoQuant` 需要尝试的参数取值。

2. calib_dataloader.py

将 `create_calib_dataloader()` 的函数体替换为你自己的实现。该函数必须返回一个 `PyTorch DataLoader`, 其每个 `batch` 的形状要与 `runtime_config["input_shape"]` 匹配。前/后处理 (`resize`、`normalize` 等) 是模型相关的, 必须与训练阶段保持一致。默认实现会下载一份 `ImageNet` 校准子集并应用标准的 `MobileNetV2 transforms`; 切换到其他模型时请整体替换。

3. eval.py

将 `evaluation(quant_graph)` 的函数体替换为你自己的实现。该函数接收 `PPQ` 量化后的图作为唯一参数, 必须返回一个 `dict`, 键为指标名, 值为标量。默认实现针对 `ImageNet` 分类任务返回 `{"top1": ..., "top5": ...}`。

如果评估需要使用 `GPU/CPU`, 请从 `runtime_config["device"]` 读取。

`AutoQuant` 使用 `runtime_config["primary_metric"]` 来排序候选模型, 因此需要保证 `eval.py` 返回的字典中包含该 `key`, 且 `config.py` 中的 `primary_metric` 与该 `key` 一致, 否则程序会直接报错。例如:

文件	修改内容
<code>eval.py</code>	实现 <code>evaluation(quant_graph)</code> , 例如返回 <code>{"map": map_score}</code>
<code>config.py</code>	设置 <code>"primary_metric": "map"</code>

3.2.4 Stage 1: 搜索与精度评估

在任何拥有校准数据和 GPU (或 CPU) 的主机上运行 Stage 1。无需 ESP 开发板。AutoQuant 会枚举所有 (策略, 参数) 组合, 对每个组合进行量化和精度评估, 随后按均衡规则保留 K 个结果 ($K = \text{runtime_config}["\text{num_of_candidates}"]$): 先在 `low_latency_pool` (`mixed_precision=False` 且 `horizontal_layer_split=False`) 里按 `runtime_config["primary_metric"]` 从高到低取 $K/2$ 个, 其余名额从其他候选按同一指标补齐。若其他候选不足, 则由 `low_latency_pool` 中的候选补齐。若其中一侧候选为空, 则退化为按 `primary_metric` 排序的普通 Top-K。

以下命令默认在 `esp-dl/examples/tutorial/how_to_quantize_model` 目录执行。若当前不在该目录, 请先 `cd` 到该目录。

```
python -m auto_quant.main # 全新运行; 会重置 summary/candidates
python -m auto_quant.main --resume # 从中断处继续
```

`--resume` 标志会复用已有的 `outputs/run/summary.json`, 跳过其中已经完成的 (策略, 参数) 组合。在程序崩溃、OOM 或人为中断后使用。

Stage 1 在 `outputs/run/` 下生成:

- `<index>/`` (如 ```0000/`、`0001/` ...) —— 每个实验一个子目录, 内含 `config.json` 和量化后的模型文件 (`.espd1`、`.info`、`.json`、`.native`)。
- `summary.json` —— 每个实验的追加日志 (精度、参数、hash)。
- `candidates.json` —— 当前的 Top-K 列表, 每个实验后都会重写一次, 按 Stage 1 的均衡规则维护, 从而保证不同策略族都有代表项, 且部分结果随时可用。

3.2.5 Stage 2: 板上延时测试

Stage 2 必须在物理连接 ESP 开发板的本地机器上运行 (远程服务器无法访问 USB 设备)。Stage 2 会依次烧录 Stage 1 产出的每个 Top-K 候选, 测量片上推理延时, 最后输出一份 Markdown 报告 `topk_report.md`, 由用户根据精度/延时折中自行选择最终部署哪个模型。

运行前

1. 确认 Stage 1 已完成、`outputs/run/candidates.json` 已存在。如果 Stage 1 是在远程机器上运行的, 把整个 `outputs/run/` 目录拷贝到本地 (例如使用 `rsync`)。
2. 接入 ESP32-P4 或 ESP32-S3 开发板, 并仔细确认 `config.py` 中的 `runtime_config["target"]` 和 `runtime_config["port"]` 与所连硬件一致。
3. 激活 ESP-IDF 环境。

运行

```
python -m auto_quant.test_candidates
python -m auto_quant.test_candidates --port /dev/ttyUSB1
python -m auto_quant.test_candidates --target esp32s3
```

可选的 `--port` 和 `--target` 参数会在单次调用里临时覆盖 `runtime_config` 里的对应字段, 方便在多个开发板之间切换而不改 `config.py`。

输出文件

- `outputs/run/candidates.json` ——与 Stage 1 相同的 Top-K 列表, 额外增加了 `latency_ms` 字段。每完成一个候选就更新一次, 因此即使中途崩溃或 Ctrl-C 也能保留已测的部分结果。
- `outputs/run/topk_report.md` ——Markdown 报告, 所有 Top-K 候选的 `primary_metric` 和 `latency_ms` 整理成两张四列表 (`index | folder | primary_metric | latency_ms`): 一张按 `primary_metric` 降序, 一张按 `latency_ms` 升序, 方便对照查看精度/延时折中。报告里不再重复展开每个候选的完整量化策略和参数, 请直接查 `candidates.json` (全量列表) 或 `<folder>/config.json` (单次实验快照)。失败的候选会单独列在报告末尾的 "Failed measurements" 小节。读完报告后, 把所选中行 `folder` 对应实验目录里的 `model.espd1` 烧录到你的固件中。

切换目标芯片

当 `runtime_config["target"]` 在两次运行之间发生变化 (例如从 `esp32p4` 切换到 `esp32s3`), Stage 2 会自动清理 `test app` 中的 `sdkconfig`、`build/`、`managed_components/` 以及 `dependencies.lock`, 不用手动清理。

3.2.6 输出目录结构

```
outputs/run/
summary.json           # Stage 1 的全部实验记录 (追加写入)
candidates.json        # Top-K; Stage 2 之后会附加 latency_ms 字段
topk_report.md         # Stage 2 之后供人工选型的 Markdown 报告
0000/
  config.json          # 该实验使用的量化配置
  model.espd1          # 量化后的模型 (可直接部署)
  model.info
  model.json
  model.native
0001/
...
```

3.2.7 故障排查

AssertionError: primary_metric=... not in metrics=...

`runtime_config["primary_metric"]` 设置的 key 与 `evaluation()` 实际返回的 key 不一致。对齐两者后重新运行即可。

某个候选没有打印出延时

- 检查 `runtime_config["port"]` (ESP32-P4 用 `/dev/ttyUSB0`, ESP32-S3 用 `/dev/ttyACM0`)。
- 确认串口没有被其他进程占用 (`idf.py monitor`、`minicom` 等)。
- 确认 `runtime_config["target"]` 与所连开发板一致。

idf.py flash 失败

- 芯片类型设错——正确设置 `runtime_config["target"]` 后重新运行; Stage 2 会自动清理过期的构建产物。
- 串口被占用——关闭所有打开的串口监视器。

找不到模型文件

- 检查 `config.py` 中的 `runtime_config["export_path"]`。
- 确认 Stage 1 已完成、对应的 `outputs/run/<index>/model.espd1` 存在。

3.3 如何加载、测试和性能分析模型

在本教程中，我们将介绍如何加载、测试和分析一个 `espd1` 模型。参考例程

- 准备工作
- 从 `rodata` 中加载模型
- 从 `partition` 中加载模型
- 从 `sdcard` 中加载模型
- 测试模型板端推理是否正确
- 分析模型内存使用情况
- 分析模型推理延迟
- 组合性能分析: `profile()` 方法

3.3.1 准备工作

1. 安装 `ESP_IDF`
2. 量化导出 `espd1` 模型

3.3.2 从 `rodata` 中加载模型

此方法将模型文件直接嵌入到应用程序 FLASH 的 `.rodata` 段中。这是最简单的方法，但缺点是每次应用程序代码更改时模型都会被重新烧录。

1. 在 `CMakeLists.txt` 中添加模型文件

要将 `.espd1` 模型文件嵌入到 `.rodata` 段，请在 `CMakeLists.txt` 中添加以下代码。前几行应放在 `idf_component_register()` 之前，最后一行放在 `idf_component_register()` 之后。

```
idf_build_get_property(component_targets __COMPONENT_TARGETS)
if ("___idf_espressif__esp-dl" IN_LIST component_targets)
    idf_component_get_property(espd1_dir espressif__esp-dl COMPONENT_DIR)
elseif ("___idf_esp-dl" IN_LIST component_targets)
    idf_component_get_property(espd1_dir esp-dl COMPONENT_DIR)
endif()
set(cmake_dir ${espd1_dir}/fbs_loader/cmake)
include(${cmake_dir}/utilities.cmake)
set(embed_files your_model_path/model_name.espd1)

idf_component_register(...)
```

(续下页)

(接上页)

```
target_add_aligned_binary_data(${COMPONENT_LIB} ${embed_files} BINARY)
```

2. 在程序中加载模型

包含头文件:

```
#include "dl_model_base.hpp"
```

声明模型符号并创建模型:

```
// 符号名由三部分组成: 前缀 "_binary_", 文件名 "model_espdl", 后缀 "_start"
extern const uint8_t model_espdl[] asm("_binary_model_espdl_start");

// 基本用法 - 使用默认参数加载模型
dl::Model *model = new dl::Model((const char *)model_espdl, fbs::MODEL_LOCATION_
↳ IN_FLASH_RODATA);

// 高级用法 - 自定义参数:
// - 将参数保留在 FLASH 中 (节省 PSRAM/内部 RAM, 但性能较低)
// - 限制内部 RAM 使用为 0 字节 (优先使用 PSRAM)
// - 使用贪婪内存管理器
// - 无加密密钥
// - param_copy = false (将参数保留在 FLASH 中)
// dl::Model *model = new dl::Model((const char *)model_espdl,
//                                  fbs::MODEL_LOCATION_IN_FLASH_RODATA,
//                                  0, // max_internal_size
//                                  dl::MEMORY_MANAGER_GREEDY,
//                                  nullptr, // key
//                                  false); // param_copy
```

备注: 性能与内存权衡:

- **烧录时间:** 使用从 *rodata* 中加载模型时, 模型文件嵌入在应用程序二进制文件中, 每次修改代码时都会重新烧录。对于大型模型, 这会增加烧录时间。考虑使用从 *partition* 中加载模型 或从 *sdcard* 中加载模型 来避免此问题。
- **内存 vs 性能:** `param_copy` 参数控制模型参数是否从 FLASH 复制到更快的内存 (PSRAM/内部 RAM)。设置 `param_copy=false` 可以节省 RAM, 但由于 FLASH 访问速度较慢, 会降低推理性能。仅在 RAM 极其紧张时才禁用参数复制。
- **应用程序分区大小:** 嵌入在 *.rodata* 中的大型模型可能需要增加 *partition.csv* 中的应用程序分区大小。

3.3.3 从 partition 中加载模型

此方法将模型存储在单独的 FLASH 分区中, 允许您独立于应用程序代码更新模型。

1. 在 partition.csv 中添加模型信息

创建或修改您的 `partition.csv` 文件以包含模型分区。有关分区表的详细信息, 请参阅 [ESP-IDF 分区表文档](#)。

```
# Name, Type, SubType, Offset, Size, Flags
factory, app, factory, 0x010000, 4000K,
model, data, spiffs, , 4000K,
```

- **Name:** 任何有意义的名称（包括空终止符最多 16 个字符）
- **Type:** data
- **SubType:** spiffs（模型存储必需）
- **Offset:** 留空以自动计算
- **Size:** 必须大于模型文件大小

2. 在 CMakeLists.txt 中添加模型烧录信息

```
idf_component_register(...)
set(image_file your_model_path/model_name.espd)
esptool_py_flash_to_partition(flash "model" "${image_file}")
```

esptool_py_flash_to_partition 中的第二个参数必须与 partition.csv 中的 Name 字段匹配。

3. 在程序中加载模型

包含头文件:

```
#include "dl_model_base.hpp"
```

创建模型实例:

```
// 基本用法 - 使用默认参数加载模型
dl::Model *model = new dl::Model("model", fbs::MODEL_LOCATION_IN_FLASH_PARTITION);

// 高级用法 - 将参数保留在 FLASH 中以节省 RAM
// dl::Model *model = new dl::Model("model",
//                                     fbs::MODEL_LOCATION_IN_FLASH_PARTITION,
//                                     0, // max_internal_size
//                                     dl::MEMORY_MANAGER_GREEDY,
//                                     nullptr, // key
//                                     false); // param_copy
```

第一个参数（分区标签）必须与 partition.csv 中的 Name 字段匹配。

备注：烧录优化：使用 idf.py app-flash 代替 idf.py flash，可以仅烧录应用程序分区而不重新烧录模型分区。这显著减少了开发期间的烧录时间。

3.3.4 从 sdcard 中加载模型

此方法从 SD 卡加载模型，当 FLASH 空间有限或需要频繁更新模型而无需重新烧录时非常有用。

1. 准备 SD 卡

- **格式：**SD 卡应格式化为 FAT32。如果未格式化，挂载时将自动格式化（数据会丢失）。
- **备份：**在使用 ESP-DL 之前，请始终备份 SD 卡数据。

2. 挂载 SD 卡

- **使用 BSP (板级支持包):**

在 menuconfig 中启用 CONFIG_BSP_SD_FORMAT_ON_MOUNT_FAIL 以允许自动格式化。

```
#include "bsp/esp-bsp.h"
ESP_ERROR_CHECK(bsp_sdcard_mount());
```

- **不使用 BSP:**

配置挂载选项, 设置 format_if_mount_failed = true。

```
#include "esp_vfs_fat.h"
#include "sdmmc_cmd.h"

esp_vfs_fat_sdmmc_mount_config_t mount_config = {
    .format_if_mount_failed = true,
    .max_files = 5,
    .allocation_unit_size = 16 * 1024
};
// 挂载 SD 卡 (具体实现取决于您的硬件)
```

3. 复制模型到 SD 卡

将您的 .espd1 模型文件复制到 SD 卡 (例如, 复制到根目录作为 model.espd1)。

4. 在程序中加载模型

包含头文件:

```
#include "dl_model_base.hpp"
```

- 如果不使用 BSP(Board Support Package)

```
// 挂载 sdcard.
const char *model_path = "/your_sdcard_mount_point/your_model_path/model_name.
↳espd1";
Model *model = new Model(model_path, fbs::MODEL_LOCATION_IN_SDCARD);
```

备注: 使用从 *sdcard* 中加载模型时, 模型加载过程将花费更长的时间, 因为模型数据需要从 *sdcard* 复制到 PSRAM 或者 internal RAM。如果你的 FLASH 空间紧张, 这种方法很有用。

3.3.5 测试模型板端推理是否正确

test() 方法通过将推理结果与模型文件中嵌入的基准真值进行比较, 验证模型是否产生正确的推理结果。

前提条件:

- .espd1 模型必须在 ESP-PPQ 中导出时启用 **** 测试输入和输出 **** (使用 export_test_values 选项)。
- 对于部署, 您可以导出一个没有测试数据的版本以减小模型大小。

API: esp_err_t dl::Model::test()

返回值: 如果所有测试通过则返回 ESP_OK, 否则返回 ESP_FAIL。

用法:

```
#include "dl_model_base.hpp"

// 创建模型后...
esp_err_t ret = model->test();
if (ret == ESP_OK) {
    ESP_LOGI(TAG, "模型测试通过!");
} else {
    ESP_LOGE(TAG, "模型测试失败!");
}

// 或使用便捷宏:
ESP_ERROR_CHECK(model->test());
```

工作原理:

1. 加载模型中嵌入的测试输入张量, 所以 test() 不需要外部输入
2. 通过所有模型层运行推理
3. 将每个输出与基准真值进行比较 (考虑量化误差的容差)
4. 报告每个输出的成功或失败

INT16 模型注意事项: 由于量化舍入误差, INT16 模型允许比较时有 ± 1 的差异。

3.3.6 分析模型内存使用情况

profile_memory() 方法打印跨不同内存类型 (内部 RAM、PSRAM、FLASH) 的内存使用详细明细。

API: void dl::Model::profile_memory()

用法:

```
#include "dl_model_base.hpp"

// 创建并测试模型后...
model->profile_memory();
```

输出包括:

名称	解释
fbs_model parameter	flatbuffers 模型, 包含一个子项, 模型参数 parameter。flatbuffers 模型除了模型参数之外, 还包括测试输入输出, 模型参数/变量的形状, 模型结构等信息。
parameter_copy	复制的模型参数, 当 flatbuffers 模型位于 FLASH 的时候, 默认情况下会复制到 PSRAM 或者 internal RAM 以提高推理性能。
variable	内存管理模块申请的内存, 模型输入/输出以及中间的计算结果都会使用这部分空间。
others	类成员变量所需要的空间, heap_caps_aligned_alloc / heap_caps_aligned_calloc 申请过程中对齐的额外部分 (很小)。

显示的内存类型: 每个类别的内部 RAM、PSRAM 和 FLASH 使用情况。

3.3.7 分析模型推理延迟

`profile_module()` 方法打印模型中每个模块（层）的详细延迟信息。

API: `void dl::Model::profile_module(bool sort_module_by_latency = false)`

参数: `-sort_module_by_latency`: 如果为 `true`, 模块按延迟排序 (最高优先)。如果为 `false` (默认), 模块按拓扑顺序显示。

用法:

```
// 默认: 拓扑顺序
model->profile_module();

// 按延迟排序 (最高优先)
model->profile_module(true);
```

输出包括: - 模块名称 - 模块类型 (操作类型) - 推理延迟 (微秒, 如果启用 `DL_LOG_LATENCY_UNIT` 则为周期数) - 末尾的总推理延迟

相关 API:

- `std::map<std::string, module_info> get_module_info()` - 以编程方式返回模块信息
- `void print_module_info(const std::map<std::string, module_info> &info, bool sort_module_by_latency = false)` - 从映射打印模块信息

3.3.8 组合性能分析: `profile()` 方法

`profile()` 方法结合了 `profile_memory()` 和 `profile_module()`, 进行综合分析。

API: `void dl::Model::profile(bool sort_module_by_latency = false)`

用法:

```
// 拓扑顺序的综合性能分析
model->profile();

// 按延迟排序的综合性能分析
model->profile(true);
```

这是获取内存和性能分析的最便捷方式。

3.4 如何进行模型推理

在本教程中, 我们将介绍最基本的模型推理流程。参考例程

- 准备工作
- 加载模型
- 获取模型输入/输出。
- 量化输入

- 量化单个值
- 量化 `dl::TensorBase`
- 反量化输出
 - 反量化单个值
 - 反量化 `dl::TensorBase`
- 模型推理

3.4.1 准备工作

安装 `ESP_IDF`

3.4.2 加载模型

如何加载模型

3.4.3 获取模型输入/输出。

```
std::map<std::string, dl::TensorBase *> model_inputs = model->get_inputs();
dl::TensorBase *model_input = model_inputs.begin()->second;
std::map<std::string, dl::TensorBase *> model_outputs = model->get_outputs();
dl::TensorBase *model_output = model_outputs.begin()->second;
```

可以通过 `get_inputs()` 和 `get_outputs()` api 获得输入/输出的名字和对应的 `dl::TensorBase`。更多信息, 请参阅 `dl::TensorBase` 文档。

备注: ESP-DL 的内存管理器会为每个模型的输入/中间结果/输出分配一整块的内存。由于它们共用这部分内存, 所以当模型进行推理的时候, 后面的结果会覆盖前面的结果。也就是说, `model_input` 中的数据, 在执行完模型推理之后, 可能就会被 `model_output` 或者其他中间结果所覆盖。

3.4.4 量化输入

8bit 和 16bit 量化的模型, 分别接受 `int8_t` 和 `int16_t` 类型的输入。float 类型的输入必须先根据 `exponent` 量化成对应的整数类型之后才能喂入模型。计算公式:

$$Q = \text{Clip} \left(\text{Round} \left(\frac{R}{\text{Scale}} \right), \text{MIN}, \text{MAX} \right)$$

$$\text{Scale} = 2^{\text{Exp}}$$

其中:

- R 是要量化的浮点数。
- Q 是量化后的整数值, 需要在 [MIN, MAX] 范围内进行裁剪。
- MIN 整数最小值, 8bit 时, MIN = -128, 16bit 时, MIN = -32768。
- MAX 整数最大值, 8bit 时, MAX = 127, 16bit 时, MAX = 32767。

量化单个值

```
float input_v = VALUE;
// Note that dl::quantize accepts inverse of scale as the second input, so we use DL_
↪RESCALE here.
int8_t quant_input_v = dl::quantize<int8_t>(input_v, DL_RESCALE(model_input->
↪exponent));
```

量化 dl::TensorBase

```
// assume that input_tensor already contains the float input data.
dl::TensorBase *input_tensor;
model_input->assign(input_tensor);
```

3.4.5 反量化输出

8bit 和 16bit 量化的模型，分别得到 int8_t 和 int16_t 类型的输出。必须根据 exponent 反量化之后才能得到浮点输出。计算公式：

$$R' = Q \times \text{Scale}$$

$$\text{Scale} = 2^{\text{Exp}}$$

其中：

- R' 是反量化后恢复的近似浮点值。
- Q 是量化后的整数值。

反量化单个值

```
int8_t quant_output_v = VALUE;
float output_v = dl::dequantize(quant_output_v, DL_SCALE(model_output->exponent));
```

反量化 dl::TensorBase

```
// create a TensorBase filled with 0 of shape [1, 1]
dl::TensorBase *output_tensor = new dl::TensorBase({1, 1}, nullptr, 0, dl::DATA_TYPE_
↪FLOAT);
output_tensor->assign(model_output);
```

3.4.6 模型推理

请参阅:

- 参考例程
- `void dl::Model::run(runtime_mode_t mode)`
- `void dl::Model::run(TensorBase *input, runtime_mode_t mode)`
- `void dl::Model::run(std::map<std::string, TensorBase*> &user_inputs, runtime_mode_t mode, std::map<std::string, TensorBase*> user_outputs)`

3.5 如何创建新模块（算子）

本教程将指导您在 `dl::module` 命名空间中创建一个新模块。Module 类是所有模块的基类，您将扩展这个基类来创建您的自定义模块。

备注: ESP-DL 中的模块接口应与 ONNX 对齐。

3.5.1 理解基类 Module

基类提供了几个必须在派生类中重写的虚方法。

- **方法:**
 - `dl::module::Module::Module()`: 构造函数，用于初始化模块。
 - `dl::module::Module::~~Module()`: 析构函数，用于释放资源。
 - `dl::module::Module::get_output_shape()`: 根据输入形状计算输出形状。
 - `dl::module::Module::forward()`: 运行模块，高级接口。
 - `dl::module::Module::forward_args()`: 运行模块，低级接口。
 - `dl::module::Module::deserialize()`: 从序列化信息创建模块实例。
 - `dl::module::Module::print()`: 打印模块信息。

更多信息，请参考 [Module Class Reference](#)。

3.5.2 创建新模块类

要创建一个新模块，您需要从 Module 基类派生一个新类并重写必要的方法。

示例: 创建 MyCustomModule 类

更多示例, 请参考 esp-dl/dl/module。

```

#include "module.h" // 包含定义 Module 类的头文件

namespace dl {
namespace module {

class MyCustomModule : public Module {
public:
    // 构造函数
    MyCustomModule(const char *name = "MyCustomModule",
                  module_inplace_t inplace = MODULE_NON_INPLACE,
                  quant_type_t quant_type = QUANT_TYPE_NONE)
        : Module(name, inplace, quant_type) {}

    // 析构函数
    virtual ~MyCustomModule() {}

    // 重写 get_output_shape 方法
    std::vector<std::vector<int>> get_output_shape(std::vector<std::vector<int>> &
↪input_shapes) override {
        // 实现根据输入形状计算输出形状的逻辑
        std::vector<std::vector<int>> output_shapes;
        // 示例: 假设输出形状与输入形状相同
        output_shapes.push_back(input_shapes[0]);
        return output_shapes;
    }

    // 重写 forward 方法
    void forward(std::vector<dl::TensorBase *> &tensors, runtime_mode_t mode =_
↪RUNTIME_MODE_AUTO) override {
        // 实现运行模块的逻辑
        // 示例: 对张量执行某些操作
        for (auto &tensor : tensors) {
            // 对每个张量执行某些操作
        }
    }

    // 重写 forward_args 方法
    void forward_args(void *args) override {
        // 实现低级接口的逻辑
        // 示例: 根据参数执行某些操作
    }

    // 从序列化信息反序列化模块实例
    static Module *deserialize(fbs::FbsModel *fbs_model, std::string node_name){
        // 实现反序列化模块实例的逻辑
        // 接口应与 ONNX 对齐
    }

    // 重写 print 方法
    void print() override {
        // 打印模块信息
        ESP_LOGI("MyCustomModule", "Module Name: %s, Quant type: %d", name.c_str(),_
↪quant_type);
    }
}

```

(续下页)

(接上页)

```
};

} // namespace module
} // namespace dl
```

注册 MyCustomModule 类

当您实现了 MyCustomModule 类后, 请在 dl_module_creator 中注册您的模块, 使其全局可用。

```
void register_dl_modules()
{
    if (creators.empty()) {
        ...
        this->register_module("MyCustomModule", MyCustomModule::deserialize);
    }
}
```

3.6 使用 AI Agent 自动实现算子

本文档介绍如何在各种 Coding Agent 工具 (如 Claude、Cursor、OpenCode 等) 中安装和使用 espd1-operator skill, 以便在 ESP-DL 框架中完成对神经网络算子的自动化实现。以下说明以 Linux 环境为例。

- 什么是 *espd1-operator skill*
- 依赖要求
- 安装/放置 Skill
- 快速开始示例
- Skill 触发使用
- 主要功能流程
- 故障排除
- 相关资源

3.6.1 什么是 espd1-operator skill

espd1-operator 是一个供 Coding Agent 使用的自动化开发 skill, 用于在 ESP-DL 框架中实现、测试和优化神经网络算子。当你向 Coding Agent (如 Claude、Cursor、OpenCode 等) 提出算子相关请求时, 该 skill 会指导 AI 自动完成以下工作:

Coding Agent 将自动为你:

1. **分析算子需求** - 解析 ONNX 算子规范, 确定算子类型和数据类型支持
2. **生成 esp-dl C++ 代码** - 自动创建 Module 层和 Base 层的头文件/实现文件
3. **修改 esp-ppq 量化配置** - 在量化工具中注册算子支持, 配置 layout pattern

4. 创建测试用例 - 生成 PyTorch/ONNX 测试模型, 配置测试参数
5. 执行构建与测试 - 运行 Docker 构建, 生成测试数据, 执行硬件测试
6. 验证结果对齐 - 确保 esp-dl 和 esp-ppq 的推理结果一致

Skill 的核心价值:

- 端到端自动化 - 从需求到可运行代码, Coding Agent 自动完成所有步骤
- 跨仓库协调 - 同时修改 esp-dl (C++) 和 esp-ppq (Python) 两个代码库
- 遵循最佳实践 - 自动应用 ESP-DL 的代码规范、目录结构和测试流程
- 增量式开发 - 支持新算子实现、数据类型扩展等多种场景

适用场景:

场景	示例请求
实现新算子	“实现 HardSwish 算子, 支持 int8 和 float32”
添加数据类型支持	“给 Tanh 算子添加 int16 支持”
量化支持	“为 Mod 添加量化支持”
结果对齐	“验证 LogSoftmax 在 esp-dl 和 esp-ppq 的结果是否一致”

3.6.2 依赖要求

在使用 `espd1-operator skill` 之前, 你需要提前安装以下依赖:

必需提前安装的依赖

依赖	用途	安装命令
Docker	用于构建和测试环境	官方安装指南
uv	Python 包管理器	<code>curl -LsSf https://astral.sh/uv/install.sh sh</code>
Git	版本控制	<code>apt install git (Ubuntu/Debian)</code>

由 Skill 自动处理的依赖

以下依赖 ** 无需手动安装 **, `espd1-operator skill` 会在执行过程中自动处理:

- **esp-ppq**: Python 量化工具包 - skill 会自动在 Docker 容器中以源码的形式安装
- **文档生成脚本**: `gen_ops_markdown.py` 等工具 - skill 会自动运行
- **Docker 镜像**: `espd1/idf-ppq` 镜像 - skill 会自动构建 (如不存在)
- **ESP-IDF**: 开发框架 - 已包含在 Docker 镜像中

验证依赖安装

安装完成后, 请验证以下命令可以正常执行:

```
# 检查 Docker
docker --version

# 检查 uv
uv --version

# 检查 Git
git --version
```

如果以上命令都能正常输出版本信息, 说明环境已就绪, 可以开始使用 skill。

3.6.3 安装/放置 Skill

项目结构

首先, 确认你的项目目录结构如下 (esp_dl_project_1 是项目根目录, 根目录目录名称可以任意)。以 opencode 为例, 目录结构全景图如下:

```
esp_dl_project_1/      <-- 项目根目录 (所有命令在此执行)
├── esp-dl/           # ESP-DL 主代码库
│   ├── esp-dl/      # 核心库源代码 (dl/, vision/, audio/ 等)
│   ├── examples/   # 示例程序
│   ├── test_apps/  # 测试应用
│   ├── tools/      # 工具脚本
│   └── ...
├── esp-ppq/         # 量化工具 (与 esp-dl 同级)
│   ├── esp_ppq/    # 主包源代码
│   ├── pyproject.toml # 项目配置文件
│   └── ...
└── .opencode/      # OpenCode 配置 (需要创建)
    └── skills/espdl-operator/ # skill 安装位置 (指向 esp-dl/tools/agents/
        ↪ skills/espdl-operator/)
```

注意: skill 的源代码位于 esp-dl/tools/agents/skills/espdl-operator/, 你需要将其复制或链接到 .opencode/skills/espdl-operator/ (或其他 Agent 对应的目录)。skill 文件包括 SKILL.md (主文件) 和 references/ (参考模板和检查清单)。

重要: 执行命令的位置

所有以下命令都必须在项目根目录 “esp_dl_project_1/” 下执行。

如果你不确定当前位置, 请先执行:

```
# 查看当前目录
pwd

# 应该输出类似: /home/username/workspace/esp_dl_project_1
# 或 /path/to/esp_dl_project_1

# 如果不在项目根目录, 请先跳转
cd /path/to/esp_dl_project_1
```

方法一：使用 npx（推荐 - 最简单）

为 OpenCode、Cursor、Claude Code 及其他兼容工具安装 skill 最简单的方法是使用 npx：

```
npx skills add espressif/esp-dl --skill espdl-operator
```

备注：npx 是 Node.js（通过 npm）自带的包执行器。如果你尚未安装 Node.js，请参考 [Node.js 安装指南](#) 进行安装。

执行命令后，skill 将被自动安装并在你的 Coding Agent 工具中可用。

方法二：手动安装

如果你偏好手动安装，或者你的工具不支持 npx，请根据你使用的具体工具按照以下说明操作：

备注：不同 Agent 工具的 skill 安装目录可能有所不同，请根据实际使用的工具选择对应的安装路径。

OpenCode

方法一：复制文件

```
# 确保你在项目根目录 esp_dl_project_1/
cd /path/to/esp_dl_project_1

# 创建 .opencode/skills 目录
mkdir -p .opencode/skills/espdl-operator

# 从 esp-dl/tools/agents/skills/espdl-operator 复制到 .opencode/skills/espdl-operator
cp -r esp-dl/tools/agents/skills/espdl-operator/* .opencode/skills/espdl-operator/
```

方法二：使用符号链接（推荐用于开发，保持同步更新）

```
# 确保你在项目根目录 esp_dl_project_1/
cd /path/to/esp_dl_project_1

# 创建 .opencode/skills 目录
mkdir -p .opencode/skills

# 创建符号链接（使用相对路径）
# 注意：从 .opencode/skills/espdl-operator 指向 esp-dl/tools/agents/skills/espdl-
↪operator
ln -s ../../esp-dl/tools/agents/skills/espdl-operator .opencode/skills/espdl-operator

# 验证链接是否成功
ls -la .opencode/skills/espdl-operator
# 应该显示 SKILL.md 和 references/ 目录
```

启动 OpenCode 后，系统会自动加载该 skill。

Cursor

方法一：复制文件

```
# 确保你在项目根目录 esp_dl_project_1/  
cd /path/to/esp_dl_project_1  
  
# 创建 Cursor skills 目录  
mkdir -p .cursor/skills/espdl-operator  
  
# 复制 skill 文件  
cp -r esp-dl/tools/agents/skills/espdl-operator/* .cursor/skills/espdl-operator/
```

方法二：使用符号链接

```
# 确保你在项目根目录 esp_dl_project_1/  
cd /path/to/esp_dl_project_1  
  
# 创建 .cursor/skills 目录  
mkdir -p .cursor/skills  
  
# 创建符号链接  
ln -s ../../esp-dl/tools/agents/skills/espdl-operator .cursor/skills/espdl-operator
```

Claude Desktop (Claude Code)

方法一：复制文件

```
# 确保你在项目根目录 esp_dl_project_1/  
cd /path/to/esp_dl_project_1  
  
# 创建 Claude skills 目录  
mkdir -p .claude/skills/espdl-operator  
  
# 复制 skill 文件  
cp -r esp-dl/tools/agents/skills/espdl-operator/* .claude/skills/espdl-operator/
```

方法二：使用符号链接

```
# 确保你在项目根目录 esp_dl_project_1/  
cd /path/to/esp_dl_project_1  
  
# 创建 .claude/skills 目录  
mkdir -p .claude/skills  
  
# 创建符号链接  
ln -s ../../esp-dl/tools/agents/skills/espdl-operator .claude/skills/espdl-operator
```

3.6.4 快速开始示例

假设你要实现一个新的算子 MyOp:

1. 确保 skill 已安装

```
ls -la .opencode/skills/espdl-operator/SKILL.md
```

2. 在 Agent 中提问

```
"帮我实现一个 MyOp 算子, 支持 int8、int16 和 float32"
```

3. Agent 会自动

- 加载 skill
- 按照 9 个阶段指导 Coding Agent 工作
- 生成必要的代码文件
- 运行 Docker 测试

3.6.5 Skill 触发使用

安装完成后, 你可以通过以下方式触发 espdl-operator skill:

自然语言触发

直接在对话中使用以下关键词:

中文触发词	英文触发词
“实现算子”	“implement operator”
“添加算子”	“add operator”
“量化支持”	“quantization support”
“算子对齐”	“operator alignment”
“添加新的算子”	“add a new op”

示例对话

```
用户: "帮我实现一个 Mod 算子"
Agent: [自动加载 espdl-operator skill 并开始指导]

用户: "添加 LogSoftmax 算子到 esp-dl"
Agent: [自动加载 skill 并提供实现步骤]
```

显式调用

如果自动触发未生效, 可以显式要求 Agent 使用该 skill:

```
"使用 espd1-operator skill 帮我实现 Softmax 算子"  
"根据 espd1-operator skill 的指导, 给 LogSoftmax 添加 int16 支持"
```

3.6.6 主要功能流程

该 skill 指导 Coding Agent 完成以下主要阶段:

Phase 1: 研究与分类

- 理解 ONNX 算子规范
- 确定算子类型 (Elementwise、Convolution、Pooling 等)
- 确定支持的数据类型 (int8、int16、float32)

Phase 2: 实现 esp-dl Module 层

- 创建算子模块头文件 (dl_module_<op>.hpp)
- 实现 get_output_shape() 和 forward() 方法
- 在 dl_module_creator.hpp 中注册算子

Phase 3: 实现 esp-dl Base 层

- 创建 C 参考实现 (dl_base_<op>.hpp/cpp)

Phase 4: esp-ppq 集成

- 在 EspdlQuantizer.py 中注册量化支持
- 在 espd1_typedef.py 中配置 layout pattern

Phase 5: 配置测试用例

- 添加 PyTorch/ONNX 测试模型构建器
- 在 op_cfg.toml 中配置测试参数

Phase 6: Docker 构建与测试

- 生成测试用例 (int8、int16、float32)
- 构建测试应用
- 在硬件上运行测试

Phase 7: SIMD 优化 (可选)

- 这一部分暂未支持, 后续会迭代优化

Phase 8: 算子对齐验证

- 确保 esp-dl 和 esp-ppq 推理结果一致

Phase 9: 更新文档

- 运行 gen_ops_markdown.py 更新算子支持状态文档

3.6.7 故障排除

Skill 未触发

- 确认 skill 目录位于正确的目录 (以 opencode 为例: .opencode/skills/espdl-operator)
- 尝试使用显式触发词:” 使用 espdl-operator skill…”
- 确认 Agent 工具支持 skill

Skill 流程未全部执行

- 中间有流程未执行, 显式调用 espdl-operator skill 执行, 例如:

```
# 未执行 docker 命令进行硬件烧录/测试, 则在对话中下达命令:
基于 espdl-operator skill 的说明, 对硬件进行烧录测试, 硬件已连接
```

算子实现效果不理想

该 skill 的作用是指导 AI 自动完成算子代码的编写。所以最终效果会受到两个因素影响:

1. 你使用的 AI 工具 (如 OpenCode、Cursor、Claude 等)
2. AI 模型本身的代码能力 (不同模型写代码的水平, 理解能力, 调用 tool 的能力有差异)

根据我们的测试经验, 以下组合在实现 C 语言版本的算子时, 效果较好 (因资源有限, 许多组合还未覆盖, 可自行尝试):

Coding Agent 工具	使用的模型
Cursor	Claude Opus 4.6
OpenCode	Kimi 2.5 + GLM 5 或 Claude Opus 4.6

如果发现生成的代码质量不好, 可以尝试:

- 换用更强的 AI 模型 (如 Claude Opus、GPT-4 等)
- 尝试下效果更好的 Coding Agent 工具
- 参考 SKILL.md 中的详细指导手动实现

Docker 问题

```
# 检查 Docker 是否运行
docker ps

# 重新构建镜像
cd esp-dl/tools/agents/skills/espdl-operator/assets/docker
docker build -t espdl/idf-ppq:latest .
```

权限问题

```
# 确保设备访问权限 (Linux)
sudo usermod -a -G dialout $USER
# 重新登录以生效
```

3.6.8 相关资源

- **SKILL.md**: esp-dl/tools/agents/skills/espdl-operator/SKILL.md - 完整的开发指南
- **模板**: esp-dl/tools/agents/skills/espdl-operator/references/esp-dl-templates.md
- **检查清单**: esp-dl/tools/agents/skills/espdl-operator/references/esp-ppq-checklist.md
- **esp-dl**: esp-dl/ - esp-dl 主代码库
- **esp-ppq**: esp-ppq/ - esp-ppq 量化工具 (与 esp-dl 同级目录)

3.7 如何部署 MobileNetV2

在本教程中, 我们介绍如何使用 ESP-PPQ 对预训练的 MobileNetV2 模型进行量化, 并使用 ESP-DL 部署量化后的 MobileNetV2 模型。

- 准备工作
- 模型量化
 - 预训练模型
 - 校准数据集
 - 8bit 默认配置量化

- 混合精度量化
- 层间均衡量化
- 模型部署
 - 图像分类基类
 - 前处理
 - 后处理

3.7.1 准备工作

1. 安装 *ESP_IDF*
2. 安装 *ESP_PPQ*

3.7.2 模型量化

备注:

- **ESP32P4**: 对 Conv 和 Gemm 算子的权重采用 **per-channel** 量化策略, 其余算子仍使用 **per-tensor**。
- **ESP32S3 及其他芯片**: 因指令集限制, 所有算子均统一采用 **per-tensor** 量化策略。

由于 **per-channel** 量化在细节保留上通常优于 **per-tensor**, 因此相同模型在 ESP32P4 上的量化精度往往高于 ESP32S3。

为演示不同的量化方法, 下文的示例主要以 ESP32S3 为目标平台进行说明。同时, 在 8bit 默认配置量化章节中也提供了 ESP32P4 的量化结果, 以供开发者对比参考。

量化脚本

预训练模型

从 torchvision 加载 MobileNet_v2 的预训练模型, 你也可以从 [ONNX models](#) 或 [TensorFlow models](#) 下载:

```
import torchvision
from torchvision.models.mobilenetv2 import MobileNet_V2_Weights

model = torchvision.models.mobilenet.mobilenet_v2(weights=MobileNet_V2_Weights.
↳ IMAGENET1K_V1)
```

校准数据集

校准数据集需要和你的模型输入格式一致，校准数据集需要尽可能覆盖你的模型输入的所有可能情况，以便更好地量化模型。这里以 ImageNet 数据集为例，演示如何准备校准数据集。

使用 torchvision 加载 ImageNet 数据集：

```
import torchvision.datasets as datasets
from torch.utils.data.dataset import Subset
dataset = datasets.ImageFolder(
    CALIB_DIR,
    transforms.Compose(
        [
            transforms.Resize(256),
            transforms.CenterCrop(224),
            transforms.ToTensor(),
            transforms.Normalize(
                mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225]
            ),
        ]
    ),
)
dataset = Subset(dataset, indices=[_ for _ in range(0, 1024)])
dataloader = DataLoader(
    dataset=dataset,
    batch_size=BATCH_SIZE,
    shuffle=False,
    num_workers=4,
    pin_memory=False,
    collate_fn=collate_fn1,
)
```

8bit 默认配置量化

ESP32P4 量化

量化设置

```
import torch.nn as nn

def convert_relu6_to_relu(model):
    for child_name, child in model.named_children():
        if isinstance(child, nn.ReLU6):
            setattr(model, child_name, nn.ReLU())
        else:
            convert_relu6_to_relu(child)
    return model

target="esp32p4"
num_of_bits=8
batch_size=32
quant_setting = QuantizationSettingFactory.espdL_setting() # default setting
# Replace ReLU6 with ReLU. When testing per-channel,
# using ReLU results in higher precision after quantization.
model = convert_relu6_to_relu(model)
```

量化结果

(接上页)

Layer	NOISE:SIGNAL POWER RATIO
/features/features.0/features.0.0/Conv:	0.879%
/features/features.1/conv/conv.0/conv.0.0/Conv:	0.526%
/features/features.16/conv/conv.2/Conv:	0.215%
/features/features.1/conv/conv.1/Conv:	0.165%
/features/features.15/conv/conv.2/Conv:	0.127%
/features/features.14/conv/conv.2/Conv:	0.110%
/classifier/classifier.1/Gemm:	0.061%
/features/features.4/conv/conv.2/Conv:	0.055%
/features/features.3/conv/conv.2/Conv:	0.048%
/features/features.3/conv/conv.1/conv.1.0/Conv:	0.043%
/features/features.11/conv/conv.2/Conv:	0.039%
/features/features.17/conv/conv.2/Conv:	0.037%
/features/features.2/conv/conv.2/Conv:	0.035%
/features/features.7/conv/conv.2/Conv:	0.034%
/features/features.2/conv/conv.1/conv.1.0/Conv:	0.033%
/features/features.13/conv/conv.2/Conv:	0.027%
/features/features.5/conv/conv.1/conv.1.0/Conv:	0.020%
/features/features.11/conv/conv.1/conv.1.0/Conv:	0.017%
/features/features.12/conv/conv.2/Conv:	0.015%
/features/features.5/conv/conv.2/Conv:	0.013%
/features/features.6/conv/conv.2/Conv:	0.012%
/features/features.8/conv/conv.2/Conv:	0.012%
/features/features.10/conv/conv.2/Conv:	0.011%
/features/features.17/conv/conv.0/conv.0.0/Conv:	0.011%
/features/features.9/conv/conv.2/Conv:	0.010%
/features/features.8/conv/conv.1/conv.1.0/Conv:	0.007%
/features/features.4/conv/conv.0/conv.0.0/Conv:	0.007%
/features/features.9/conv/conv.1/conv.1.0/Conv:	0.006%
/features/features.16/conv/conv.1/conv.1.0/Conv:	0.006%
/features/features.6/conv/conv.1/conv.1.0/Conv:	0.006%
/features/features.15/conv/conv.1/conv.1.0/Conv:	0.006%
/features/features.14/conv/conv.0/conv.0.0/Conv:	0.005%
/features/features.7/conv/conv.0/conv.0.0/Conv:	0.005%
/features/features.2/conv/conv.0/conv.0.0/Conv:	0.005%
/features/features.16/conv/conv.0/conv.0.0/Conv:	0.004%
/features/features.12/conv/conv.1/conv.1.0/Conv:	0.004%
/features/features.4/conv/conv.1/conv.1.0/Conv:	0.004%
/features/features.11/conv/conv.0/conv.0.0/Conv:	0.003%
/features/features.10/conv/conv.1/conv.1.0/Conv:	0.003%
/features/features.13/conv/conv.1/conv.1.0/Conv:	0.003%
/features/features.15/conv/conv.0/conv.0.0/Conv:	0.003%
/features/features.3/conv/conv.0/conv.0.0/Conv:	0.002%
/features/features.13/conv/conv.0/conv.0.0/Conv:	0.002%
/features/features.12/conv/conv.0/conv.0.0/Conv:	0.002%
/features/features.7/conv/conv.1/conv.1.0/Conv:	0.002%
/features/features.18/features.18.0/Conv:	0.002%
/features/features.5/conv/conv.0/conv.0.0/Conv:	0.001%
/features/features.14/conv/conv.1/conv.1.0/Conv:	0.001%
/features/features.9/conv/conv.0/conv.0.0/Conv:	0.001%
/features/features.8/conv/conv.0/conv.0.0/Conv:	0.001%
/features/features.10/conv/conv.0/conv.0.0/Conv:	0.001%
/features/features.6/conv/conv.0/conv.0.0/Conv:	0.001%
/features/features.17/conv/conv.1/conv.1.0/Conv:	0.000%

* Prec@1 71.150 Prec@5 89.350

(接上页)

/features/features.14/conv/conv.2/Conv:	■	0.136%
/features/features.13/conv/conv.1/conv.1.0/Conv:	■	0.105%
/features/features.6/conv/conv.1/conv.1.0/Conv:	■	0.104%
/features/features.8/conv/conv.1/conv.1.0/Conv:	■	0.083%
/features/features.5/conv/conv.1/conv.1.0/Conv:	■	0.076%
/features/features.3/conv/conv.2/Conv:	■	0.075%
/features/features.16/conv/conv.2/Conv:	■	0.073%
/features/features.13/conv/conv.0/conv.0.0/Conv:	■	0.072%
/features/features.7/conv/conv.2/Conv:	■	0.071%
/features/features.15/conv/conv.2/Conv:	■	0.066%
/features/features.4/conv/conv.2/Conv:	■	0.064%
/features/features.11/conv/conv.2/Conv:	■	0.063%
/classifier/classifier.1/Gemm:	■	0.063%
/features/features.2/conv/conv.0/conv.0.0/Conv:	■	0.054%
/features/features.13/conv/conv.2/Conv:	■	0.050%
/features/features.10/conv/conv.1/conv.1.0/Conv:	■	0.042%
/features/features.17/conv/conv.0/conv.0.0/Conv:	■	0.040%
/features/features.2/conv/conv.2/Conv:	■	0.039%
/features/features.4/conv/conv.0/conv.0.0/Conv:	■	0.034%
/features/features.17/conv/conv.2/Conv:	■	0.030%
/features/features.14/conv/conv.0/conv.0.0/Conv:	■	0.025%
/features/features.16/conv/conv.0/conv.0.0/Conv:	■	0.024%
/features/features.10/conv/conv.2/Conv:	■	0.022%
/features/features.11/conv/conv.0/conv.0.0/Conv:	■	0.021%
/features/features.9/conv/conv.2/Conv:	■	0.021%
/features/features.14/conv/conv.1/conv.1.0/Conv:	■	0.020%
/features/features.7/conv/conv.1/conv.1.0/Conv:	■	0.020%
/features/features.5/conv/conv.2/Conv:	■	0.019%
/features/features.1/conv/conv.0/conv.0.0/Conv:	■	0.017%
/features/features.12/conv/conv.2/Conv:	■	0.017%
/features/features.6/conv/conv.2/Conv:	■	0.014%
/features/features.7/conv/conv.0/conv.0.0/Conv:	■	0.014%
/features/features.3/conv/conv.0/conv.0.0/Conv:	■	0.013%
/features/features.8/conv/conv.2/Conv:	■	0.012%
/features/features.12/conv/conv.0/conv.0.0/Conv:	■	0.009%
/features/features.15/conv/conv.0/conv.0.0/Conv:	■	0.008%
/features/features.5/conv/conv.0/conv.0.0/Conv:	■	0.006%
/features/features.6/conv/conv.0/conv.0.0/Conv:	■	0.005%
/features/features.9/conv/conv.0/conv.0.0/Conv:	■	0.003%
/features/features.18/features.18.0/Conv:	■	0.002%
/features/features.10/conv/conv.0/conv.0.0/Conv:	■	0.002%
/features/features.8/conv/conv.0/conv.0.0/Conv:	■	0.002%

* Prec@1 69.225 Prec@5 88.700

量化误差分析

将之前误差最大的层替换为 16 位量化后, 可以观察到模型准确度明显提升, 量化后的 top1 准确率为 69.225%, 和 float 模型的准确率 (71.878%) 比较接近。该模型的最后一层 /classifier/classifier.1/Gemm 的累计误差为 10.289%。

(接上页)

```

/features/features.15/conv/conv.1/conv.1.0/Conv: | 0.025%
/features/features.10/conv/conv.2/Conv: | 0.022%
/features/features.14/conv/conv.0/conv.0.0/Conv: | 0.020%
/features/features.16/conv/conv.1/conv.1.0/Conv: | 0.019%
/features/features.4/conv/conv.1/conv.1.0/Conv: | 0.018%
/features/features.13/conv/conv.1/conv.1.0/Conv: | 0.014%
/features/features.13/conv/conv.0/conv.0.0/Conv: | 0.014%
/features/features.16/conv/conv.0/conv.0.0/Conv: | 0.012%
/features/features.17/conv/conv.0/conv.0.0/Conv: | 0.012%
/features/features.12/conv/conv.0/conv.0.0/Conv: | 0.011%
/features/features.12/conv/conv.1/conv.1.0/Conv: | 0.011%
/features/features.9/conv/conv.2/Conv: | 0.009%
/features/features.8/conv/conv.2/Conv: | 0.009%
/features/features.2/conv/conv.0/conv.0.0/Conv: | 0.008%
/features/features.7/conv/conv.0/conv.0.0/Conv: | 0.007%
/features/features.10/conv/conv.1/conv.1.0/Conv: | 0.006%
/features/features.10/conv/conv.0/conv.0.0/Conv: | 0.006%
/features/features.15/conv/conv.0/conv.0.0/Conv: | 0.006%
/features/features.3/conv/conv.0/conv.0.0/Conv: | 0.004%
/features/features.6/conv/conv.0/conv.0.0/Conv: | 0.004%
/features/features.11/conv/conv.0/conv.0.0/Conv: | 0.003%
/features/features.5/conv/conv.0/conv.0.0/Conv: | 0.003%
/features/features.9/conv/conv.1/conv.1.0/Conv: | 0.003%
/features/features.14/conv/conv.1/conv.1.0/Conv: | 0.002%
/features/features.7/conv/conv.1/conv.1.0/Conv: | 0.002%
/features/features.17/conv/conv.1/conv.1.0/Conv: | 0.002%
/features/features.8/conv/conv.0/conv.0.0/Conv: | 0.001%
/features/features.9/conv/conv.0/conv.0.0/Conv: | 0.001%
/features/features.18/features.18.0/Conv: | 0.001%

```

* Prec@1 69.800 Prec@5 88.400

量化误差分析

注意到对 8bit 量化应用层间均衡有助于降低量化损失，量化后的 top1 准确率为 69.80%，和 float 模型的准确率 (71.878%) 更加接近，比混合精度量化的量化精度更高。

备注： 如果想进一步降低量化误差，可以尝试使用 QAT (Auantization Aware Training)。具体方法请参考 [PPQ QAT example](#)。

3.7.3 模型部署

参考示例

图像分类基类

- dl_cls_base.hpp
- dl_cls_base.cpp

前处理

ImagePreprocessor 类中封装了常用的图像前处理流程, 包括 color conversion, crop, resize, normalization, quantize。

- dl_image_preprocessor.hpp
- dl_image_preprocessor.cpp

后处理

- dl_cls_postprocessor.hpp
- dl_cls_postprocessor.cpp
- imagenet_cls_postprocessor.hpp
- imagenet_cls_postprocessor.cpp

3.8 如何部署 YOLO11n

在本教程中, 我们介绍如何使用 ESP-PPQ 对预训练的 YOLO11n 模型进行量化, 并使用 ESP-DL 部署量化后的 YOLO11n 模型。

- 准备工作
- 模型量化
 - 预训练模型
 - 校准数据集
 - 8bit 默认配置量化
 - 混合精度 + 算子分裂量化
 - 量化感知训练
- 模型部署
 - 目标检测基类
 - 前处理
 - 后处理

(接上页)

/model.23/cv3.1/cv3.1.1/cv3.1.1.1/conv/Conv:	██████████	18.913%
/model.22/m.0/m/m.1/cv2/conv/Conv:	██████████	18.645%
/model.22/cv2/conv/Conv:	██████████	18.628%
/model.23/cv2.1/cv2.1.1/conv/Conv:	██████████	17.980%
/model.8/m.0/cv2/conv/Conv:	██████████	16.247%
/model.23/cv2.0/cv2.0.1/conv/Conv:	██████████	15.602%
/model.10/m/m.0/attn/qkv/conv/Conv:	██████████	14.666%
/model.10/m/m.0/attn/pe/conv/Conv:	██████████	14.556%
/model.23/cv2.1/cv2.1.0/conv/Conv:	██████████	14.302%
/model.22/cv1/conv/Conv:	██████████	13.921%
/model.10/m/m.0/attn/MatMul_1:	██████████	13.905%
/model.10/cv1/conv/Conv:	██████████	13.494%
/model.23/cv3.1/cv3.1.0/cv3.1.0.1/conv/Conv:	██████████	11.800%
/model.19/m.0/cv2/conv/Conv:	██████████	11.515%
/model.22/m.0/m/m.0/cv2/conv/Conv:	██████████	11.286%
/model.20/conv/Conv:	██████████	10.930%
/model.13/m.0/cv2/conv/Conv:	██████████	10.882%
/model.23/cv3.2/cv3.2.1/cv3.2.1.0/conv/Conv:	██████████	10.692%
/model.23/cv2.2/cv2.2.2/Conv:	██████████	10.113%
/model.10/cv2/conv/Conv:	██████████	9.720%
/model.8/cv2/conv/Conv:	██████████	9.598%
/model.8/m.0/cv1/conv/Conv:	██████████	9.470%
/model.19/cv2/conv/Conv:	██████████	9.314%
/model.22/m.0/m/m.0/cv1/conv/Conv:	██████████	9.068%
/model.23/cv3.0/cv3.0.0/cv3.0.0.1/conv/Conv:	██████████	9.065%
/model.8/cv1/conv/Conv:	██████████	9.051%
/model.8/m.0/cv3/conv/Conv:	██████████	9.044%
/model.6/m.0/cv2/conv/Conv:	██████████	8.811%
/model.22/m.0/m/m.1/cv1/conv/Conv:	██████████	8.781%
/model.13/cv2/conv/Conv:	██████████	8.687%
/model.8/m.0/m/m.0/cv1/conv/Conv:	██████████	8.503%
/model.8/m.0/m/m.0/cv2/conv/Conv:	██████████	8.470%
/model.19/cv1/conv/Conv:	██████████	8.199%
/model.10/m/m.0/attn/MatMul:	██████████	8.117%
/model.8/m.0/m/m.1/cv1/conv/Conv:	██████████	7.964%
/model.13/cv1/conv/Conv:	██████████	7.734%
/model.19/m.0/cv1/conv/Conv:	██████████	7.661%
/model.22/m.0/cv1/conv/Conv:	██████████	7.490%
/model.13/m.0/cv1/conv/Conv:	██████████	7.162%
/model.8/m.0/m/m.1/cv2/conv/Conv:	██████████	7.145%
/model.23/cv2.0/cv2.0.0/conv/Conv:	██████████	7.041%
/model.23/cv2.1/cv2.1.2/Conv:	██████████	6.917%
/model.23/cv2.0/cv2.0.2/Conv:	██████████	6.778%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:	██████████	6.641%
/model.17/conv/Conv:	██████████	6.125%
/model.16/m.0/cv2/conv/Conv:	██████████	5.937%
/model.6/cv2/conv/Conv:	██████████	5.838%
/model.6/m.0/cv3/conv/Conv:	██████████	5.832%
/model.6/cv1/conv/Conv:	██████████	5.688%
/model.7/conv/Conv:	██████████	5.612%
/model.9/cv2/conv/Conv:	██████████	5.367%
/model.10/m/m.0/ffn/ffn.0/conv/Conv:	██████████	5.158%
/model.6/m.0/m/m.0/cv1/conv/Conv:	██████████	5.143%
/model.16/m.0/cv1/conv/Conv:	██████████	5.137%
/model.23/cv3.1/cv3.1.0/cv3.1.0.0/conv/Conv:	██████████	5.087%
/model.16/cv2/conv/Conv:	██████████	4.989%
/model.2/cv2/conv/Conv:	██████████	4.547%

(续下页)

(接上页)

/model.6/cv2/conv/Conv:	■	0.011%
/model.13/m.0/cv2/conv/Conv:	■	0.010%
/model.5/conv/Conv:		0.010%
/model.19/m.0/cv2/conv/Conv:		0.009%
/model.6/m.0/m.m.1/cv1/conv/Conv:		0.009%
/model.23/cv3.0/cv3.0.0/cv3.0.0.1/conv/Conv:		0.008%
/model.23/cv2.2/cv2.2.0/conv/Conv:		0.008%
/model.23/cv2.1/cv2.1.1/conv/Conv:		0.008%
/model.9/cv1/conv/Conv:		0.008%
/model.23/cv2.0/cv2.0.1/conv/Conv:		0.007%
/model.16/m.0/cv1/conv/Conv:		0.007%
/model.17/conv/Conv:		0.007%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:		0.007%
/model.10/m.m.0/ffn/ffn.1/conv/Conv:		0.007%
/model.23/cv2.0/cv2.0.2/Conv:		0.006%
/model.8/m.0/cv1/conv/Conv:		0.006%
/model.23/cv2.2/cv2.2.2/Conv:		0.005%
/model.23/cv2.1/cv2.1.2/Conv:		0.005%
/model.22/m.0/cv3/conv/Conv:		0.005%
/model.23/cv3.1/cv3.1.0/cv3.1.0.1/conv/Conv:		0.005%
/model.7/conv/Conv:		0.005%
/model.8/cv2/conv/Conv:		0.004%
/model.22/cv2/conv/Conv:		0.004%
/model.6/m.0/cv3/conv/Conv:		0.004%
/model.10/m.m.0/ffn/ffn.0/conv/Conv:		0.004%
/model.8/m.0/m.m.1/cv2/conv/Conv:		0.004%
/model.22/m.0/m.m.1/cv1/conv/Conv:		0.004%
/model.8/m.0/m.m.1/cv1/conv/Conv:		0.004%
/model.23/cv3.1/cv3.1.1/cv3.1.1.1/conv/Conv:		0.003%
/model.10/m.m.0/attn/proj/conv/Conv:		0.003%
/model.22/m.0/m.m.0/cv2/conv/Conv:		0.003%
/model.22/m.0/cv1/conv/Conv:		0.003%
/model.8/m.0/cv3/conv/Conv:		0.003%
/model.6/m.0/m.m.0/cv1/conv/Conv:		0.003%
/model.23/cv3.0/cv3.0.0/cv3.0.0.0/conv/Conv:		0.003%
/model.23/cv3.2/cv3.2.1/cv3.2.1.0/conv/Conv:		0.002%
/model.6/m.0/m.m.1/cv2/conv/Conv:		0.002%
/model.8/m.0/m.m.0/cv2/conv/Conv:		0.002%
/model.23/cv3.2/cv3.2.1/cv3.2.1.1/conv/Conv:		0.002%
/model.10/m.m.0/attn/MatMul_1:		0.002%
/model.22/m.0/m.m.1/cv2/conv/Conv:		0.001%
/model.6/m.0/m.m.0/cv2/conv/Conv:		0.001%
/model.23/cv3.0/cv3.0.1/cv3.0.1.0/conv/Conv:		0.001%
/model.8/m.0/m.m.0/cv1/conv/Conv:		0.001%
/model.23/cv3.2/cv3.2.0/cv3.2.0.1/conv/Conv:		0.001%
/model.23/cv3.0/cv3.0.1/cv3.0.1.1/conv/Conv:		0.001%
/model.6/m.0/cv1/conv/Conv:		0.001%
/model.23/cv3.2/cv3.2.2/Conv:		0.001%
/model.20/conv/Conv:		0.001%
/model.23/cv3.1/cv3.1.2/Conv:		0.001%
/model.23/cv3.2/cv3.2.0/cv3.2.0.0/conv/Conv:		0.001%
/model.6/m.0/cv2/conv/Conv:		0.001%
/model.23/cv3.0/cv3.0.2/Conv:		0.000%
/model.10/m.m.0/attn/MatMul:		0.000%
/model.23/cv3.1/cv3.1.0/cv3.1.0.0/conv/Conv:		0.000%
/model.8/m.0/cv2/conv/Conv:		0.000%
/model.22/m.0/cv2/conv/Conv:		0.000%

(接上页)

/model.23/cv3.2/cv3.2.1/cv3.2.1.1/conv/Conv:	██████████	13.683%
/model.22/cv2/conv/Conv:	██████████	13.654%
/model.22/m.0/cv2/conv/Conv:	██████████	13.514%
/model.23/cv3.1/cv3.1.1/cv3.1.1.1/conv/Conv:	██████████	12.885%
/model.23/cv2.1/cv2.1.1/conv/Conv:	██████████	10.865%
/model.23/cv2.0/cv2.0.1/conv/Conv:	██████████	9.875%
/model.23/cv2.1/cv2.1.0/conv/Conv:	██████████	9.658%
/model.22/cv1/conv/Conv:	██████████	8.917%
/model.10/m/m.0/attn/MatMul_1:	██████████	8.368%
/model.23/cv2.2/cv2.2.2/Conv:	██████████	8.156%
/model.22/m.0/m/m.0/cv2/conv/Conv:	██████████	8.056%
/model.10/m/m.0/attn/qkv/conv/Conv:	██████████	7.948%
/model.23/cv3.1/cv3.1.0/cv3.1.0.1/conv/Conv:	██████████	7.824%
/model.13/m.0/cv2/conv/Conv:	██████████	7.504%
/model.19/m.0/cv2/conv/Conv:	██████████	7.290%
/model.20/conv/Conv:	██████████	6.986%
/model.10/m/m.0/attn/pe/conv/Conv:	██████████	6.926%
/model.23/cv3.0/cv3.0.0/cv3.0.0.1/conv/Conv:	██████████	6.771%
/model.23/cv3.2/cv3.2.1/cv3.2.1.0/conv/Conv:	██████████	6.756%
/model.22/m.0/m/m.1/cv1/conv/Conv:	██████████	6.465%
/model.22/m.0/m/m.0/cv1/conv/Conv:	██████████	6.274%
/model.19/cv2/conv/Conv:	██████████	6.116%
/model.10/cv1/conv/Conv:	██████████	5.868%
/model.13/cv2/conv/Conv:	██████████	5.815%
/model.10/cv2/conv/Conv:	██████████	5.664%
/model.19/cv1/conv/Conv:	██████████	5.178%
/model.8/m.0/cv2/conv/Conv:	██████████	4.970%
/model.19/m.0/cv1/conv/Conv:	██████████	4.919%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:	██████████	4.864%
/model.22/m.0/cv1/conv/Conv:	██████████	4.844%
/model.10/m/m.0/attn/MatMul:	██████████	4.650%
/model.13/cv1/conv/Conv:	██████████	4.564%
/model.23/cv2.0/cv2.0.0/conv/Conv:	██████████	4.389%
/model.13/m.0/cv1/conv/Conv:	██████████	4.243%
/model.23/cv2.0/cv2.0.2/Conv:	██████████	4.232%
/model.23/cv2.1/cv2.1.2/Conv:	██████████	4.222%
/model.6/m.0/cv2/conv/Conv:	██████████	4.023%
/model.17/conv/Conv:	██████████	3.754%
/model.16/m.0/cv2/conv/Conv:	██████████	3.511%
/model.8/m.0/cv1/conv/Conv:	██████████	3.277%
/model.16/m.0/cv1/conv/Conv:	██████████	3.158%
/model.23/cv3.0/cv3.0.1/cv3.0.1.0/conv/Conv:	██████████	3.155%
/model.23/cv3.1/cv3.1.0/cv3.1.0.0/conv/Conv:	██████████	3.152%
/model.8/cv2/conv/Conv:	██████████	3.119%
/model.8/m.0/m/m.1/cv1/conv/Conv:	██████████	3.106%
/model.8/m.0/cv3/conv/Conv:	██████████	3.083%
/model.6/m.0/cv3/conv/Conv:	██████████	3.068%
/model.8/cv1/conv/Conv:	██████████	3.035%
/model.16/cv2/conv/Conv:	██████████	3.002%
/model.2/cv2/conv/Conv:	██████████	2.992%
/model.8/m.0/m/m.0/cv2/conv/Conv:	██████████	2.971%
/model.6/cv1/conv/Conv:	██████████	2.819%
/model.8/m.0/m/m.0/cv1/conv/Conv:	██████████	2.809%
/model.10/m/m.0/ffn/ffn.0/conv/Conv:	██████████	2.760%
/model.2/cv1/conv/Conv:	██████████	2.683%
/model.6/cv2/conv/Conv:	██████████	2.630%
/model.8/m.0/m/m.1/cv2/conv/Conv:	██████████	2.615%

(续下页)

(接上页)

/model.13/m.0/cv1/conv/Conv:	█	0.011%
/model.8/cv1/conv/Conv:	█	0.010%
/model.13/m.0/cv2/conv/Conv:	█	0.010%
/model.5/conv/Conv:	█	0.010%
/model.6/m.0/m/m.1/cv1/conv/Conv:	█	0.009%
/model.23/cv3.0/cv3.0.0/cv3.0.0.1/conv/Conv:	█	0.008%
/model.23/cv2.2/cv2.2.0/conv/Conv:	█	0.008%
/model.23/cv2.1/cv2.1.1/conv/Conv:	█	0.008%
/model.19/m.0/cv2/conv/Conv:	█	0.008%
/model.8/cv2/conv/Conv:	█	0.008%
/model.9/cv1/conv/Conv:	█	0.008%
/model.23/cv2.0/cv2.0.1/conv/Conv:	█	0.007%
/model.16/m.0/cv1/conv/Conv:	█	0.007%
/model.17/conv/Conv:	█	0.007%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:	█	0.007%
/model.10/m/m.0/ffn/ffn.1/conv/Conv:	█	0.007%
/model.22/m.0/cv1/conv/Conv:		0.006%
/model.10/cv2/conv/Conv:		0.006%
/model.23/cv2.0/cv2.0.2/Conv:		0.006%
/model.23/cv2.2/cv2.2.2/Conv:		0.005%
/model.23/cv2.1/cv2.1.2/Conv:		0.005%
/model.22/m.0/cv3/conv/Conv:		0.005%
/model.23/cv3.1/cv3.1.0/cv3.1.0.1/conv/Conv:		0.005%
/model.22/cv2/conv/Conv:		0.005%
/model.7/conv/Conv:		0.004%
/model.6/m.0/cv3/conv/Conv:		0.004%
/model.10/m/m.0/ffn/ffn.0/conv/Conv:		0.004%
/model.8/m.0/m/m.1/cv2/conv/Conv:		0.004%
/model.22/m.0/m/m.1/cv1/conv/Conv:		0.004%
/model.8/m.0/m/m.1/cv1/conv/Conv:		0.004%
/model.23/cv3.1/cv3.1.1/cv3.1.1.1/conv/Conv:		0.003%
/model.8/m.0/cv1/conv/Conv:		0.003%
/model.10/m/m.0/attn/proj/conv/Conv:		0.003%
/model.22/m.0/m/m.0/cv2/conv/Conv:		0.003%
PPQ_Operation_2:		0.003%
/model.8/m.0/cv3/conv/Conv:		0.003%
/model.6/m.0/m/m.0/cv1/conv/Conv:		0.003%
/model.23/cv3.2/cv3.2.1/cv3.2.1.0/conv/Conv:		0.002%
/model.6/m.0/m/m.1/cv2/conv/Conv:		0.002%
/model.8/m.0/m/m.0/cv2/conv/Conv:		0.002%
/model.23/cv3.0/cv3.0.0/cv3.0.0.0/conv/Conv:		0.002%
/model.23/cv3.2/cv3.2.1/cv3.2.1.1/conv/Conv:		0.002%
/model.10/m/m.0/attn/MatMul_1:		0.002%
/model.22/m.0/m/m.1/cv2/conv/Conv:		0.001%
/model.6/m.0/m/m.0/cv2/conv/Conv:		0.001%
/model.8/m.0/m/m.0/cv1/conv/Conv:		0.001%
/model.23/cv3.0/cv3.0.1/cv3.0.1.0/conv/Conv:		0.001%
/model.23/cv3.2/cv3.2.0/cv3.2.0.1/conv/Conv:		0.001%
/model.2/cv2/conv/Conv:		0.001%
/model.23/cv3.0/cv3.0.1/cv3.0.1.1/conv/Conv:		0.001%
/model.6/m.0/cv1/conv/Conv:		0.001%
/model.23/cv3.2/cv3.2.2/Conv:		0.001%
/model.20/conv/Conv:		0.001%
/model.23/cv3.1/cv3.1.2/Conv:		0.001%
/model.23/cv3.2/cv3.2.0/cv3.2.0.0/conv/Conv:		0.001%
/model.6/m.0/cv2/conv/Conv:		0.001%
/model.23/cv3.0/cv3.0.2/Conv:		0.000%

(续下页)

(接上页)

/model.8/cv2/conv/Conv:	██████	7.952%
/model.8/m.0/cv1/conv/Conv:	██████	7.697%
/model.13/cv2/conv/Conv:	██████	7.557%
/model.19/cv2/conv/Conv:	██████	7.443%
/model.10/cv2/conv/Conv:	██████	7.403%
/model.6/m.0/cv2/conv/Conv:	██████	7.099%
/model.8/cv1/conv/Conv:	██████	6.996%
/model.19/cv1/conv/Conv:	██████	6.912%
/model.8/m.0/m.0/cv1/conv/Conv:	██████	6.908%
/model.8/m.0/cv3/conv/Conv:	██████	6.755%
/model.23/cv3.2/cv3.2.1/cv3.2.1.0/conv/Conv:	██████	6.746%
/model.8/m.0/m.0/cv2/conv/Conv:	██████	6.743%
/model.8/m.0/m.1/cv1/conv/Conv:	██████	6.638%
/model.13/cv1/conv/Conv:	██████	6.361%
/model.2/m.0/cv2/conv/Conv:	██████	6.274%
/model.13/m.0/cv1/conv/Conv:	██████	6.261%
/model.19/m.0/cv1/conv/Conv:	██████	6.191%
/model.22/m.0/m.0/cv1/conv/Conv:	██████	6.036%
/model.23/cv2.2/cv2.2.2/Conv:	██████	5.999%
/model.22/m.0/m.1/cv1/conv/Conv:	██████	5.899%
/model.23/cv2.0/cv2.0.0/conv/Conv:	██████	5.618%
/model.8/m.0/m.1/cv2/conv/Conv:	██████	5.560%
/model.22/m.0/cv1/conv/Conv:	██████	5.336%
/model.16/m.0/cv2/conv/Conv:	██████	5.316%
/model.17/conv/Conv:	██████	5.113%
/model.6/m.0/cv3/conv/Conv:	██████	5.103%
/model.16/m.0/cv1/conv/Conv:	██████	5.101%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:	██████	5.052%
/model.2/cv2/conv/Conv:	██████	5.003%
/model.6/cv2/conv/Conv:	██████	4.968%
/model.6/cv1/conv/Conv:	██████	4.792%
/model.23/cv2.1/cv2.1.2/Conv:	██████	4.543%
/model.7/conv/Conv:	██████	4.520%
/model.3/conv/Conv:	██████	4.362%
/model.16/cv2/conv/Conv:	██████	4.028%
/model.23/cv2.0/cv2.0.2/Conv:	██████	4.001%
/model.23/cv3.1/cv3.1.0/cv3.1.0.0/conv/Conv:	██████	3.954%
/model.9/cv2/conv/Conv:	██████	3.901%
/model.6/m.0/m.0/cv1/conv/Conv:	██████	3.891%
/model.10/m.0/m.0/ffn/ffn.0/conv/Conv:	██████	3.791%
/model.23/cv3.0/cv3.0.1/cv3.0.1.0/conv/Conv:	██████	3.711%
/model.4/cv1/conv/Conv:	██████	3.673%
/model.6/m.0/m.0/cv2/conv/Conv:	██████	3.620%
/model.6/m.0/m.1/cv1/conv/Conv:	██████	3.513%
/model.4/cv2/conv/Conv:	██████	3.421%
/model.5/conv/Conv:	██████	3.320%
/model.6/m.0/cv1/conv/Conv:	██████	3.073%
/model.2/cv1/conv/Conv:	██████	3.021%
/model.16/cv1/conv/Conv:	██████	2.764%
/model.6/m.0/m.1/cv2/conv/Conv:	██████	2.454%
/model.4/m.0/cv1/conv/Conv:	██████	2.408%
/model.4/m.0/cv2/conv/Conv:	██████	1.689%
/model.2/m.0/cv1/conv/Conv:	██████	1.602%
/model.9/cv1/conv/Conv:	██████	1.568%
/model.1/conv/Conv:	██████	1.205%
/model.23/cv3.0/cv3.0.0/cv3.0.0.0/conv/Conv:	██████	1.091%
/model.23/cv3.2/cv3.2.2/Conv:	██████	0.746%

(续下页)

(接上页)

/model.10/m/m.0/ffn/ffn.0/conv/Conv:		0.011%
/model.23/cv3.2/cv3.2.0/cv3.2.0.1/conv/Conv:		0.011%
/model.8/cv2/conv/Conv:		0.011%
/model.23/cv2.1/cv2.1.2/Conv:		0.010%
/model.22/m.0/cv3/conv/Conv:		0.010%
/model.23/cv2.1/cv2.1.1/conv/Conv:		0.008%
/model.10/m/m.0/ffn/ffn.1/conv/Conv:		0.008%
/model.23/cv2.0/cv2.0.1/conv/Conv:		0.007%
/model.10/m/m.0/attn/proj/conv/Conv:		0.007%
/model.8/m.0/cv1/conv/Conv:		0.007%
/model.22/m.0/m/m.0/cv2/conv/Conv:		0.006%
/model.8/m.0/m/m.1/cv2/conv/Conv:		0.005%
/model.22/cv2/conv/Conv:		0.005%
/model.20/conv/Conv:		0.005%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:		0.005%
/model.6/m.0/m/m.0/cv1/conv/Conv:		0.005%
/model.8/m.0/m/m.0/cv1/conv/Conv:		0.004%
/model.23/cv3.1/cv3.1.1/cv3.1.1.1/conv/Conv:		0.003%
/model.8/m.0/m/m.0/cv2/conv/Conv:		0.003%
/model.23/cv3.0/cv3.0.0/cv3.0.0.0/conv/Conv:		0.003%
/model.6/m.0/cv1/conv/Conv:		0.003%
/model.23/cv3.2/cv3.2.2/Conv:		0.003%
/model.23/cv3.2/cv3.2.1/cv3.2.1.0/conv/Conv:		0.003%
/model.6/m.0/m/m.1/cv2/conv/Conv:		0.003%
/model.23/cv3.2/cv3.2.1/cv3.2.1.1/conv/Conv:		0.002%
/model.22/m.0/m/m.1/cv2/conv/Conv:		0.002%
/model.6/m.0/m/m.0/cv2/conv/Conv:		0.002%
/model.23/cv3.0/cv3.0.1/cv3.0.1.0/conv/Conv:		0.002%
/model.10/m/m.0/attn/MatMul_1:		0.002%
/model.23/cv3.0/cv3.0.2/Conv:		0.001%
/model.23/cv3.1/cv3.1.2/Conv:		0.001%
/model.23/cv3.0/cv3.0.1/cv3.0.1.1/conv/Conv:		0.001%
/model.23/cv3.1/cv3.1.0/cv3.1.0.0/conv/Conv:		0.001%
/model.23/cv3.2/cv3.2.0/cv3.2.0.0/conv/Conv:		0.001%
/model.6/m.0/cv2/conv/Conv:		0.000%
/model.10/m/m.0/attn/MatMul:		0.000%
/model.8/m.0/cv2/conv/Conv:		0.000%
/model.22/m.0/cv2/conv/Conv:		0.000%

量化误差分析

在对 8-bit 量化应用量化感知训练后, 在相同输入下, 量化后的模型在 COCO val2017 上的 mAP50:95 提升至 36.0%; 同时输出层的累计误差大幅减少。相比前两种量化方式, 量化感知训练后的 8-bit 量化模型可以在最快的推理速度下达到最高的量化精度。

模型的输出层 /model.23/cv3.2/cv3.2.2/Conv, /model.23/cv2.2/cv2.2.2/Conv, /model.23/cv3.1/cv3.1.2/Conv, /model.23/cv2.1/cv2.1.2/Conv, /model.23/cv3.0/cv3.0.2/Conv 和 /model.23/cv2.0/cv2.0.2/Conv 的累计误差分别为 0.746%, 5.999%, 0.480%, 4.543%, 0.386% 和 4.001%。

备注: 如果想要更快的模型推理速度, 并且可以接受一定程度的精度损失, 可以考虑在量化 YOLO11N 的时候将输入大小设置为 320x320。不同分辨率下的模型推理速度可以在 [README.md](#) 中找到。

3.8.3 模型部署

参考示例

目标检测基类

- dl_detect_base.hpp
- dl_detect_base.cpp

前处理

ImagePreprocessor 类中封装了常用的图像前处理流程, 包括 color conversion, crop, resize, normalization, quantize。

- dl_image_preprocessor.hpp
- dl_image_preprocessor.cpp

后处理

- dl_detect_postprocessor.hpp
- dl_detect_postprocessor.cpp
- dl_detect_yolo11_postprocessor.hpp
- dl_detect_yolo11_postprocessor.cpp

3.9 如何部署 YOLO11n-pose

在本教程中, 我们介绍如何使用 ESP-PPQ 对预训练的 YOLO11n-pose 模型进行量化, 并使用 ESP-DL 部署量化后的 YOLO11n-pose 模型。

- 准备工作
- 模型量化
 - 预训练模型
 - 校准数据集
 - 8bit 默认配置量化
 - 量化感知训练
- 模型部署
 - 目标检测基类
 - 前处理
 - 后处理

(接上页)

/model.23/cv2.1/cv2.1.1/conv/Conv:		21.489%
/model.23/cv4.0/cv4.0.1/conv/Conv:		21.445%
/model.23/cv3.1/cv3.1.0/cv3.1.0.1/conv/Conv:		20.528%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:		20.083%
/model.23/cv3.1/cv3.1.0/cv3.1.0.0/conv/Conv:		20.066%
/model.13/m.0/cv2/conv/Conv:		20.042%
/model.22/m.0/cv3/conv/Conv:		19.737%
/model.10/m/m.0/ffn/ffn.1/conv/Conv:		19.585%
/model.23/cv3.1/cv3.1.1/cv3.1.1.1/conv/Conv:		19.392%
/model.23/cv3.0/cv3.0.1/cv3.0.1.0/conv/Conv:		18.773%
/model.23/cv3.2/cv3.2.1/cv3.2.1.1/conv/Conv:		18.688%
/model.22/cv1/conv/Conv:		18.579%
/model.19/cv2/conv/Conv:		18.494%
/model.22/m.0/m/m.1/cv2/conv/Conv:		17.576%
/model.17/conv/Conv:		17.224%
/model.19/cv1/conv/Conv:		17.140%
/model.22/cv2/conv/Conv:		16.785%
/model.23/cv4.2/cv4.2.1/conv/Conv:		16.375%
/model.23/cv4.2/cv4.2.0/conv/Conv:		16.167%
/model.23/cv2.1/cv2.1.0/conv/Conv:		15.655%
/model.23/cv3.0/cv3.0.1/cv3.0.1.1/conv/Conv:		15.504%
/model.23/cv2.2/cv2.2.0/conv/Conv:		15.431%
/model.10/m/m.0/attn/proj/conv/Conv:		15.251%
/model.23/cv3.0/cv3.0.0/cv3.0.0.1/conv/Conv:		15.171%
/model.22/m.0/m/m.0/cv2/conv/Conv:		15.006%
/model.19/m.0/cv1/conv/Conv:		14.692%
/model.23/cv2.2/cv2.2.1/conv/Conv:		14.548%
/model.22/m.0/m/m.0/cv1/conv/Conv:		13.065%
/model.16/m.0/cv2/conv/Conv:		12.980%
/model.22/m.0/m/m.1/cv1/conv/Conv:		12.921%
/model.10/m/m.0/attn/pe/conv/Conv:		12.745%
/model.23/cv4.1/cv4.1.2/Conv:		12.498%
/model.13/cv2/conv/Conv:		11.932%
/model.23/cv4.2/cv4.2.2/Conv:		11.797%
/model.13/m.0/cv1/conv/Conv:		11.777%
/model.16/cv2/conv/Conv:		10.892%
/model.13/cv1/conv/Conv:		10.760%
/model.23/cv2.0/cv2.0.0/conv/Conv:		10.352%
/model.23/cv4.0/cv4.0.2/Conv:		10.325%
/model.22/m.0/cv1/conv/Conv:		10.257%
/model.8/m.0/cv2/conv/Conv:		9.687%
/model.10/m/m.0/ffn/ffn.0/conv/Conv:		8.997%
/model.10/cv1/conv/Conv:		8.787%
/model.16/m.0/cv1/conv/Conv:		8.629%
/model.10/m/m.0/attn/qkv/conv/Conv:		8.600%
/model.8/m.0/cv3/conv/Conv:		8.328%
/model.10/m/m.0/attn/MatMul_1:		8.293%
/model.16/cv1/conv/Conv:		7.947%
/model.10/cv2/conv/Conv:		7.824%
/model.8/cv2/conv/Conv:		7.696%
/model.23/cv3.0/cv3.0.0/cv3.0.0.0/conv/Conv:		7.615%
/model.8/m.0/m/m.1/cv2/conv/Conv:		7.145%
/model.8/m.0/m/m.0/cv2/conv/Conv:		7.033%
/model.10/m/m.0/attn/MatMul:		6.707%
/model.8/m.0/m/m.1/cv1/conv/Conv:		6.376%
/model.23/cv2.1/cv2.1.2/Conv:		6.321%
/model.8/cv1/conv/Conv:		6.296%

(续下页)

(接上页)

/model.6/cv1/conv/Conv:		0.021%
/model.19/cv1/conv/Conv:		0.020%
/model.23/cv2.1/cv2.1.1/conv/Conv:		0.018%
/model.23/cv4.1/cv4.1.0/conv/Conv:		0.017%
/model.9/cv1/conv/Conv:		0.015%
/model.23/cv4.2/cv4.2.1/conv/Conv:		0.014%
/model.10/m/m.0/attn/qkv/conv/Conv:		0.014%
/model.19/cv2/conv/Conv:		0.014%
/model.16/m.0/cv2/conv/Conv:		0.014%
/model.23/cv4.2/cv4.2.0/conv/Conv:		0.014%
/model.6/m.0/m/m.0/cv1/conv/Conv:		0.013%
/model.22/m.0/cv3/conv/Conv:		0.013%
/model.23/cv3.2/cv3.2.0/cv3.2.0.1/conv/Conv:		0.013%
/model.23/cv4.0/cv4.0.0/conv/Conv:		0.013%
/model.23/cv3.1/cv3.1.1/cv3.1.1.1/conv/Conv:		0.013%
/model.22/m.0/m/m.1/cv1/conv/Conv:		0.012%
/model.6/m.0/cv3/conv/Conv:		0.012%
/model.10/m/m.0/attn/pe/conv/Conv:		0.012%
/model.23/cv4.1/cv4.1.1/conv/Conv:		0.011%
/model.8/m.0/m/m.1/cv1/conv/Conv:		0.011%
/model.13/m.0/cv1/conv/Conv:		0.011%
/model.22/m.0/m/m.0/cv1/conv/Conv:		0.011%
/model.6/m.0/m/m.1/cv1/conv/Conv:		0.011%
/model.23/cv3.2/cv3.2.1/cv3.2.1.1/conv/Conv:		0.011%
/model.8/m.0/cv3/conv/Conv:		0.010%
/model.7/conv/Conv:		0.010%
/model.17/conv/Conv:		0.009%
/model.8/m.0/m/m.0/cv1/conv/Conv:		0.009%
/model.13/m.0/cv2/conv/Conv:		0.009%
/model.10/m/m.0/attn/MatMul:		0.009%
/model.19/m.0/cv1/conv/Conv:		0.008%
/model.16/m.0/cv1/conv/Conv:		0.008%
/model.23/cv2.2/cv2.2.1/conv/Conv:		0.008%
/model.8/m.0/m/m.1/cv2/conv/Conv:		0.008%
/model.8/m.0/cv1/conv/Conv:		0.008%
/model.10/cv2/conv/Conv:		0.007%
/model.23/cv2.0/cv2.0.2/Conv:		0.007%
/model.22/m.0/cv1/conv/Conv:		0.007%
/model.6/m.0/cv1/conv/Conv:		0.007%
/model.23/cv2.0/cv2.0.0/conv/Conv:		0.006%
/model.23/cv2.1/cv2.1.0/conv/Conv:		0.006%
/model.22/m.0/m/m.1/cv2/conv/Conv:		0.006%
/model.23/cv3.2/cv3.2.1/cv3.2.1.0/conv/Conv:		0.005%
/model.8/m.0/m/m.0/cv2/conv/Conv:		0.005%
/model.23/cv2.1/cv2.1.2/Conv:		0.005%
/model.23/cv3.2/cv3.2.0/cv3.2.0.0/conv/Conv:		0.005%
/model.23/cv2.2/cv2.2.2/Conv:		0.005%
/model.22/cv1/conv/Conv:		0.004%
/model.10/m/m.0/attn/proj/conv/Conv:		0.004%
/model.23/cv4.2/cv4.2.2/Conv:		0.004%
/model.23/cv4.1/cv4.1.2/Conv:		0.004%
/model.22/m.0/m/m.0/cv2/conv/Conv:		0.004%
/model.23/cv2.2/cv2.2.0/conv/Conv:		0.003%
/model.6/m.0/m/m.1/cv2/conv/Conv:		0.003%
/model.23/cv4.0/cv4.0.1/conv/Conv:		0.003%
/model.6/m.0/m/m.0/cv2/conv/Conv:		0.003%
/model.10/m/m.0/attn/MatMul_1:		0.002%

(续下页)

(接上页)

/model.23/cv4.0/cv4.0.2/Conv:		0.002%
/model.10/m/m.0/ffn/ffn.1/conv/Conv:		0.002%
/model.20/conv/Conv:		0.002%
/model.23/cv2.0/cv2.0.1/conv/Conv:		0.002%
/model.10/m/m.0/ffn/ffn.0/conv/Conv:		0.001%
/model.23/cv3.1/cv3.1.0/cv3.1.0.1/conv/Conv:		0.001%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:		0.001%
/model.23/cv3.2/cv3.2.2/Conv:		0.001%
/model.23/cv3.0/cv3.0.1/cv3.0.1.0/conv/Conv:		0.001%
/model.23/cv3.1/cv3.1.0/cv3.1.0.0/conv/Conv:		0.001%
/model.23/cv3.1/cv3.1.2/Conv:		0.000%
/model.23/cv3.0/cv3.0.2/Conv:		0.000%
/model.23/cv3.0/cv3.0.0/cv3.0.0.1/conv/Conv:		0.000%
/model.6/m.0/cv2/conv/Conv:		0.000%
/model.23/cv3.0/cv3.0.1/cv3.0.1.1/conv/Conv:		0.000%
/model.23/cv3.0/cv3.0.0/cv3.0.0.0/conv/Conv:		0.000%
/model.8/m.0/cv2/conv/Conv:		0.000%
/model.22/m.0/cv2/conv/Conv:		0.000%

量化误差分析

在相同输入下, 量化后的模型在 COCO 上的 Pose mAP50:95 仅为 43.1%, 低于浮点模型 (50.0%), 存在一定的精度损失。

量化感知训练

为了进一步提高量化模型的精度, 可以采用量化感知训练。本示例基于 8-bit 量化方式进行量化感知训练。

量化设置

- `yolo11n-pose_qat.py`
- `trainer.py`

量化结果

Layer	NOISE:SIGNAL POWER RATIO
/model.22/m.0/cv2/conv/Conv:	██████████ 27.739%
/model.23/cv3.2/cv3.2.0/cv3.2.0.1/conv/Conv:	██████████ 26.872%
/model.23/cv4.1/cv4.1.0/conv/Conv:	██████████ 26.229%
/model.23/cv2.1/cv2.1.1/conv/Conv:	██████████ 25.300%
/model.23/cv3.2/cv3.2.1/cv3.2.1.0/conv/Conv:	██████████ 24.625%
/model.23/cv2.0/cv2.0.1/conv/Conv:	██████████ 23.751%
/model.20/conv/Conv:	██████████ 23.320%
/model.23/cv3.2/cv3.2.0/cv3.2.0.0/conv/Conv:	██████████ 22.901%
/model.23/cv4.1/cv4.1.1/conv/Conv:	██████████ 22.516%
/model.10/m/m.0/ffn/ffn.1/conv/Conv:	██████████ 22.035%
/model.19/m.0/cv2/conv/Conv:	██████████ 21.569%
/model.23/cv4.0/cv4.0.0/conv/Conv:	██████████ 21.199%
/model.23/cv3.1/cv3.1.0/cv3.1.0.1/conv/Conv:	██████████ 20.785%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:	██████████ 20.597%
/model.23/cv3.1/cv3.1.1/cv3.1.1.1/conv/Conv:	██████████ 20.329%
/model.23/cv4.0/cv4.0.1/conv/Conv:	██████████ 20.179%
/model.23/cv3.1/cv3.1.0/cv3.1.0.0/conv/Conv:	██████████ 19.983%
/model.22/m.0/cv3/conv/Conv:	██████████ 19.919%
/model.13/m.0/cv2/conv/Conv:	██████████ 19.424%
/model.23/cv3.0/cv3.0.1/cv3.0.1.0/conv/Conv:	██████████ 18.893%

(续下页)

(接上页)

/model.19/cv2/conv/Conv:	██████████	18.055%
/model.23/cv3.2/cv3.2.1/cv3.2.1.1/conv/Conv:	██████████	17.915%
/model.22/m.0/m.1/cv2/conv/Conv:	██████████	17.796%
/model.22/cv1/conv/Conv:	██████████	17.777%
/model.23/cv4.2/cv4.2.1/conv/Conv:	██████████	17.573%
/model.19/cv1/conv/Conv:	██████████	17.116%
/model.17/conv/Conv:	██████████	16.869%
/model.22/cv2/conv/Conv:	██████████	16.750%
/model.23/cv2.2/cv2.2.1/conv/Conv:	██████████	16.540%
/model.10/m.0/m.0/attn/proj/conv/Conv:	██████████	16.491%
/model.23/cv2.2/cv2.2.0/conv/Conv:	██████████	16.421%
/model.23/cv2.1/cv2.1.0/conv/Conv:	██████████	16.205%
/model.23/cv4.2/cv4.2.0/conv/Conv:	██████████	16.116%
/model.23/cv3.0/cv3.0.1/cv3.0.1.1/conv/Conv:	██████████	15.400%
/model.22/m.0/m.0/cv2/conv/Conv:	██████████	15.251%
/model.23/cv3.0/cv3.0.0/cv3.0.0.1/conv/Conv:	██████████	14.851%
/model.10/m.0/m.0/attn/pe/conv/Conv:	██████████	14.659%
/model.19/m.0/cv1/conv/Conv:	██████████	14.289%
/model.22/m.0/m.1/cv1/conv/Conv:	██████████	13.038%
/model.16/m.0/cv2/conv/Conv:	██████████	12.941%
/model.22/m.0/m.0/cv1/conv/Conv:	██████████	12.791%
/model.23/cv4.2/cv4.2.2/Conv:	██████████	12.508%
/model.23/cv4.1/cv4.1.2/Conv:	██████████	12.226%
/model.13/cv1/conv/Conv:	██████████	11.821%
/model.13/cv2/conv/Conv:	██████████	11.612%
/model.13/m.0/cv1/conv/Conv:	██████████	11.515%
/model.10/m.0/m.0/attn/MatMul_1:	██████████	11.303%
/model.16/cv2/conv/Conv:	██████████	11.028%
/model.10/m.0/m.0/attn/qkv/conv/Conv:	██████████	10.951%
/model.10/cv1/conv/Conv:	██████████	10.755%
/model.23/cv2.0/cv2.0.0/conv/Conv:	██████████	10.684%
/model.22/m.0/cv1/conv/Conv:	██████████	10.164%
/model.10/m.0/m.0/ffn/ffn.0/conv/Conv:	██████████	9.968%
/model.16/m.0/cv1/conv/Conv:	██████████	9.656%
/model.23/cv4.0/cv4.0.2/Conv:	██████████	9.566%
/model.8/m.0/cv2/conv/Conv:	██████████	9.521%
/model.10/cv2/conv/Conv:	██████████	8.068%
/model.16/cv1/conv/Conv:	██████████	7.989%
/model.23/cv2.1/cv2.1.2/Conv:	██████████	7.969%
/model.8/m.0/cv3/conv/Conv:	██████████	7.725%
/model.23/cv3.0/cv3.0.0/cv3.0.0.0/conv/Conv:	██████████	7.570%
/model.8/m.0/m.0/cv2/conv/Conv:	██████████	7.339%
/model.8/m.0/m.1/cv2/conv/Conv:	██████████	7.283%
/model.8/cv2/conv/Conv:	██████████	7.092%
/model.10/m.0/m.0/attn/MatMul:	██████████	6.654%
/model.8/cv1/conv/Conv:	██████████	6.492%
/model.8/m.0/m.1/cv1/conv/Conv:	██████████	6.451%
/model.23/cv2.0/cv2.0.2/Conv:	██████████	5.990%
/model.23/cv2.2/cv2.2.2/Conv:	██████████	5.902%
/model.6/m.0/m.0/cv2/conv/Conv:	██████████	5.898%
/model.6/m.0/cv2/conv/Conv:	██████████	5.881%
/model.6/m.0/cv3/conv/Conv:	██████████	5.402%
/model.8/m.0/cv1/conv/Conv:	██████████	5.210%
/model.23/cv3.2/cv3.2.2/Conv:	██████████	5.126%
/model.6/cv1/conv/Conv:	██████████	4.983%
/model.9/cv2/conv/Conv:	██████████	4.616%
/model.9/cv1/conv/Conv:	██████████	3.934%

(续下页)

(接上页)

/model.10/cv2/conv/Conv:	0.037%
/model.19/cv1/conv/Conv:	0.037%
/model.8/cv2/conv/Conv:	0.036%
/model.13/cv1/conv/Conv:	0.036%
/model.6/m.0/m.m.1/cv1/conv/Conv:	0.033%
/model.22/m.0/cv3/conv/Conv:	0.031%
/model.19/m.0/cv1/conv/Conv:	0.027%
/model.23/cv3.2/cv3.2.0/cv3.2.0.1/conv/Conv:	0.026%
/model.8/m.0/cv1/conv/Conv:	0.025%
/model.19/m.0/cv2/conv/Conv:	0.025%
/model.8/m.0/cv3/conv/Conv:	0.024%
/model.10/m.m.0/attn/qkv/conv/Conv:	0.023%
/model.8/m.0/m.m.0/cv1/conv/Conv:	0.023%
/model.22/m.0/cv1/conv/Conv:	0.021%
/model.6/m.0/m.m.0/cv1/conv/Conv:	0.021%
/model.23/cv2.0/cv2.0.0/conv/Conv:	0.020%
/model.6/m.0/cv1/conv/Conv:	0.020%
/model.23/cv4.0/cv4.0.0/conv/Conv:	0.019%
/model.9/cv1/conv/Conv:	0.018%
/model.23/cv4.1/cv4.1.2/Conv:	0.018%
/model.23/cv2.1/cv2.1.1/conv/Conv:	0.018%
/model.13/m.0/cv1/conv/Conv:	0.016%
/model.23/cv2.1/cv2.1.0/conv/Conv:	0.016%
/model.23/cv4.1/cv4.1.1/conv/Conv:	0.016%
/model.16/m.0/cv2/conv/Conv:	0.015%
/model.10/m.m.0/attn/proj/conv/Conv:	0.013%
/model.23/cv3.1/cv3.1.1/cv3.1.1.1/conv/Conv:	0.013%
/model.8/m.0/m.m.0/cv2/conv/Conv:	0.013%
/model.16/m.0/cv1/conv/Conv:	0.012%
/model.23/cv2.2/cv2.2.0/conv/Conv:	0.011%
/model.20/conv/Conv:	0.011%
/model.22/m.0/m.m.0/cv1/conv/Conv:	0.011%
/model.23/cv3.2/cv3.2.1/cv3.2.1.1/conv/Conv:	0.011%
/model.8/m.0/m.m.1/cv2/conv/Conv:	0.010%
/model.23/cv2.0/cv2.0.2/Conv:	0.009%
/model.10/m.m.0/attn/MatMul:	0.009%
/model.22/cv1/conv/Conv:	0.009%
/model.13/m.0/cv2/conv/Conv:	0.008%
/model.23/cv2.2/cv2.2.1/conv/Conv:	0.008%
/model.23/cv2.1/cv2.1.2/Conv:	0.007%
/model.23/cv3.2/cv3.2.1/cv3.2.1.0/conv/Conv:	0.007%
/model.22/m.0/m.m.1/cv2/conv/Conv:	0.007%
/model.6/m.0/m.m.0/cv2/conv/Conv:	0.006%
/model.22/m.0/m.m.0/cv2/conv/Conv:	0.006%
/model.23/cv4.0/cv4.0.1/conv/Conv:	0.005%
/model.23/cv3.2/cv3.2.0/cv3.2.0.0/conv/Conv:	0.005%
/model.23/cv4.0/cv4.0.2/Conv:	0.004%
/model.6/m.0/m.m.1/cv2/conv/Conv:	0.004%
/model.23/cv3.0/cv3.0.0/cv3.0.0.1/conv/Conv:	0.004%
/model.10/m.m.0/ffn/ffn.1/conv/Conv:	0.003%
/model.23/cv3.2/cv3.2.2/Conv:	0.003%
/model.10/m.m.0/attn/MatMul_1:	0.002%
/model.10/m.m.0/ffn/ffn.0/conv/Conv:	0.002%
/model.23/cv3.1/cv3.1.0/cv3.1.0.1/conv/Conv:	0.002%
/model.23/cv2.0/cv2.0.1/conv/Conv:	0.002%
/model.23/cv3.1/cv3.1.1/cv3.1.1.0/conv/Conv:	0.001%
/model.23/cv3.0/cv3.0.2/Conv:	0.001%

(续下页)

(接上页)

/model.23/cv3.1/cv3.1.2/Conv:		0.001%
/model.23/cv3.0/cv3.0.1/cv3.0.1.0/conv/Conv:		0.001%
/model.23/cv3.1/cv3.1.0/cv3.1.0.0/conv/Conv:		0.001%
/model.23/cv3.0/cv3.0.0/cv3.0.0.0/conv/Conv:		0.000%
/model.6/m.0/cv2/conv/Conv:		0.000%
/model.23/cv3.0/cv3.0.1/cv3.0.1.1/conv/Conv:		0.000%
/model.8/m.0/cv2/conv/Conv:		0.000%
/model.22/m.0/cv2/conv/Conv:		0.000%

量化误差分析

在对 8-bit 量化应用量化感知训练后，在相同输入下，量化后的模型在 COCO 上的 Pose mAP50:95 提升至 44.9%；同时输出层的累计误差大幅减少。相比 8-bit 后量化方式，量化感知训练后的 8-bit 量化模型可以在相同的推理速度下达到最高的量化精度。

3.9.3 模型部署

example

目标检测基类

- dl_detect_base.hpp
- dl_detect_base.cpp

前处理

ImagePreprocessor 类中封装了常用的图像前处理流程，包括 color conversion, crop, resize, normalization, quantize。

- dl_image_preprocessor.hpp
- dl_image_preprocessor.cpp

后处理

- dl_detect_postprocessor.hpp
- dl_detect_postprocessor.cpp
- dl_pose_yolo11_postprocessor.hpp
- dl_pose_yolo11_postprocessor.cpp

3.10 如何部署流式模型

时间序列模型如今被应用在许多领域，例如，音频领域。而音频模型在部署时通常有两种模式：

- **Offline** 模式：模型需要一次性接收完整的音频数据（例如整个语音文件），然后进行整体处理。
- **Streaming** 模式：流式模式下，模型逐帧（逐块）接收音频数据，实时处理并输出中间结果。

在本教程中，我们来介绍如何使用 ESP-PPQ 量化流式模型，并使用 ESP-DL 部署量化后的流式模型。

- 准备工作
- 模型量化
 - 自动流式转换
 - 自动流式转换的工作原理
 - 手动流式缓存配置
- 模型部署

3.10.1 准备工作

1. 安装 `ESP_IDF`
2. 安装 `ESP_PPQ`

3.10.2 模型量化

参考示例

时间序列模型种类繁多，这里仅以 Temporal Convolutional Network(TCN) 为例，不熟悉的可自行查找资料了解，这里不过多介绍其细节。其它模型需根据自身情况，量体裁衣。

该示例代码中构建了一个 TCN 模型：`models.py` (模型非完整，仅用于演示)。

ESP-PPQ 提供了自动流式转换功能，可以简化创建流式模型的过程。通过 `auto_streaming=True` 参数，ESP-PPQ 自动处理流式推理所需的模型转换。

备注：

- **Offline** 模式，模型输入是一段完整数据，`input shape` 在时间维度上的 `size` 一般比较大（例如 `[1, 16, 15]`）。
- **Streaming** 模式，模型输入是连续数据，在时间维度上的 `size` 较小，匹配实时处理的块大小（例如 `[1, 16, 3]`）。

自动流式转换

ESP-PPQ 通过量化过程中的 `auto_streaming=True` 参数提供自动流式转换功能。启用此标志后, ESP-PPQ 会自动转换模型以支持流式推理:

1. 分析模型结构以识别适当的分块点
2. 创建内部状态管理以在块之间保持上下文
3. 生成适合流式场景的优化代码

自动流式转换的工作原理

ESP-PPQ 中的自动流式转换会分析模型图, 并在关键位置插入 `StreamingCache` 节点以实现时间上下文保持。转换过程遵循以下原则:

1. 算子分类

- **支持流式的算子**: 需要时间上下文的卷积、池化和转置卷积操作 (例如 `Conv`、`AveragePool`、`MaxPool`、`ConvTranspose`)。
- **绕过算子**: 不需要时间上下文的激活函数、数学运算、量化节点和其他操作 (例如 `Relu`、`Add`、`MatMul`、`LayerNorm`)。

2. 窗口大小计算

对于支持流式的算子, ESP-PPQ 根据以下因素计算所需的缓存窗口大小: - Kernel size and dilation rates - Padding configuration - Stride values

窗口大小决定了需要缓存多少历史帧才能正确计算当前帧。

3. StreamingCache 节点插入

ESP-PPQ 在支持流式的算子之前插入 `StreamingCache` 节点。这些节点: - 维护历史帧的滑动窗口缓冲区 - 调整张量形状以容纳缓存窗口 - 保留原始操作的量化配置 - 管理帧轴对齐以进行正确的时间处理

4. 填充调整

对于流式操作, ESP-PPQ 调整填充配置: - 移除底部填充以防止前瞻到未来帧 - 保持对称或仅顶部填充以实现因果处理

限制和注意事项

- 自动转换开箱即用支持基于卷积的时间操作
- 自定义操作或复杂的时间依赖关系可能需要手动配置流式表
- 转换假设时间维度沿轴 1 (可通过 `streaming_table` 配置)

以下是如何使用自动流式功能的示例:

```
# 导出非流式模型
quant_ppq_graph = espdl_quantize_torch(
    model=model,
    espdl_export_file=ESPD_MODEL_PATH,
    calib_data_loader=data_loader,
    calib_steps=32, # 校准步数
    input_shape=INPUT_SHAPE, # 离线模式的输入形状
    inputs=None,
    target=TARGET, # 量化目标类型
    num_of_bits=NUM_OF_BITS, # 量化位数
    dispatching_override=None,
    device=DEVICE,
    error_report=True,
```

(续下页)

(接上页)

```

skip_export=False,
export_test_values=True,
verbose=1, # 输出详细日志信息
)

# 使用自动转换导出流式模型
quant_ppq_graph = espdl_quantize_torch(
    model=model,
    espdl_export_file=ESPDL_STEAMING_MODEL_PATH,
    calib_data_loader=dataloader,
    calib_steps=32,
    input_shape=INPUT_SHAPE,
    inputs=None,
    target=TARGET,
    num_of_bits=NUM_OF_BITS,
    dispatching_override=None,
    device=DEVICE,
    error_report=True,
    skip_export=False,
    export_test_values=False,
    verbose=1,
    auto_streaming=True, # 启用自动流式转换
    streaming_input_shape=[1, 16, 3], # 流式模式的输入形状
    streaming_table=None,
)

```

手动流式缓存配置

对于 ESP-PPQ 流式转换功能不自动支持的算子（例如 Transpose、Reshape、Slice 等），您可以使用 `insert_streaming_cache_on_var` 函数手动插入 `StreamingCache` 节点。该函数允许您为无法自动插入 `streamingCache` 的变量指定缓存属性。

`insert_streaming_cache_on_var` 函数的签名如下：

```

def insert_streaming_cache_on_var(
    var_name: str,
    window_size: int,
    op_name: str = None,
    frame_axis: int = 1
) -> Dict[str, Any]

```

参数说明：- `var_name`：需要插入流式缓存的变量名称 - `window_size`：缓存窗口大小（需要缓存的帧数）
- `op_name`：（可选）与变量关联的算子名称 - `frame_axis`：（可选）表示时间维度的轴，默认为 1

该函数返回一个包含流式缓存配置的字典，应将其添加到 `streaming_table` 列表中并传递给 `espdl_quantize_torch` 函数。

使用示例：

```

streaming_table = []
# 为无法自动插入 streamingCache 的变量手动指定缓存属性
streaming_table.append(
    insert_streaming_cache_on_var("/out_conv/Conv_output_0", output_frame_size - 1)
)
streaming_table.append(insert_streaming_cache_on_var("PPQ_Variable_0", 1, "/Slice"))

```

(续下页)

(接上页)

```

quant_ppq_graph = espdl_quantize_torch(
    model=model,
    espdl_export_file=ESPDL_STEAMING_MODEL_PATH,
    calib_data_loader=dataloader,
    calib_steps=32,
    input_shape=INPUT_SHAPE,
    inputs=None,
    target=TARGET,
    num_of_bits=NUM_OF_BITS,
    dispatching_override=None,
    device=DEVICE,
    error_report=True,
    skip_export=False,
    export_test_values=False,
    verbose=1,
    auto_streaming=True,
    streaming_input_shape=[1, 16, 3],
    streaming_table=streaming_table, # 传递手动配置的流式表
)

```

3.10.3 模型部署

参考示例, 该示例使用预生成的数据来模拟实时数据流。

备注: 基础的模型加载和推理方法, 可参考其它文档, 这里不再赘述:

- 如何加载和测试模型
- 如何进行模型推理

在流式模式下, 模型按时间接收数据块, 而不是要求一次性获得整个输入。流式模型依次处理这些块, 同时在块之间保持内部状态。部署代码负责将输入分解为适当的块并将其馈送到模型。见 `app_main.cpp` 如下代码块:

```

dl::TensorBase *run_streaming_model(dl::Model *model, dl::TensorBase *test_input)
{
    std::map<std::string, dl::TensorBase *> model_inputs = model->get_inputs();
    dl::TensorBase *model_input = model_inputs.begin()->second;
    std::map<std::string, dl::TensorBase *> model_outputs = model->get_outputs();
    dl::TensorBase *model_output = model_outputs.begin()->second;

    if (!test_input) {
        ESP_LOGE(TAG,
            "Model input doesn't have a corresponding test input. Please enable_
↪export_test_values option "
            "in esp-ppq when export espdl model.");
        return nullptr;
    }

    int test_input_size = test_input->get_bytes();
    uint8_t *test_input_ptr = (uint8_t *)test_input->data;
    int model_input_size = model_input->get_bytes();
    uint8_t *model_input_ptr = (uint8_t *)model_input->data;
    int chunks = test_input_size / model_input_size;

```

(续下页)

(接上页)

```
for (int i = 0; i < chunks; i++) {  
    // assign chunk data to model input  
    memcpy(model_input_ptr, test_input_ptr + i * model_input_size, model_input_  
↪size);  
    model->run(model_input);  
}  
  
return model_output;  
}
```

这种方法允许模型通过将长序列分解为更小、更易管理的块来高效处理。每个块依次馈送到模型中，内部状态自动维护以确保跨块的连续性。

备注:

- 块的数量是根据完整输入大小与流式模型输入大小的比率计算的。
- ESP-DL 流式模型自动处理内部状态管理，使部署变得简单。
- 流式模型的输出应与等效离线模型输出的最后部分匹配。

3.11 使用 TQT 量化模型

本文档说明为何以及如何如何在 ESP-PPQ 中使用 Trained Quantization Thresholds (TQT) 做量化，以获得更高的量化精度。注意 ESP-PPQ 必须是 1.2.7 版本或更新版本。

- 为何使用 *TQT*
 - 后量化 (*PTQ*) 的局限性
 - 量化感知训练 (*QAT*) 的复杂性
 - *TQT* 能带来什么
- 如何使用 *TQT*
 - 快速开始
 - *TQTSetting* 常用参数
- *TQT* 量化示例
 - *YOLO26n* 量化
 - *MobileNetV2* 量化
- 常见问题
 - 如何加速 *TQT*?
 - 可以不修改 *scale* 只训练模型权重吗?
 - *TQT* 可以和 *Weight Equalization* 一起使用吗?

3.11.1 为何使用 TQT

TQT (Trained Quantization Thresholds) 出自论文 [Trained Quantization Thresholds for Accurate and Efficient Fixed-Point Inference of Deep Neural Networks](#) (Sambhav R. Jain, Albert Gural, Michael Wu, Chris H. Dick), 发表于 **MLSys 2020**。由于 ESP32 系列芯片目前仅支持 **Per-Tensor + Symmetric + Power-of-2** 量化策略生成的 `.espd1` 模型, 我们在 ESP-PPQ 中引入了 TQT 的核心思想, 通过 **标准反向传播与梯度下降联合优化** 量化阈值 (scale) 和模型权重, 以适配硬件约束。

在 ESP-PPQ 中, TQT 以 `TrainedQuantizationThresholdPass` 的形式实现, 我们在 log 域对 scale 进行优化, 并结合 ESP-DL 的 **Power-of-2** 约束做了适配 (如 `int_lambda`, `STE` 等)。

后量化 (PTQ) 的局限性

方式	优点	局限
PTQ	实现简单、无需训练	scale 由统计得到, 未必最优; 对敏感结构 (如 YOLO26n head) 误差较大
TQT	微调 scale/权重	需要一定校准数据与算力

量化感知训练 (QAT) 的复杂性

量化感知训练 (QAT) 通常在完整训练或微调阶段插入伪量化, 能获得较好精度, 但需要:

- **完整标注数据与较长训练、微调流程** (对于不使用标注数据进行训练的模型, QAT 也往往花费较长的训练、微调流程);
- **较大算力与调参成本**;
- 与现有 PTQ/校准工具链的衔接往往需要额外工程。

TQT 则 **无需标签**: 损失为浮点输出和量化输出的 MSE, 只需校准集即可, 操作简单, 在保持 Power-of-2 约束下联合微调权重与 $\log_2(\text{scale})$, 通常用较少步数即可明显提升精度。

TQT 能带来什么

- **可学习的 scale (仍为 2 的 k 次方)**: 在 log 域优化 $\alpha = \log_2(\text{scale})$, 数值稳定且自然满足 `POWER_OF_2`。但如果不满足条件或被禁用, 则不会训练。
- **STE 与 range-precision 折中**: 对阈值梯度使用 STE 并合理设计前向/反向行为, 可在 **表示范围**与 **精度** 之间取得更好折中; ESP-PPQ 中通过 `alpha_ste` 实现。
- **可学习权重**: 权重可训练, 可与 scale 一起做块级微调, 进一步减小量化误差。
- **与 ESP-DL 对齐**: 使用默认的 `alpha_ste` 前向时, 前向用的就是整数 exponent; 配合 `int_lambda` 可让学到的 alpha 更接近整数, 便于导出与芯片端一致。
- **无需标签**: 只需校准集即可做 scale/权重微调。

3.11.2 如何使用 TQT

快速开始

1. 使用 ESP-DL 默认 setting 并开启 TQT

```

from esp_ppq import QuantizationSettingFactory
from esp_ppq.api import espdl_quantize_onnx

quant_setting = QuantizationSettingFactory.espdl_setting()
quant_setting.tqt_optimization = True
quant_setting.tqt_optimization_setting.int_lambda = 0.25 # 可选: 让 exponent_
→更接近整数
quant_setting.tqt_optimization_setting.steps = 500
quant_setting.tqt_optimization_setting.lr = 1e-5
quant_setting.tqt_optimization_setting.collecting_device = 'cuda' # or 'cpu'
quant_setting.tqt_optimization_setting.block_size = 2

```

2. 传入 setting 做量化

```

from esp_ppq.api import ENABLE_CUDA_KERNEL
with ENABLE_CUDA_KERNEL():
    quant_ppq_graph = espdl_quantize_onnx(
        onnx_import_file=ONNX_PATH,
        espdl_export_file=ESPDL_MODLE_PATH,
        calib_dataloader=dataloader,
        calib_steps=32,
        input_shape=[1] + INPUT_SHAPE,
        target=TARGET,
        num_of_bits=NUM_OF_BITS,
        collate_fn=collate_fn,
        setting=quant_setting,
        device=DEVICE,
        error_report=False,
        skip_export=False,
        export_test_values=True,
        verbose=0,
        inputs=None,
    )

```

3. 导出与返回值

- 返回的 `quant_ppq_graph` 里 `scale` 为 TQT 微调后的 2 的整数次方, 可用于模型精度评估。
- 导出的 `espdl` 文件: `exponent` 来自 `int(log2(config.scale))`, 导出与芯片端推理会使用 TQT 优化后的 `scale`。

TQTSetting 常用参数

参数	类型	默认值	说明
interested_layers	List[str]	[]	指定要微调的 op 名称; 空则对满足条件的 Conv/Gemm 等全部微调
steps	int	500	每个 block 的微调步数
lr	float	1e-5	学习率
block_size	int	4	图划分 block 的大小, 影响稳定性与速度
is_scale_trainable	bool	True	是否微调 scale (需 POWER_OF_2 + LINEAR + SYMMETRICAL)。True: 同时训练 scale 与 weights; False: 仅训练 weights
gamma	float	0.0	正则: MSE(浮点输出, 量化输出)
int_lambda	float	0.0	正则: 让 alpha 靠近 round(alpha), 取值范围为 [0.0, 1.0]
collecting_device	str	cpu	缓存校准数据的设备, 有 GPU 可设为 cuda

3.11.3 TQT 量化示例

以下示例展示如何对 **YOLO26n** (检测) 和 **MobileNetV2** (分类) 模型启用 TQT 量化并导出 ESP-DL。校准数据与模型路径需按实际环境修改。

YOLO26n 量化

准备

项目	说明
模型名称	YOLO26n (one2one 分支 end2end 推理)
任务	目标检测
输入形状	[1, 3, 640, 640] (NCHW)
ONNX	yolo26n_o2o.onnx
校准集	calib_yolo26n.zip
备注	也可以通过 one2many 分支做非 end2end 推理, ONNX 为 yolo26n_o2m.onnx, 量化方法与 end2end 模型一致

量化脚本

```
import os
from esp_ppq import QuantizationSettingFactory
from esp_ppq.api import espdL_quantize_onnx
from torch.utils.data import DataLoader
import torch
from torch.utils.data import Dataset
from torchvision import transforms
from PIL import Image
from onnxsim import simplify
import onnx
import zipfile
import urllib.request
```

(续下页)

(接上页)

```

from esp_ppq.api import ENABLE_CUDA_KERNEL

class CaliDataset(Dataset):
    def __init__(self, path, img_shape=640):
        super().__init__()
        self.transform = transforms.Compose(
            [
                transforms.Resize((img_shape, img_shape)),
                transforms.ToTensor(),
                transforms.Normalize(mean=[0, 0, 0], std=[1, 1, 1]),
            ]
        )

        self.imgs_path = []
        self.path = path
        for img_name in os.listdir(self.path):
            img_path = os.path.join(self.path, img_name)
            self.imgs_path.append(img_path)

    def __len__(self):
        return len(self.imgs_path)

    def __getitem__(self, idx):
        img = Image.open(self.imgs_path[idx])
        if img.mode == 'L':
            img = img.convert('RGB')
        img = self.transform(img)
        return img

def report_hook(blocknum, blocksize, total):
    downloaded = blocknum * blocksize
    percent = downloaded / total * 100
    print(f"\rDownloading calibration dataset: {percent:.2f}%", end="")

def quant_yolo26n(imgsz):
    BATCH_SIZE = 32
    INPUT_SHAPE = [3, imgsz, imgsz]
    DEVICE = "cpu"
    TARGET = "esp32p4" # or "esp32s3"
    NUM_OF_BITS = 8

    yolo26n_onnx_url = "https://dl.espressif.com/public/yolo26n_o2o.onnx"
    ONNX_PATH = "yolo26n_o2o.onnx"
    urllib.request.urlretrieve(
        yolo26n_onnx_url, "yolo26n_o2o.onnx", reporthook=report_hook
    )

    ESPDL_MODULE_PATH = "yolo26n_o2o_ptq_fq_tqt_p4_640.espd1"

    yolo26n_caib_url = "https://dl.espressif.com/public/calib_yolo26n.zip"
    CALIB_DIR = "calib_yolo26n"
    urllib.request.urlretrieve(
        yolo26n_caib_url, "calib_yolo26n.zip", reporthook=report_hook
    )
    with zipfile.ZipFile("calib_yolo26n.zip", "r") as zip_file:

```

(续下页)

(接上页)

```

zip_file.extractall("./")

model = onnx.load(ONNX_PATH)
sim = True
if sim:
    model, check = simplify(model)
    assert check, "Simplified ONNX model could not be validated"
onnx.save(onnx.shape_inference.infer_shapes(model), ONNX_PATH)

calibration_dataset = CaliDataset(CALIB_DIR, img_shape=imgsz)
dataloader = DataLoader(
    dataset=calibration_dataset, batch_size=BATCH_SIZE, shuffle=False, num_
↪workers=8,
)

def collate_fn(batch: torch.Tensor) -> torch.Tensor:
    return batch.to(DEVICE)

# default setting
quant_setting = QuantizationSettingFactory.espdn_setting()
# activation calibration algo
quant_setting.quantize_activation_setting.calib_algorithm = "percentile" # kl -->
↪ percentile to get better mAP
# focused logits quantization
quant_setting.quant_config_modify = True
# o2o
quant_setting.quant_config_modify_setting.custom_config = {
    "/model.23/one2one_cv3.0/one2one_cv3.0.2/Conv": 0.0625,
    "/model.23/one2one_cv3.1/one2one_cv3.1.2/Conv": 0.0625,
    "/model.23/one2one_cv3.2/one2one_cv3.2.2/Conv": 0.0625,
}

# o2m
# quant_setting.quant_config_modify_setting.custom_config = {
#     "/model.23/cv3.0/cv3.0.2/Conv": 0.0625,
#     "/model.23/cv3.1/cv3.1.2/Conv": 0.0625,
#     "/model.23/cv3.2/cv3.2.2/Conv": 0.0625,
# }

# TQT
quant_setting.tqt_optimization = True
quant_setting.tqt_optimization_setting.collecting_device = "cpu"
quant_setting.tqt_optimization_setting.steps = 200 #300 for o2m
quant_setting.tqt_optimization_setting.block_size = 2
quant_setting.tqt_optimization_setting.lr = 1e-5

quant_ppq_graph = espdn_quantize_onnx(
    onnx_import_file=ONNX_PATH,
    espdn_export_file=ESPDN_MODLE_PATH,
    calib_dataloader=dataloader,
    calib_steps=32,
    input_shape=[1] + INPUT_SHAPE,
    target=TARGET,
    num_of_bits=NUM_OF_BITS,
    collate_fn=collate_fn,
    setting=quant_setting,

```

(续下页)

```

device=DEVICE,
error_report=True,
skip_export=False,
export_test_values=False,
verbose=0,
inputs=None,
)
return quant_ppq_graph

if __name__ == "__main__":
    quant_yolo26n(imgsz=640)

```

Focused Logits Quantization

目标检测模型中，与 backbone/neck 不同，**检测头 (head) 输出**对量化误差非常敏感。同时，对于类似 YOLO26n 这样的检测模型，其 classification 分支的输出 logits 为 **Sigmoid 的输入**，经过 Sigmoid 可以得到类别置信度或分数。此时，若将 logits 与其它层使用同一套校准得到的 scale，容易得到 **过大的 scale (量化步长过粗)**，导致类别分数失真，mAP 下降。

- **Sigmoid 的有效输入范围有限**： $\sigma(x) = 1/(1+\exp(-x))$ 在 x 的绝对值较大时很快饱和。 x 很负时输出趋近 0，很正时趋近 1，只有中间一段（大致在 -4 到 4 内）变化明显。也就是说，对输出真正有区分度的 logits 只分布在这段 **有限范围**内，超出部分进入饱和区、几乎不再变化。
- **两边都有饱和区，检测中更关注正样本**：在目标检测里，我们更关心 **正样本** (有目标、高置信度)，对应 logits 偏正、Sigmoid 输出接近 1 的一侧。但无论正负两侧，只要 logits 被量化得过粗 (scale 过大、步长太大)，多个不同的 logits 会被舍入到同一量化档位：在饱和区附近会成片地被压成 0 或 1，在中间有效区则丢失细粒度区分，导致分数失真，mAP 掉点。

因此，对 **送入 Sigmoid 的 logits** (即示例中 one2one_cv3.0/1/2.2 这三个 head 分支的最后一层 Conv 输出) 做 **Focused logits quantization**：为这几层 **单独指定更细的 scale** (如 $0.0625 = 2^{-4}$ 次方)，在保持 Power-of-2 的前提下减小量化步长，在 Sigmoid 的有效输入范围内保留更多档位，避免过早进入饱和、并保留正样本侧的区分度，从而稳定或提升 mAP。若使用其他检测模型或不同导出结构，可根据 Sigmoid 的输入来定位对应层并调整 custom_config。

ESP32-P4 上模型量化精度测试结果

在 COCO val2017 评估 yolo26n(size=640 pixels) 量化后的精度：

配置	mAP50-95(o2m)	mAP50-95(e2e)	备注
PTQ (无 TQT)	0.342	0.332	仅校准
PTQ + TQT	0.371	0.363	校准 + TQT 微调 scale/权重

测试结果表明，无论是对 end2end 推理模型，还是非 end2end 推理模型，TQT 均可以显著提升量化后模型的 mAP。

备注：如果追求更快的模型推理速度，可以使用 size=512 pixels。实验结果表明，在 size=512 pixels 和 end2end 推理模式下，8-bit PTQ 量化模型的 mAP50-95 为 0.315，经过 TQT 后 mAP50-95 可提升至 0.341，精度提升了 2.6 个百分点。

MobileNetV2 量化

量化脚本

```

import os
import subprocess
from typing import Iterable, Tuple, List, Tuple

import torch
from datasets.imagenet_util import (
    evaluate_ppq_module_with_imagenet,
    load_imagenet_from_directory,
)
from esp_ppq import QuantizationSettingFactory, QuantizationSetting
from esp_ppq.api import espdl_quantize_onnx, get_target_platform
from torch.utils.data import DataLoader
import torchvision.datasets as datasets
import torchvision.transforms as transforms
from torch.utils.data.dataset import Subset
import urllib.request
import zipfile

def quant_setting_mobilenet_v2(
    onnx_path: str,
    optim_quant_method: List[str] = None,
) -> Tuple[QuantizationSetting, str]:
    '''
    Quantize onnx model with optim_quant_method.

    Args:
        optim_quant_method (List[str]): support 'MixedPrecision_quantization',
        ↳ 'LayerwiseEqualization_quantization'
        ↳ -'MixedPrecision_quantization': if some layers in model have larger errors in
        ↳ 8-bit quantization, dispatching
        ↳ the layers to 16-bit quantization. You can
        ↳ remove or add layers according to your
        ↳ needs.
        ↳ -'LayerwiseEqualization_quantization': using weight equalization strategy,
        ↳ which is proposed by Markus Nagel.
        ↳ Refer to paper https://openaccess.thecvf.com/content\_ICCV\_2019/papers/Nagel\_Data-Free\_Quantization\_Through\_Weight\_Equalization\_and\_Bias\_Correction\_ICCV\_2019\_paper.pdf for more information.
        ↳ Since ReLU6 exists in MobilenetV2,
        ↳ convert ReLU6 to ReLU for better precision.

    Returns:
        [tuple]: [QuantizationSetting, str]
    '''
    quant_setting = QuantizationSettingFactory.espdl_setting()
    if optim_quant_method is not None:
        if "MixedPrecision_quantization" in optim_quant_method:
            # These layers have larger errors in 8-bit quantization, dispatching to
            ↳ 16-bit quantization.
            # You can remove or add layers according to your needs.
            quant_setting.dispatching_table.append(
                "/features/features.1/conv/conv.0/conv.0.0/Conv",

```

(续下页)

```

        get_target_platform(TARGET, 16),
    )
    quant_setting.dispatching_table.append(
        "/features/features.1/conv/conv.0/conv.0.2/Clip",
        get_target_platform(TARGET, 16),
    )
    elif "LayerwiseEqualization_quantization" in optim_quant_method:
        # layerwise equalization
        quant_setting.equalization = True
        quant_setting.equalization_setting.iterations = 4
        quant_setting.equalization_setting.value_threshold = 0.4
        quant_setting.equalization_setting.opt_level = 2
        quant_setting.equalization_setting.interested_layers = None
        # replace ReLU6 with ReLU
        onnx_path = onnx_path.replace("mobilenet_v2.onnx", "mobilenet_v2_relu.onnx
↪")
    else:
        raise ValueError(
            "Please set optim_quant_method correctly. Support 'MixedPrecision_
↪quantization', 'LayerwiseEqualization_quantization'"
        )

    return quant_setting, onnx_path

def collate_fn1(x: Tuple) -> torch.Tensor:
    return torch.cat([sample[0].unsqueeze(0) for sample in x], dim=0)

def collate_fn2(batch: torch.Tensor) -> torch.Tensor:
    return batch.to(DEVICE)

def report_hook(blocknum, blocksize, total):
    downloaded = blocknum * blocksize
    percent = downloaded / total * 100
    print(f"\rDownloading calibration dataset: {percent:.2f}%", end="")

if __name__ == "__main__":
    BATCH_SIZE = 32
    INPUT_SHAPE = [3, 224, 224]
    DEVICE = "cpu" # 'cuda' or 'cpu', if you use cuda, please make sure that cuda_
↪is available
    TARGET = "esp32p4" # 'c', 'esp32s3' or 'esp32p4'
    NUM_OF_BITS = 8
    ONNX_PATH = "./models/torch/mobilenet_v2.onnx" #'models/onnx/mobilenet_v2.onnx'
    ESPDL_MODEL_PATH = "models/onnx/mobilenet_v2.espd1"
    CALIB_DIR = "./imagenet"

    # Download mobilenet_v2 model from onnx models and dataset
    imagenet_url = "https://dl.espressif.com/public/imagenet_calib.zip"
    os.makedirs(CALIB_DIR, exist_ok=True)
    if not os.path.exists("imagenet_calib.zip"):
        urllib.request.urlretrieve(
            imagenet_url, "imagenet_calib.zip", reporthook=report_hook
        )

```

(接上页)

```

if not os.path.exists(os.path.join(CALIB_DIR, "calib")):
    with zipfile.ZipFile("imagenet_calib.zip", "r") as zip_file:
        zip_file.extractall(CALIB_DIR)
CALIB_DIR = os.path.join(CALIB_DIR, "calib")

# -----
# Prepare Calibration Dataset
# -----
if os.path.exists(CALIB_DIR):
    print(f"load imagenet calibration dataset from directory: {CALIB_DIR}")
    dataset = datasets.ImageFolder(
        CALIB_DIR,
        transforms.Compose(
            [
                transforms.Resize(256),
                transforms.CenterCrop(224),
                transforms.ToTensor(),
                transforms.Normalize(
                    mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225]
                ),
            ]
        ),
    )
    dataset = Subset(dataset, indices=[_ for _ in range(0, 1024)])
    dataloader = DataLoader(
        dataset=dataset,
        batch_size=BATCH_SIZE,
        shuffle=False,
        num_workers=4,
        pin_memory=False,
        collate_fn=collate_fn1,
    )
else:
    # Random calibration dataset only for debug
    print("load random calibration dataset")

    def load_random_calibration_dataset() -> Iterable:
        return [torch.rand(size=INPUT_SHAPE) for _ in range(BATCH_SIZE)]

    # Load training data for creating a calibration dataloader.
    dataloader = DataLoader(
        dataset=load_random_calibration_dataset(),
        batch_size=BATCH_SIZE,
        shuffle=False,
    )

# -----
# Quantize ONNX Model.
# -----

quant_setting = QuantizationSettingFactory.espdn_setting()
# TQT
quant_setting.tqt_optimization = True
quant_setting.tqt_optimization_setting.collecting_device = "cpu"
quant_setting.tqt_optimization_setting.steps = 500
quant_setting.tqt_optimization_setting.block_size = 4
quant_setting.tqt_optimization_setting.lr = 1e-4

```

(续下页)

```

quant_ppq_graph = espdl_quantize_onnx(
    onnx_import_file=ONNX_PATH,
    espdl_export_file=ESPDL_MODEL_PATH,
    calib_dataloader=dataloader,
    calib_steps=32,
    input_shape=[1] + INPUT_SHAPE,
    target=TARGET,
    num_of_bits=NUM_OF_BITS,
    collate_fn=collate_fn2,
    setting=quant_setting,
    device=DEVICE,
    error_report=True,
    skip_export=False,
    export_test_values=False,
    verbose=1,
)

# -----
# Evaluate Quantized Model.
# -----

evaluate_ppq_module_with_imagenet(
    model=quant_ppq_graph,
    imagenet_validation_dir=CALIB_DIR,
    batchsize=BATCH_SIZE,
    device=DEVICE,
    verbose=1,
)

```

量化结果及分析

在采用 TQT 策略对 8-bit 模型进行量化优化的情况下，模型的 Top-1 准确率为 71.775%，相比 ESP32P4 本身的 per-channel 的量化模型（Top-1 准确率为 70.850%），TQT 在 Top-1 精度上提升了 1.306 个百分点，与 float32 模型的精度（71.878%）更加接近，表明即使在严格的 8-bit 量化约束下，仅通过对量化阈值进行可学习优化，TQT 也能够有效缓解量化误差带来的性能下降，使量化模型的精度接近浮点模型水平。

3.11.4 常见问题

如何加速 TQT ?

有 GPU 时，把校准数据放到 GPU 上可加快速度。

先将 `collecting_device` 设为 `cuda`：

```
quant_setting.tqt_optimization_setting.collecting_device = "cuda"
```

再在 `ENABLE_CUDA_KERNEL` 中执行量化：

```

with ENABLE_CUDA_KERNEL():
    quant_ppq_graph = espdl_quantize_onnx(
        onnx_import_file=ONNX_PATH,
        espdl_export_file=ESPDL_MODEL_PATH,

```

(接上页)

```

    calib_data_loader=dataloader,
    calib_steps=32,
    input_shape=[1] + INPUT_SHAPE,
    target=TARGET,
    num_of_bits=NUM_OF_BITS,
    collate_fn=collate_fn,
    setting=quant_setting,
    device=DEVICE,
    error_report=False,
    skip_export=False,
    export_test_values=False,
    verbose=0,
    inputs=None,
)

```

同时适当增大 `block_size` (如 2~4) 可减少块数、缩短总时间, 但过大可能不稳定, 建议先试 4。

可以不修改 `scale` 只训练模型权重吗?

可以。将 `is_scale_trainable` 设置为 `False` 即可。

```
quant_setting.tqt_optimization_setting.is_scale_trainable = False
```

TQT 可以和 `Weight Equalization` 一起使用吗?

可以。特别是对于 `espdet_pico` 这样的模型, 二者可以使用可以进一步提升量化效果。

```

quant_setting.tqt_optimization_setting.equalization = True
quant_setting.tqt_optimization_setting.tqt_optimization = True

```


4.1 Tensor API Reference

Tensor is the fundamental data type in esp-dl, used for storing multi-type data such as int8, int16, float, etc., similar to the tensor in PyTorch. We have implemented some common tensor operations. Please refer to the following APIs for details.

4.1.1 Header File

- `esp-dl/dl/tensor/include/dl_tensor_base.hpp`

4.1.2 Classes

class **ExponentInfo**

Exponent info for per-tensor / per-channel quantization.

Per-tensor: `m_exponent` only, zero heap allocation. Per-channel: heap-allocated `m_exponents` array. Provides operator `int()` so existing `DL_SCALE(tensor->exponent)` code works unchanged.

Public Functions

inline **ExponentInfo** (int exponent = 0)

Construct a per-tensor *ExponentInfo* with a single exponent value.

参数

exponent –The exponent value for per-tensor quantization. Default is 0.

inline **ExponentInfo** (const std::vector<int> &exponents)

Construct an *ExponentInfo* from a vector of exponents.

参数

exponents –Vector of exponent values. If size <= 1, uses per-tensor mode; otherwise uses per-channel mode with heap allocation.

inline **~ExponentInfo** ()

Destroy the *ExponentInfo* object and free heap-allocated memory.

inline **ExponentInfo** (const *ExponentInfo* &other)

Copy constructor.

参数

other –The *ExponentInfo* object to copy from.

inline *ExponentInfo* &**operator=** (const *ExponentInfo* &other)

Copy assignment operator.

参数

other –The *ExponentInfo* object to assign from.

返回

Reference to this object.

inline **ExponentInfo** (*ExponentInfo* &&other) noexcept

Move constructor.

参数

other –The *ExponentInfo* object to move from.

inline *ExponentInfo* &**operator=** (*ExponentInfo* &&other) noexcept

Move assignment operator.

参数

other –The *ExponentInfo* object to move from.

返回

Reference to this object.

inline *ExponentInfo* &**operator=** (int value)

Assign a single integer value as per-tensor exponent.

参数

value –The exponent value to assign.

返回

Reference to this object.

inline int **get** (int ch = -1) const

Get exponent value.

参数

ch –Channel index. -1 (default) returns per-tensor value. $ch \geq 0$ returns per-channel value if available, otherwise per-tensor.

返回

The exponent value for the specified channel or per-tensor exponent.

```
inline operator int () const
```

Implicit conversion to int, returns the per-tensor exponent value.

返回

The per-tensor exponent value.

```
inline bool is_per_channel () const
```

Check if using per-channel quantization.

返回

true if per-channel exponents are stored, false otherwise.

```
inline int channel_size () const
```

Get the number of channels.

返回

Number of channels for per-channel mode, or 1 for per-tensor mode.

```
inline const int *data () const
```

Get pointer to exponent data.

返回

Pointer to per-channel exponents array if available, otherwise pointer to per-tensor exponent.

```
inline bool operator==(const ExponentInfo &other) const
```

Compare two *ExponentInfo* objects for equality.

参数

other –The *ExponentInfo* object to compare with.

返回

true if both have the same exponent values, false otherwise.

```
inline bool operator!=(const ExponentInfo &other) const
```

Compare two *ExponentInfo* objects for inequality.

参数

other –The *ExponentInfo* object to compare with.

返回

true if exponent values differ, false if equal.

```
class TensorBase
```

This class is designed according to PyTorch Tensor. *TensorBase* is required to ensure that the first address are aligned to 16 bytes and the memory size should be a multiple of 16 bytes.

TODO:: Implement more functions

Public Functions

TensorBase (std::vector<int> shape, const void *element, int exponent = 0, dtype_t dtype = DATA_TYPE_FLOAT, bool deep = true, uint32_t caps = MALLOC_CAP_DEFAULT)

Construct a *TensorBase* object.

参数

- **shape** –Shape of tensor
- **element** –Pointer of data
- **exponent** –Exponent of tensor, default is 0
- **dtype** –Data type of element, default is float
- **deep** –True: malloc memory and copy data, false: use the pointer directly
- **caps** –Bitwise OR of MALLOC_CAP_* flags indicating the type of memory to be returned

TensorBase (std::vector<int> shape, const void *element, const std::vector<int> &exponents, dtype_t dtype = DATA_TYPE_FLOAT, bool deep = true, uint32_t caps = MALLOC_CAP_DEFAULT)

Construct a *TensorBase* object with per-channel exponents.

参数

- **shape** –Shape of tensor
- **element** –Pointer of data
- **exponents** –Per-channel exponents
- **dtype** –Data type of element, default is float
- **deep** –True: malloc memory and copy data, false: use the pointer directly
- **caps** –Bitwise OR of MALLOC_CAP_* flags indicating the type of memory to be returned

inline virtual ~**TensorBase** ()

Destroy the *TensorBase* object.

bool **assign** (*TensorBase* *tensor)

Assign tensor to this tensor.

参数

tensor –

返回

true if assign successfully, otherwise false.

bool **assign** (std::vector<int> shape, const void *element, int exponent, dtype_t dtype)

Assign data to this tensor.

参数

- **shape** –
- **element** –
- **exponent** –
- **dtype** –

返回

true if assign successfully, otherwise false.

```
inline int get_size ()
```

Get the size of Tensor.

返回

the size of Tensor.

```
inline int get_aligned_size ()
```

Get the aligned size of Tensor.

返回

the aligned size of Tensor.

```
inline size_t get_dtype_bytes ()
```

Get the dtype size, in bytes.

返回

the size of dtype.

```
inline const char *get_dtype_string ()
```

Get the dtype string of Tensor.

返回

the string of Tensor' s dtype.

```
inline int get_bytes ()
```

Get the bytes of Tensor.

返回

the bytes of Tensor.

```
inline int get_aligned_bytes ()
```

Get the bytes of Tensor.

返回

the bytes of Tensor.

```
inline virtual void *get_element_ptr ()
```

Get data pointer. If *cache(preload data pointer)* is not null, return cache pointer, otherwise return data pointer.

返回

the pointer of Tensor' s data

```
template<typename T>
```

```
inline T *get_element_ptr ()
```

Get data pointer by the specified template. If *cache(preload data pointer)* is not null, return cache pointer, otherwise return data pointer.

返回

the pointer of Tensor' s data

```
TensorBase &set_element_ptr (void *data)
```

Set the data pointer of Tensor.

参数

data –point to data memory

返回

TensorBase& self

```
inline std::vector<int> get_shape ()
```

Get the shape of Tensor.

返回

std::vector<int> the shape of Tensor

```
TensorBase &set_shape (const std::vector<int> shape)
```

Set the shape of Tensor.

参数

shape –the shape of Tensor.

返回

Tensor.

```
inline int get_exponent ()
```

Get the exponent of Tensor.

返回

int the exponent of Tensor

```
inline dtype_t get_dtype ()
```

Get the data type of Tensor.

返回

dtype_t the data type of Tensor

```
inline uint32_t get_caps ()
```

Get the memory flags of Tensor.

返回

uint32_t the memory flags of Tensor

```
TensorBase *reshape (std::vector<int> shape)
```

Change a new shape to the Tensor without changing its data.

参数

shape –the target shape

返回

TensorBase *self

```
template<typename T>
```

```
TensorBase *flip (const std::vector<int> &axes)
```

Flip the input Tensor along the specified axes.

参数

axes –the specified axes

返回

TensorBase& self

```
TensorBase *transpose (TensorBase *input, std::vector<int> perm = {})
```

Reverse or permute the axes of the input Tensor.

参数

- **input** –the input Tensor
- **perm** –the new arrangement of the dims. if perm == {}, the dims arrangement will be reversed.

返回*TensorBase* *self

```
template<typename T>
TensorBase *transpose (T *input_element, std::vector<int> &input_shape, std::vector<int>
                        &input_axis_offset, std::vector<int> &perm)
```

Reverse or permute the axes of the input Tensor.

参数

- **input_element** –the input data pointer
- **input_shape** –the input data shape
- **input_axis_offset** –the input data axis offset
- **perm** –the new arrangement of the dims. if perm == {}, the dims arrangement will be reversed.

返回*TensorBase* *self

```
bool is_same_shape (TensorBase *tensor)
```

Check the shape is the same as the shape of input.

参数

tensor –Input tensor pointer

返回

- true: same shape
- false: not

```
bool equal (TensorBase *tensor, float epsilon = 1e-6, bool verbose = false)
```

Compare the shape and data of two Tensor.

参数

- **tensor** –Input tensor
- **epsilon** –The max error of two element
- **verbose** –If true, print the detail of results

返回

true if two tensor is equal otherwise false

```
TensorBase *slice (const std::vector<int> &start, const std::vector<int> &end, const std::vector<int> &axes =
                  {}, const std::vector<int> &step = {})
```

Produces a slice of the this tensor along multiple axes.

警告: The length of start, end and step must be same as the shape of input tensor

参数

- **start** –Starting indices
- **end** –Ending indices
- **axes** –Axes that starts and ends apply to.
- **step** –Slice step, step = 1 if step is not specified

返回

TensorBase* Output tensor pointer, created by this slice function

```
template<typename T>
TensorBase *pad (T *input_element, const std::vector<int> &input_shape, const std::vector<int> &pads, const
padding_mode_t mode, TensorBase *const_value = nullptr)
```

Pad input tensor.

参数

- **input_element** –Data pointer of input tensor
- **input_shape** –Shape of input tensor
- **pads** –The number of padding elements to add, pads format should be: [x1_begin, x2_begin, ..., x1_end, x2_end, ...]
- **mode** –Supported modes: constant(default), reflect, edge
- **const_value** –(Optional) A scalar value to be used if the mode chosen is constant

返回

Output tensor pointer

```
TensorBase *pad (TensorBase *input, const std::vector<int> &pads, const padding_mode_t mode, TensorBase
*const_value = nullptr)
```

Pad input tensor.

参数

- **input** –Input tensor pointer
- **pads** –Padding elements to add, pads format should be: [x1_begin, x2_begin, ..., x1_end, x2_end, ...]
- **mode** –Supported modes: constant(default), reflect, edge
- **const_value** –(Optional) A scalar value to be used if the mode chosen is constant

返回

Output tensor pointer

```
template<typename T>
bool compare_elements (const T *gt_elements, float epsilon = 1e-6, bool verbose = false)
```

Compare the elements of two Tensor.

参数

- **gt_elements** –The ground truth elements
- **epsilon** –The max error of two element
- **verbose** –If true, print the detail of results

返回

true if all elements are equal otherwise false

```
int get_element_index (const std::vector<int> &axis_index)
```

Get the index of element.

参数

axis_index –The coordinates of element

返回

int the index of element

std::vector<int> **get_element_coordinates** (int index)

Get the coordinates of element.

参数

index –The index of element

返回

The coordinates of element

template<typename **T**>

T **get_element** (int index)

Get a element of Tensor by index.

参数

index –The index of element

返回

The element of tensor

template<typename **T**>

T **get_element** (const std::vector<int> &axis_index)

Get a element of Tensor.

参数

axis_index –The index of element

返回

The element of tensor

void **memset** (int value)

Set a element of Tensor by index.

参数

value –The value of element

void **rand** ()

Fill tensor data with random bytes from hardware RNG.

Uses `esp_fill_random()` to fill the tensor' s internal buffer with random bytes. Requires ESP-IDF (`esp_random.h`).

size_t **set_preload_addr** (void *addr, size_t size)

Set preload address of Tensor.

参数

- **addr** –The address of preload data
- **size** –Size of preload data

返回

The size of preload data

inline virtual void **preload** ()

Preload the data of Tensor.

void **reset_bias_layout** (quant_type_t op_quant_type, bool is_depthwise)

Reset the layout of Tensor.

警告: Only available for Convolution. Don't use it unless you know exactly what it does.

参数

- **op_quant_type** –The quant type of operation
- **is_depthwise** –Whether is depthwise convolution

void **push** (*TensorBase* *new_tensor, int dim)

Push new_tensor to current tensor. The time series dimension size of new tensor must is lesser or equal than that of the current tensor.” .

参数

- **new_tensor** –The new tensor will be pushed
- **dim** –Specify the dimension on which to perform streaming stack pushes

virtual void **print** (bool print_data = false)

print the information of *TensorBase*

参数

- **print_data** –Whether print the data

Public Members

int **size**

size of element including padding

std::vector<int> **shape**

shape of Tensor

dtype_t **dtype**

data type of element

ExponentInfo **exponent**

exponent of element (per-tensor or per-channel)

bool **auto_free**

free element when object destroy

std::vector<int> **axis_offset**

element offset of each axis

void ***data**

data pointer

void ***cache**

cache pointer, used for preload and do not need to free

uint32_t **caps**

flags indicating the type of memory

Public Static Functions

```
static void slice (TensorBase *input, TensorBase *output, const std::vector<int> &start, const std::vector<int>
&end, const std::vector<int> &axes = {}, const std::vector<int> &step = {})
```

Produces a slice along multiple axes.

警告: The length of start, end and step must be same as the shape of input tensor

参数

- **input** –Input Tensor
- **output** –Output Tensor
- **start** –Starting indices
- **end** –Ending indices
- **axes** –Axes that starts and ends apply to.
- **step** –Slice step, step = 1 if step is not specified

4.2 Module API Reference

The Module is the base class for operators in esp-dl, and all operators inherit from this base class. This base class defines the basic interfaces for operators, enabling the model layer to automatically execute operators and manage memory planning.

4.2.1 Header File

- esp-dl/dl/module/include/dl_module_base.hpp

4.2.2 Classes

class **Module**

Base class for module.

Public Functions

Module (const char *name = NULL, module_inplace_t inplace = MODULE_NON_INPLACE, quant_type_t quant_type = QUANT_TYPE_NONE)

Construct a new *Module* object.

参数

- **name** –Name of module.
- **inplace** –Inplace operation mode
- **quant_type** –Quantization type

virtual ~**Module** ()

Destroy the *Module* object. Return resource.

inline virtual std::vector<int> **get_outputs_index** ()

Get the tensor index of this module' s outputs.

返回

Tensor index of model' s tensors

virtual std::vector<std::vector<int>> **get_output_shape** (std::vector<std::vector<int>> &input_shapes) = 0

Calculate output shape by input shape.

参数

input_shapes –Input shapes

返回

outputs shapes

virtual void **forward** (*ModelContext* *context, runtime_mode_t mode = RUNTIME_MODE_AUTO) = 0

Build the module, high-level interface for *Module* layer.

参数

- **context** –*Model* context including all inputs and outputs and other runtime information
- **mode** –Runtime mode, default is RUNTIME_MODE_AUTO

inline virtual void **forward_args** (void *args)

Run the module, Low-level interface for base layer and multi-core processing.

参数

args –ArgsType, arithArgsType, resizeArgsType and so on

inline virtual void **print** ()

print module information

inline virtual void **set_preload_addr** (void *addr, size_t size)

set preload RAM pointer

参数

- **addr** –Internal RAM address, should be aligned to 16 bytes
- **size** –The size of RAM address

inline virtual void **preload** ()

Perform a preload operation.

警告: Not implemented

inline virtual void **reset** ()

reset all state of module, include inputs, outputs and preload cache setting

virtual void **run** (*TensorBase* *input, *TensorBase* *output, runtime_mode_t mode =
RUNTIME_MODE_SINGLE_CORE)

Run the module with single input and single output.

参数

- **input** –Input tensor
- **output** –Output tensor
- **mode** –Runtime mode

virtual void **run** (std::vector<dl::*TensorBase**> inputs, std::vector<dl::*TensorBase**> outputs, runtime_mode_t
mode = RUNTIME_MODE_SINGLE_CORE)

Run the module by inputs and outputs.

参数

- **inputs** –Input tensors
- **outputs** –Output tensors
- **mode** –Runtime mode

Public Members

char ***name**

Name of module.

module_inplace_t **inplace**

Inplace type.

quant_type_t **quant_type**

Quantization type.

std::vector<int> **m_inputs_index**

Tensor index of model' s tensors that used for inputs.

std::vector<int> **m_outputs_index**

Tensor index of model' s tensors that used for outputs.

Public Static Functions

static inline *Module* ***deserialize** (fbs::*FbsModel* *fbs_model, std::string node_name)
create module instance by node serialization information

参数

- **fbs_model** –Flatbuffer’ s model
- **node_name** –The node name in model’ s graph

返回

The pointer of module instance

4.2.3 Header File

- esp-dl/dl/module/include/dl_module_creator.hpp

4.2.4 Classes

class **ModuleCreator**

Singleton class for registering modules.

Public Types

using **Creator** = std::function<*Module**(fbs::*FbsModel**, std::string)>
Module creator function type.

Public Functions

inline void **register_module** (const std::string &op_type, *Creator* creator)

Register a module creator to the module creator map This function allows for the dynamic registration of new module types and their corresponding creator functions at runtime. By associating the module type name with the creator function, the system can flexibly create instances of various modules.

参数

- **op_type** –The module type name, used as the key in the map
- **creator** –The module creator function, used to create modules of a specific type

inline *Module* ***create** (fbs::*FbsModel* *fbs_model, const std::string &op_type, const std::string name)
Create module instance pointer.

参数

- **fbs_model** –Flatbuffer model pointer
- **op_type** –Module/Operator type
- **name** –*Module* name

返回

Module instance pointer

```
inline void register_dl_modules ()
    Pre-register the already implemented modules.

inline void print ()
    Print all modules has been registered.

inline void clear ()
    Clear all modules has been registered.
```

Public Static Functions

```
static inline ModuleCreator *get_instance ()
    Get instance of ModuleCreator by this function. It is only safe method to get instance of ModuleCreator
    because ModuleCreator is a singleton class.
```

返回
ModuleCreator instance pointer

4.3 Model API Reference

This section covers model loading and static memory planning, making it convenient for users to directly load and run ESPDL models.

4.3.1 Header File

- [esp-dl/dl/model/include/dl_model_base.hpp](#)

4.3.2 Macros

```
DL_LOG_INFER_LATENCY_INIT_WITH_SIZE (size)
DL_LOG_INFER_LATENCY_INIT ()
DL_LOG_INFER_LATENCY_START ()
DL_LOG_INFER_LATENCY_END ()
DL_LOG_INFER_LATENCY_PRINT (prefix, key)
DL_LOG_INFER_LATENCY_END_PRINT (prefix, key)
DL_LOG_INFER_LATENCY_ARRAY_INIT_WITH_SIZE (n, size)
DL_LOG_INFER_LATENCY_ARRAY_INIT (n)
DL_LOG_INFER_LATENCY_ARRAY_START (i)
DL_LOG_INFER_LATENCY_ARRAY_END (i)
DL_LOG_INFER_LATENCY_ARRAY_PRINT (i, prefix, key)
DL_LOG_INFER_LATENCY_ARRAY_END_PRINT (i, prefix, key)
```

4.3.3 Classes

class **Model**

Neural Network *Model*.

Public Functions

Model (const char *rodata_address_or_partition_label_or_path, fbs::model_location_type_t location = fbs::MODEL_LOCATION_IN_FLASH_RODATA, int max_internal_size = 0, memory_manager_t mm_type = MEMORY_MANAGER_GREEDY, const uint8_t *key = nullptr, bool param_copy = true)

Create the *Model* object by rodata address or partition label.

参数

- **rodata_address_or_partition_label_or_path** –The address of model data while location is MODEL_LOCATION_IN_FLASH_RODATA. The label of partition while location is MODEL_LOCATION_IN_FLASH_PARTITION. The path of model while location is MODEL_LOCATION_IN_SDCARD.
- **location** –The model location.
- **max_internal_size** –In bytes. Limit the max internal size usage. Only take effect when there's a PSRAM, and you want to alloc memory on internal RAM first.
- **mm_type** –Type of memory manager
- **key** –The key of encrypted model.
- **param_copy** –Set to false to avoid copy model parameters from FLASH to PSRAM. Only set this param to false when your PSRAM resource is very tight. This saves PSRAM and sacrifices the performance of model inference because the frequency of PSRAM is higher than FLASH. Only takes effect when MODEL_LOCATION_IN_FLASH_RODATA(CONFIG_SPIRAM_RODATA not set) or MODEL_LOCATION_IN_FLASH_PARTITION.

Model (const char *rodata_address_or_partition_label_or_path, int model_index, fbs::model_location_type_t location = fbs::MODEL_LOCATION_IN_FLASH_RODATA, int max_internal_size = 0, memory_manager_t mm_type = MEMORY_MANAGER_GREEDY, const uint8_t *key = nullptr, bool param_copy = true)

Create the *Model* object by rodata address or partition label.

参数

- **rodata_address_or_partition_label_or_path** –The address of model data while location is MODEL_LOCATION_IN_FLASH_RODATA. The label of partition while location is MODEL_LOCATION_IN_FLASH_PARTITION. The path of model while location is MODEL_LOCATION_IN_SDCARD.
- **model_index** –The model index of packed models.
- **location** –The model location.
- **max_internal_size** –In bytes. Limit the max internal size usage. Only take effect when there's a PSRAM, and you want to alloc memory on internal RAM first.
- **mm_type** –Type of memory manager
- **key** –The key of encrypted model.

- **param_copy** –Set to false to avoid copy model parameters from FLASH to PSRAM. Only set this param to false when your PSRAM resource is very tight. This saves PSRAM and sacrifices the performance of model inference because the frequency of PSRAM is higher than FLASH. Only takes effect when `MODEL_LOCATION_IN_FLASH_RODATA(CONFIG_SPIRAM_RODATA not set)` or `MODEL_LOCATION_IN_FLASH_PARTITION`.

Model (const char *rodata_address_or_partition_label_or_path, const char *model_name, fbs::model_location_type_t location = fbs::MODEL_LOCATION_IN_FLASH_RODATA, int max_internal_size = 0, memory_manager_t mm_type = MEMORY_MANAGER_GREEDY, const uint8_t *key = nullptr, bool param_copy = true)

Create the *Model* object by rodata address or partition label.

参数

- **rodata_address_or_partition_label_or_path** –The address of model data while location is `MODEL_LOCATION_IN_FLASH_RODATA`. The label of partition while location is `MODEL_LOCATION_IN_FLASH_PARTITION`. The path of model while location is `MODEL_LOCATION_IN_SDCARD`.
- **model_name** –The model name of packed models.
- **location** –The model location.
- **max_internal_size** –In bytes. Limit the max internal size usage. Only take effect when there's a PSRAM, and you want to alloc memory on internal RAM first.
- **mm_type** –Type of memory manager
- **key** –The key of encrypted model.
- **param_copy** –Set to false to avoid copy model parameters from FLASH to PSRAM. Only set this param to false when your PSRAM resource is very tight. This saves PSRAM and sacrifices the performance of model inference because the frequency of PSRAM is higher than FLASH. Only takes effect when `MODEL_LOCATION_IN_FLASH_RODATA(CONFIG_SPIRAM_RODATA not set)` or `MODEL_LOCATION_IN_FLASH_PARTITION`.

Model (fbs::FbsModel *fbs_model, int internal_size = 0, memory_manager_t mm_type = MEMORY_MANAGER_GREEDY)

Create the *Model* object by fbs_model.

参数

- **fbs_model** –The fbs model.
- **internal_size** –Internal ram size, in bytes
- **mm_type** –Type of memory manager

virtual ~**Model** ()

Destroy the *Model* object.

virtual esp_err_t **load** (const char *rodata_address_or_partition_label_or_path, fbs::model_location_type_t location = fbs::MODEL_LOCATION_IN_FLASH_RODATA, const uint8_t *key = nullptr, bool param_copy = true)

Load model graph and parameters from FLASH or sdcard.

参数

- **rodata_address_or_partition_label_or_path** –The address of model data while location is `MODEL_LOCATION_IN_FLASH_RODATA`. The label of partition

while location is `MODEL_LOCATION_IN_FLASH_PARTITION`. The path of model while location is `MODEL_LOCATION_IN_SDCARD`.

- **location** –The model location.
- **key** –The key of encrypted model.
- **param_copy** –Set to false to avoid copy model parameters from FLASH to PSRAM. Only set this param to false when your PSRAM resource is very tight. This saves PSRAM and sacrifices the performance of model inference because the frequency of PSRAM is higher than FLASH. Only takes effect when `MODEL_LOCATION_IN_FLASH_RODATA(CONFIG_SPIRAM_RODATA not set)` or `MODEL_LOCATION_IN_FLASH_PARTITION`.

返回

- `ESP_OK` Success
- `ESP_FAIL` Failed

```
virtual esp_err_t load (const char *rodata_address_or_partition_label_or_path, fbs::model_location_type_t
    location = fbs::MODEL_LOCATION_IN_FLASH_RODATA, int model_index = 0,
    const uint8_t *key = nullptr, bool param_copy = true)
```

Load model graph and parameters from FLASH or sdcard.

参数

- **rodata_address_or_partition_label_or_path** –The address of model data while location is `MODEL_LOCATION_IN_FLASH_RODATA`. The label of partition while location is `MODEL_LOCATION_IN_FLASH_PARTITION`. The path of model while location is `MODEL_LOCATION_IN_SDCARD`.
- **location** –The model location.
- **model_index** –The model index of packed models.
- **key** –The key of encrypted model.
- **param_copy** –Set to false to avoid copy model parameters from FLASH to PSRAM. Only set this param to false when your PSRAM resource is very tight. This saves PSRAM and sacrifices the performance of model inference because the frequency of PSRAM is higher than FLASH. Only takes effect when `MODEL_LOCATION_IN_FLASH_RODATA(CONFIG_SPIRAM_RODATA not set)` or `MODEL_LOCATION_IN_FLASH_PARTITION`.

返回

- `ESP_OK` Success
- `ESP_FAIL` Failed

```
virtual esp_err_t load (const char *rodata_address_or_partition_label_or_path, fbs::model_location_type_t
    location = fbs::MODEL_LOCATION_IN_FLASH_RODATA, const char
    *model_name = nullptr, const uint8_t *key = nullptr, bool param_copy = true)
```

Load model graph and parameters from FLASH or sdcard.

参数

- **rodata_address_or_partition_label_or_path** –The address of model data while location is `MODEL_LOCATION_IN_FLASH_RODATA`. The label of partition while location is `MODEL_LOCATION_IN_FLASH_PARTITION`. The path of model while location is `MODEL_LOCATION_IN_SDCARD`.

- **location** –The model location.
- **model_name** –The model name of packed models.
- **key** –The key of encrypted model.
- **param_copy** –Set to false to avoid copy model parameters from FLASH to PSRAM. Only set this param to false when your PSRAM resource is very tight. This saves PSRAM and sacrifices the performance of model inference because the frequency of PSRAM is higher than FLASH. Only takes effect when `MODEL_LOCATION_IN_FLASH_RODATA(CONFIG_SPIRAM_RODATA` not set) or `MODEL_LOCATION_IN_FLASH_PARTITION`.

返回

- ESP_OK Success
- ESP_FAIL Failed

virtual esp_err_t **load** (fbs::FbsModel *fbs_model)

Load model graph and parameters from Flatbuffers model.

参数

fbs_model –The FlatBuffers model

返回

- ESP_OK Success
- ESP_FAIL Failed

virtual void **build** (size_t max_internal_size, memory_manager_t mm_type = MEMORY_MANAGER_GREEDY, bool preload = false)

Allocate memory for the model.

参数

- **max_internal_size** –In bytes. Limit the max internal size usage. Only take effect when there's a PSRAM, and you want to alloc memory on internal RAM first.
- **mm_type** –Type of memory manager
- **preload** –Whether to preload the model's parameters to internal ram (not implemented yet)

virtual void **run** (runtime_mode_t mode = RUNTIME_MODE_SINGLE_CORE)

Run the model module by module.

参数

mode –Runtime mode.

virtual void **run** (TensorBase *input, runtime_mode_t mode = RUNTIME_MODE_SINGLE_CORE)

Run the model module by module.

参数

- **input** –The model input.
- **mode** –Runtime mode.

virtual void **run** (std::map<std::string, TensorBase*> &user_inputs, runtime_mode_t mode = RUNTIME_MODE_SINGLE_CORE, std::map<std::string, TensorBase*> user_outputs = {})

Run the model module by module.

参数

- **user_inputs** –The model inputs.
- **mode** –Runtime mode.
- **user_outputs** –It’ s for debug to specify the output of the intermediate layer; Under normal use, there is no need to pass a value to this parameter. If no parameter is passed, the default is the graphical output, which can be obtained through *Model::get_outputs()*.

void **minimize** ()

Minimize the model.

esp_err_t **test** ()

Test whether the model inference result is correct. The model should contain test_inputs and test_outputs. Enable export_test_values option in esp-ppq to use this api.

返回

esp_err_t

std::map<std::string, mem_info_t> **get_memory_info** ()

Get memory info.

返回

Memory usage statistics on internal and PSRAM.

std::map<std::string, module_info> **get_module_info** ()

Get module info.

返回

return Type and latency of each module.

void **print_module_info** (const std::map<std::string, module_info> &info, bool sort_module_by_latency = false)

Print the module info obtained by get_module_info function.

参数

- **info** –
- **sort_module_by_latency** –

void **profile_memory** ()

Print model memory summary.

void **profile_module** (bool sort_module_by_latency = false)

Print module info summary. (Name, Type, Latency)

参数

sort_module_by_latency –True The module is printed in latency decreasing sort. False The module is printed in ONNX topological sort.

void **profile** (bool sort_module_by_latency = false)

Combination of profile_memory & profile_module.

参数

sort_module_by_latency –True The module is printed in latency decreasing sort. False The module is printed in ONNX topological sort.

virtual std::map<std::string, *TensorBase**> &**get_inputs** ()

Get inputs of model.

返回

The map of model input’ s name and TensorBase*

virtual *TensorBase* ***get_input** ()

Get the only input of model.

返回

*TensorBase**

virtual *TensorBase* ***get_input** (const std::string &name)

Get input of model by name.

参数

name –input name

返回

*TensorBase**

virtual *TensorBase* ***get_intermediate** (const std::string &name)

Get intermediate *TensorBase* of model.

备注: When using memory manager, the content of *TensorBase*' s data may be overwritten by the outputs of other

参数

name –The name of intermediate Tensor. operators.

返回

The intermediate *TensorBase**.

virtual std::map<std::string, *TensorBase**> &**get_outputs** ()

Get outputs of model.

返回

The map of model output' s name and *TensorBase**

virtual *TensorBase* ***get_output** ()

Get the only output of model.

返回

*TensorBase**

virtual *TensorBase* ***get_output** (const std::string &name)

Get output of model by name.

参数

name –output name

返回

*TensorBase**

std::string **get_metadata_prop** (const std::string &key)

Get the model' s metadata prop.

参数

key –The key of metadata prop

返回

The value of metadata prop

virtual void **print** ()

Print the model.

virtual void **reset** ()

Reset the model. This function resets the model context and all modules, but does not reload the model or rebuild the execution plan.

inline virtual fbs::*FbsModel* ***get_fbs_model** ()

Get the fbs model instance.

返回

fbs::FbsModel *

4.3.4 Header File

- esp-dl/dl/model/include/dl_model_context.hpp

4.3.5 Macros

CONTEXT_PARAMETER_OFFSET

Offset for parameter tensors

4.3.6 Classes

class **ModelContext**

Model Context class including variable tensors and parameters.

Public Functions

inline **ModelContext** ()

Constructor for *ModelContext*. Initializes the PSRAM and internal root pointers to nullptr.

inline **~ModelContext** ()

Destructor for *ModelContext*. Clears all resources and tensors.

int **add_tensor** (const std::string name, bool is_paramter = false, *TensorBase* *tensor = nullptr)

Adds a tensor to the parameter or variable list.

参数

- **name** –The name of the tensor.
- **is_paramter** –Whether the tensor is a parameter (default: false).
- **tensor** –Pointer to the *TensorBase* object (default: nullptr).

返回

int Returns the index of the added tensor.

int **push_back_tensor** (*TensorBase* *tensor, bool is_paramter = false)

Push back a tensor.

参数

- **tensor** –Pointer to the *TensorBase* object.
- **is_paramter** –Whether the tensor is a parameter (default: false).

返回

int Returns the index of the added tensor.

void **update_tensor** (int index, *TensorBase* *tensor)

Updates the tensor at the specified index.

参数

- **index** –The index of the tensor to update.
- **tensor** –Pointer to the new *TensorBase* object.

TensorBase ***get_tensor** (int index)

Gets the tensor by its index.

参数

index –The index of the tensor.

返回

*TensorBase** Returns the pointer to the *TensorBase* object, or nullptr if the index is invalid.

TensorBase ***get_tensor** (const std::string &name)

Gets the tensor by its name.

参数

name –The name of the tensor.

返回

*TensorBase** Returns the pointer to the *TensorBase* object, or nullptr if the name is not found.

int **get_tensor_index** (const std::string &name)

Gets the tensor index by its name.

参数

name –The name of the tensor.

返回

int Returns index if the name is found, else -1

int **get_variable_index** (const std::string &name)

Gets the variable tensor index by its name.

参数

name –The name of the tensor.

返回

int Returns index if the name is found and is variable tensor, else -1

inline int **get_variable_count** ()

Gets the count of variable tensors.

返回

int Returns the number of variable tensors.

inline int **get_parameter_count** ()

Gets the count of parameter tensors.

返回

int Returns the number of parameter tensors.

bool **root_alloc** (size_t internal_size, size_t psram_size, int alignment = 16)

Allocates memory for PSRAM and internal roots.

参数

- **internal_size** –The size of the internal memory in bytes.
- **psram_size** –The size of the PSRAM memory in bytes.
- **alignment** –The alignment of the memory in bytes.

返回

Bool Return true if the allocation is successful, false otherwise.

inline void ***get_psram_root** ()

Gets the pointer to the PSRAM root.

返回

Void* Returns the pointer to the PSRAM root.

inline void ***get_internal_root** ()

Gets the pointer to the internal root.

返回

Void* Returns the pointer to the internal root.

size_t **get_parameter_memory_size** (mem_info_t &mem_info, bool copy)

Gets the size of the parameters in bytes.

参数

- **mem_info** –The size of the memory used by the parameters in bytes, filtered by copy option.
- **copy** –Filter the parameters by auto_free.

返回

size_t Returns the total size of the parameters memory in bytes.

size_t **get_variable_memory_size** (mem_info_t &mem_info)

Get the variable memory size object.

参数

mem_info –The size of the memory used by the variables in bytes.

返回

size_t Returns the total size of the variables memory in bytes.

inline void **root_free** ()

Frees the memory allocated for PSRAM and internal roots. This function ensures proper cleanup of allocated memory.

inline void **root_reset** ()

Resets the context by clearing variables, parameters, and name-to-index map. This is used to reset the state of the context for a new inference or after an inference.

```
inline void minimize ()
```

Minimizes the context by clearing the name-to-index map. This is used to free unnecessary intermediate variables during the inference.

```
inline void clear ()
```

Clears all resources and tensors in the context. This includes clearing variables, parameters, name-to-index map, and freeing memory.

Public Members

```
std::vector<TensorBase*> m_variables
```

Variable tensors of model, the first one is nullptr

```
std::vector<TensorBase*> m_parameters
```

Parameters of model, the first one is nullptr

4.3.7 Header File

- `esp-dl/dl/model/include/dl_memory_manager.hpp`

4.3.8 Classes

class **MemoryManagerBase**

Memory manager base class, each model has its own memory manager TODO: share memory manager with different models.

Subclassed by `dl::MemoryManagerGreedy`

Public Functions

```
inline MemoryManagerBase (int alignment = 16)
```

Construct a new Memory Manager Base object.

参数

alignment –Memory address alignment

```
inline virtual ~MemoryManagerBase ()
```

Destroy the MemoryManager object. Return resource.

```
virtual bool alloc (fbs::FbsModel *fbs_model, std::vector<dl::module::Module*> &execution_plan,  
                  ModelContext *context) = 0
```

Allocate memory for each tensor, include all input and output tensors.

参数

- **fbs_model** –FlatBuffer' s *Model*
- **execution_plan** –Topological sorted module list
- **context** –*Model* context

返回

Bool Return true if the allocation is successful, false otherwise.

Public Members**int alignment**

The root pointer needs to be aligned must be a power of two

class **TensorInfo**

Tensor info, include tensor name, shape, dtype, size, time range and call times, which is used to plan model memory.

Public Functions

TensorInfo (std::string &name, int time_begin, int time_end, std::vector<int> shape, dtype_t dtype, int exponent, bool is_internal = false)

Construct a new Tensor Info object.

参数

- **name** –Tensor name
- **time_begin** –Tensor lifetime begin
- **time_end** –Tensor lifetime end
- **shape** –Tensor shape
- **dtype** –Tensor dtype
- **exponent** –Tensor exponent
- **is_internal** –Is tensor in internal RAM or not

inline **~TensorInfo** ()

Destroy the Tensor Info object.

void **set_inplace_leader_tensor** (*TensorInfo* *tensor)

Set the inplace leader tensor object.

参数

tensor –Inplace leader tensor

inline void **set_inplace_follower_dirty_tensor** (*TensorInfo* *tensor)

Set the inplace follower dirty tensor object.

参数

tensor –Inplace follower dirty tensor

inline void **set_inplace_follower_clean_tensor** (*TensorInfo* *tensor)

Set the inplace follower clean tensor object.

参数

tensor –Inplace follower clean tensor

```
inline std::pair<TensorInfo*, TensorInfo*> get_inplace_follower_tensor ()
```

Get the inplace follower tensor object.

返回

std::pair<TensorInfo *, TensorInfo *>

```
void update_time (int new_time)
```

Update Tensor lifetime.

参数

new_time –new tensor lifetime

```
TensorBase *create_tensor (void *internal_root, void *psram_root)
```

Create a *TensorBase* object according to *TensorInfo*.

参数

- **internal_root** –Internal RAM root pointer
- **psram_root** –PSRAM root pointer

返回

TensorBase*

```
inline bool is_inplaced ()
```

Is inplaced or not.

返回

true if inplaced else false

```
inline uint32_t get_offset ()
```

Get the tensor offset.

返回

uint32_t

```
inline void set_offset (uint32_t offset)
```

Set the tensor offset.

参数

offset –

```
inline uint32_t get_internal_offset ()
```

Get the internal offset.

返回

uint32_t

```
inline bool get_internal_state ()
```

Get the internal state.

返回

true if is internal else false

```
inline void set_internal_state (bool is_internal)
```

Set the internal state.

参数

is_internal –

```
inline void set_internal_offset (uint32_t offset)
```

Set the internal offset.

参数

offset –

```
inline int get_time_end ()
```

Get the lifetime end.

返回

int

```
inline int get_time_begin ()
```

Get the lifetime begin.

返回

int

```
inline size_t get_size ()
```

Get the tensor size.

返回

size_t

```
inline std::string get_name ()
```

Get the tensor name.

返回

std::string

```
inline std::vector<int> get_shape ()
```

Get the tensor shape.

返回

std::vector<int>

```
inline void print ()
```

print tensor info

class **MemoryChunk**

Memory chunk, include size, is free, offset, alignment and tensor, which is used to simulate memory allocation.

Public Functions

MemoryChunk (size_t size, int is_free, int alignment = 16)

Construct a new Memory Chunk object.

参数

- **size** –Memory chunk size
- **is_free** –Whether free or not
- **alignment** –Memory chunk alignment

MemoryChunk (*TensorInfo* *tensor, int alignment = 16)

Construct a new Memory Chunk object.

参数

- **tensor** – *TensorInfo*
- **alignment** – Memory chunk alignment

inline ~**MemoryChunk** ()

Destroy the Memory Chunk object.

MemoryChunk ***merge_free_chunk** (*MemoryChunk* *chunk)

Merge continuous free chunk.

参数

chunk –

返回

*MemoryChunk**

MemoryChunk ***insert** (*TensorInfo* *tensor)

Insert tensor into free chunk.

参数

tensor –

返回

*MemoryChunk**

MemoryChunk ***extend** (*TensorInfo* *tensor)

Extend free chunk and insert tensor.

参数

tensor –

返回

*MemoryChunk**

inline void **free** ()

Free memory chunk, set `is_free` to true and set `tensor` to nullptr.

Public Members

size_t **size**

Memory chunk size

bool **is_free**

Whether memory chunk is free or not

int **offset**

Offset relative to root pointer

int **alignment**

Memory address alignment

TensorInfo ***tensor**

Info of the tensor which occupies the memory

4.3.9 Header File

- `esp-dl/dl/model/include/dl_memory_manager_greedy.hpp`

4.3.10 Classes

class **MemoryManagerGreedy** : public `dl::MemoryManagerBase`

Greedy memory manager that allocates memory for tensors in execution order, prioritizing internal RAM allocation first.

Public Functions

inline **MemoryManagerGreedy** (int max_internal_size, int alignment = 16)

Constructs a greedy memory manager with specified constraints.

参数

- **max_internal_size** –Maximum allowed internal RAM usage in bytes
- **alignment** –Memory address alignment requirement (default: 16 bytes)

inline **~MemoryManagerGreedy** ()

Destructor that releases all managed memory resources.

virtual bool **alloc** (`fbs::FbsModel` *fbs_model, `std::vector<dl::module::Module*>` &execution_plan, `ModelContext` *context)

Allocates memory for all network tensors following greedy strategy.

参数

- **fbs_model** –FlatBuffer model containing network architecture
- **execution_plan** –Execution graph ordered by computation dependencies
- **context** –Device-specific runtime configuration

返回

bool True if successful allocation, false if memory insufficient

void **free** ()

Releases all allocated memory including tensor buffers and memory pools.

4.4 Fbs API Reference

The esp-dl model utilizes FlatBuffers to store information about parameters and the computation graph. Taking into account the encryption requirements of some models, this part has not been open-sourced. However, we provide a set of APIs to facilitate users in loading and parsing esp-dl models.

4.4.1 Header File

- `esp-dl/fbs_loader/include/fbs_loader.hpp`

4.4.2 Classes

class **FbsLoader**

Class for parser the flatbuffers.

Public Functions

FbsLoader (const char *rodata_address_or_partition_label_or_path = nullptr, model_location_type_t location = MODEL_LOCATION_IN_FLASH_RODATA)

Construct a new *FbsLoader* object.

参数

- **rodata_address_or_partition_label_or_path** –The address of model data while location is MODEL_LOCATION_IN_FLASH_RODATA. The label of partition while location is MODEL_LOCATION_IN_FLASH_PARTITION. The path of model while location is MODEL_LOCATION_IN_SDCARD.
- **location** –The model location.

~FbsLoader ()

Destroy the *FbsLoader* object.

FbsModel ***load** (const uint8_t *key = nullptr, bool param_copy = true)

Load the model. If there are multiple sub-models, the first sub-model will be loaded.

参数

- **key** –NULL or a 128-bit AES key, like {0x00, 0x01, 0x02, 0x03, 0x04, 0x05, 0x06, 0x07, 0x08, 0x09, 0x0a, 0x0b, 0x0c, 0x0d, 0x0e, 0x0f}
- **param_copy** –Set to false to avoid copy model parameters from FLASH to PSRAM. Only set this param to false when your PSRAM resource is very tight. This saves PSRAM and sacrifices the performance of model inference because the frequency of PSRAM is higher than FLASH. Only takes effect when MODEL_LOCATION_IN_FLASH_RODATA(CONFIG_SPIRAM_RODATA not set) or MODEL_LOCATION_IN_FLASH_PARTITION.

返回

Return nullptr if loading fails. Otherwise return the pointer of *FbsModel*.

FbsModel ***load** (const int model_index, const uint8_t *key = nullptr, bool param_copy = true)

Load the model by model index.

参数

- **model_index** –The index of model.
- **key** –NULL or a 128-bit AES key, like {0x00, 0x01, 0x02, 0x03, 0x04, 0x05, 0x06, 0x07, 0x08, 0x09, 0x0a, 0x0b, 0x0c, 0x0d, 0x0e, 0x0f}.

- **param_copy** –Set to false to avoid copy model parameters from FLASH to PSRAM. Only set this param to false when your PSRAM resource is very tight. This saves PSRAM and sacrifices the performance of model inference because the frequency of PSRAM is higher than FLASH. Only takes effect when MODEL_LOCATION_IN_FLASH_RODATA(CONFIG_SPIRAM_RODATA not set) or MODEL_LOCATION_IN_FLASH_PARTITION.

返回

Return nullptr if loading fails. Otherwise return the pointer of *FbsModel*.

FbsModel ***load** (const char *model_name, const uint8_t *key = nullptr, bool param_copy = true)

Load the model by model name.

参数

- **model_name** –The name of model.
- **key** –NULL or a 128-bit AES key, like {0x00, 0x01, 0x02, 0x03, 0x04, 0x05, 0x06, 0x07, 0x08, 0x09, 0x0a, 0x0b, 0x0c, 0x0d, 0x0e, 0x0f}
- **param_copy** –Set to false to avoid copy model parameters from FLASH to PSRAM. Only set this param to false when your PSRAM resource is very tight. This saves PSRAM and sacrifices the performance of model inference because the frequency of PSRAM is higher than FLASH. Only takes effect when MODEL_LOCATION_IN_FLASH_RODATA(CONFIG_SPIRAM_RODATA not set) or MODEL_LOCATION_IN_FLASH_PARTITION.

返回

Return nullptr if loading fails. Otherwise return the pointer of *FbsModel*.

int **get_model_num** ()

Get the number of models.

返回

The number of models

void **list_models** ()

List all model' s name.

const char ***get_model_location_string** ()

Get the model location string.

返回

The model location string.

4.4.3 Header File

- esp-dl/fbs_loader/include/fbs_model.hpp

4.4.4 Classes

class **FbsModel**

Flatbuffer model object.

Public Functions

FbsModel (const void *data, size_t size, model_location_type_t location, bool encrypt, bool rodata_move, bool auto_free, bool param_copy)

Construct a new *FbsModel* object.

参数

- **data** –The data of model flatbuffers.
- **size** –The size of model flatbuffers in bytes.
- **location** –The location of model flatbuffers.
- **encrypt** –Whether the model flatbuffers is encrypted or not.
- **rodata_move** –Whether the model flatbuffers is moved from FLASH rodata to PSRAM.
- **auto_free** –Whether to free the model flatbuffers data when destroy this class instance.
- **param_copy** –Whether to copy the parameter in flatbuffers.

~FbsModel ()

Destroy the *FbsModel* object.

void **print** ()

Print the model information.

std::vector<std::string> **topological_sort** ()

Return vector of node name in the order of execution.

返回

topological sort of node name.

esp_err_t **get_operation_attribute** (std::string node_name, std::string attribute_name, int &ret_value)

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **attribute_name** –The name of attribute.
- **ret_value** –The attribute value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

esp_err_t **get_operation_attribute** (std::string node_name, std::string attribute_name, float &ret_value)

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **attribute_name** –The name of attribute.

- **ret_value** –The attribute value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

esp_err_t **get_operation_attribute** (std::string node_name, std::string attribute_name, std::string &ret_value)

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **attribute_name** –The name of attribute.
- **ret_value** –The attribute value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

esp_err_t **get_operation_attribute** (std::string node_name, std::string attribute_name, std::vector<int> &ret_value)

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **attribute_name** –The name of attribute.
- **ret_value** –The attribute value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

esp_err_t **get_operation_attribute** (std::string node_name, std::string attribute_name, std::vector<float> &ret_value)

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **attribute_name** –The name of attribute.
- **ret_value** –The attribute value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

esp_err_t **get_operation_attribute** (std::string node_name, std::string attribute_name, dl::quant_type_t &ret_value)

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **attribute_name** –The name of attribute.
- **ret_value** –The attribute value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

```
esp_err_t get_operation_attribute (std::string node_name, std::string attribute_name,
                                  dl::activation_type_t &ret_value)
```

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **attribute_name** –The name of attribute.
- **ret_value** –The attribute value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

```
esp_err_t get_operation_attribute (std::string node_name, std::string attribute_name,
                                  dl::resize_mode_t &ret_value)
```

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **attribute_name** –The name of attribute.
- **ret_value** –The attribute value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

```
esp_err_t get_operation_attribute (std::string node_name, std::string attribute_name, dl::TensorBase
                                  *&ret_value)
```

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **attribute_name** –The name of attribute.
- **ret_value** –The attribute value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

```
esp_err_t get_operation_input_shape (std::string node_name, int index, std::vector<int> &ret_value)
```

Get operation input shape.

参数

- **node_name** –The name of operation.
- **index** –The index of inputs
- **ret_value** –Return shape value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

```
esp_err_t get_operation_output_shape (std::string node_name, int index, std::vector<int>
                                       &ret_value)
```

Get operation output shape.

参数

- **node_name** –The name of operation.
- **index** –The index of outputs
- **ret_value** –Return shape value.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

esp_err_t **get_operation_inputs_and_outputs** (std::string node_name, std::vector<std::string> &inputs, std::vector<std::string> &outputs)

Get the attribute of node.

参数

- **node_name** –The name of operation.
- **inputs** –The vector of operation inputs.
- **outputs** –The vector of operation outputs.

返回

esp_err_t Return ESP_OK if get successfully. Otherwise return ESP_FAIL.

std::string **get_operation_type** (std::string node_name)

Get operation type, “Conv” , “Linear” etc.

参数

node_name –The name of operation

返回

The type of operation.

dl::TensorBase* **get_operation_parameter** (std::string node_name, int index = 1, uint32_t caps = MALLOC_CAP_DEFAULT)

Return if the variable is a parameter.

参数

- **node_name** –The name of operation
- **index** –The index of the variable
- **caps** –Bitwise OR of MALLOC_CAP_* flags indicating the type of memory to be returned

返回

dl::TensorBase*

dl::TensorBase* **get_operation_lut** (std::string node_name, uint32_t caps = MALLOC_CAP_DEFAULT, std::string attribute_name = ”lut”)

Get LUT(Look Up Table) if the operation has LUT.

参数

- **node_name** –The name of operation
- **caps** –Bitwise OR of MALLOC_CAP_* flags indicating the type of memory to be returned
- **attribute_name** –The name of LUT attribute

返回

dl::TensorBase*

bool **is_parameter** (std::string name)

return true if the variable is a parameter

参数

name –Variable name

返回

true if the variable is a parameter else false

const void ***get_tensor_raw_data** (std::string tensor_name)

Get the raw data of FlatBuffers::DL::Tensor.

参数

tensor_name –The name of Tensor.

返回

uint8_t * The pointer of raw data.

dl::dtype_t **get_tensor_dtype** (std::string tensor_name)

Get the element type of tensor tensor.

参数

tensor_name –The tensor name.

返回

FlatBuffers::DL::TensorDataType

std::vector<int> **get_tensor_shape** (std::string tensor_name)

Get the shape of tensor.

参数

tensor_name –The name of tensor.

返回

std::vector<int> The shape of tensor.

std::vector<int> **get_tensor_exponents** (std::string tensor_name)

Get the exponents of tensor.

警告: When quantization is PER_CHANNEL, the size of exponents is same as out_channels. When quantization is PER_TENSOR, the size of exponents is 1.

参数

tensor_name –The name of tensor.

返回

The exponents of tensor.

dl::dtype_t **get_value_info_dtype** (std::string var_name)

Get the element type of value_info.

参数

var_name –The value_info name.

返回

dl::dtype_t

`std::vector<int> get_value_info_shape (std::string var_name)`

Get the shape of value_info.

参数

var_name –The value_info name.

返回

the shape of value_info.

`int get_value_info_exponent (std::string var_name)`

Get the exponent of value_info. Only support PER_TENSOR quantization.

参数

var_name –The value_info name.

返回

the exponent of value_info

`const void *get_test_input_tensor_raw_data (std::string tensor_name)`

Get the raw data of test input tensor.

参数

tensor_name –The name of test input tensor.

返回

uint8_t * The pointer of raw data.

`const void *get_test_output_tensor_raw_data (std::string tensor_name)`

Get the raw data of test output tensor.

参数

tensor_name –The name of test output tensor.

返回

uint8_t * The pointer of raw data.

`dl::TensorBase *get_test_input_tensor (std::string tensor_name)`

Get the test input tensor.

参数

tensor_name –The name of test input tensor.

返回

The pointer of tensor.

`dl::TensorBase *get_test_output_tensor (std::string tensor_name)`

Get the test output tensor.

参数

tensor_name –The name of test output tensor.

返回

The pointer of tensor.

`std::vector<std::string> get_test_outputs_name ()`

Get the name of test outputs.

返回

the name of test outputs

`std::vector<std::string> get_graph_inputs ()`

Get the graph inputs.

返回

the name of inputs

`std::vector<std::string> get_graph_outputs ()`

Get the graph outputs.

返回

the name of ounputs

void `clear_map ()`

Clear all map.

void `load_map ()`

Load all map.

`std::string get_model_name ()`

Get the model name.

返回

the name of model

`int64_t get_model_version ()`

Get the model version.

返回

The version of model

`std::string get_model_doc_string ()`

Get the model doc string.

返回

The doc string of model

`std::string get_model_metadata_prop (const std::string &key)`

Get the model' s metadata prop.

参数

key –The key of metadata prop

返回

The value of metadata prop

void `get_model_size (size_t *internal_size, size_t *psram_size, size_t *psram_rodata_size, size_t *flash_size)`

Get the model size.

参数

- **internal_size** –Flatbuffers model internal RAM usage
- **psram_size** –Flatbuffers model PSRAM usage
- **psram_rodata_size** –Flatbuffers model PSRAM rodate usage. If `CONFIG_SPIRAM_RODATA` option is on, \ Flatbuffers model in FLASH rodata will be copied to PSRAM
- **flash_size** –Flatbuffers model FLASH usage

Public Members

bool `m_param_copy`
copy flatbuffers param or not.