April 9th 2019 - 3PM PDT

- Dial-in info and Edge WG discussion thread
- Call attendees
    -
- Agenda
    - ros2-tensorflow 2 package updated by author Alberto (**@alsora**) at iRobot
        - Works with: ros2 eloquent, tensorflow 2.0, ubuntu 18.04
        - The dockerfile has been updated as well, so you can use it to get a system with all the dependencies.
        - Exposes some bugs in eloquent related to displaying images, more info here
    - ros2 cv_bridge, Lewis (**@LewisLiuPub**) at Intel discussed with maintainer Ethan Gao, and tested **the guideline**, it works and fits the latest ROS2 system. It likely needn't update for the general use. Furthermore, Lewis updated its code to support OpenCV 4.x and submitted **one PR**, which is under code review now.
    -
    -
    - hands/face
- Actions
    -

March 26th 2019 - 3PM PDT

- Dial-in info and Edge WG discussion thread
- Call attendees
    - Joe Speed
    - Steve
    - Katherine
    - Lewis
    - Alex Tyshka
    - Edward
    - Zhen Ju
    - Amit Goel
- Agenda
    - Please list packages using ML here https://discourse.ros.org/t/ros-2-packages-using-ml/13137/7
    - ML for perception, planners, controllers, et al. (Steve?)
        - Pick right areas
        - Make it useful … (watchdogs, autotraining, et al.)

        - General perception is a good place to start …
            - repo w detailed instructions re how to go from zero to deployed
    - Amit - looking at how NVIDIA could contribute …
        - semantic understanding of world, really good models that run in real-time.
        - Interested in standardized messages for perception, can create sample nodes for that
        - Bounding boxes, pixel level segmentation, pose estimation, et al.
        - Did some work on ROS1 nodes for perception
        - Q: what work has been done for perception messages, but what's been used? …
        - Q: OpenCV is often the starting point, but eventually pipeline. Running everything in OpenCV can limit you when you get to performance optimization … resizing not accelerated, AI not accelerated, etc. Looking at how to provide bindings under the hood, but also option to use independent. Making use of gstreamer … OpenCV to gstreamer binding to integrate DeepStream?
        - Q: how is data exchanged along the data pipeline … w/o too many copies
            - Katherine: partially solved, zero copy RMWs. Good to accelerate w HW though should also focus on things that lower memory, expedite processing, etc. People throwing GPU cards at problem vs optimizating it. Need tractable for safety applications
            - Steve - little machine learning bits inside of larger algorithms
            - Katherine: no real standard for going from pixel to semantic mapping. "Roll your own message" makes difficult for novice users. **Think about more extensible messaging format, json schema in ROS message …** training in simulation, how to do that labeling. **(missing tooling for collecting & training data)**
            - Steve - lacking step-by-step zero to deployed
            - Katherine - challenge is getting reasonably well labeled data set together. Not the training & deployment.
            - Amit - is there value in pre-trained models in ROS? (model zoo(s))?
                - Robot tasks are often quite specific …
                - Models would be helpful but isn't limiting factor
                - Has some datasets created for JetBot, robocars.
                - Has simulator to create photo-realistic data for a sensor in a specific place. Most is done using simulation data. It already has ground truth, is labeled data.
                - **Happy to provide synthetic data, whatever required camera type, orientation, et al.**
            - Katherine: **action: let's get a couple example nodes out there.** Get some core utilities for perception handled by a vendor …

        - Amit has sample ROS node for perception. **Action: Can create, manage and keep it up to date.** https://docs.nvidia.com/isaac/isaac/apps/samples/navigation_rosbridge/doc/ros.html https://github.com/dusty-nv/ros_deep_learning
        - Katherine: want C++ and python which is what most prototype this.
- Alex -
    - Gap between open source & commercial, is working with TensorRT
    - **More focus on training**. Not just "take from model zoo". How to create model, retrain model.
    - Make sure your custom model works on accelerators. Object detection models have custom operations, some are supported, some aren't.
    - Python great for prototyping but for deterministic want C/C++ for reliable system
    - Is open to more coding over summer. Possibly open source work DataSpeed is doing.
    - Alex mentioned in discourse that GPU acceleration in OpenCV … frustrated in ROS1, must useful thing to do w/o copying to CPU. But have to rebuilt 30 ros packages from source if you customize OpenCV. **Make OpenCV GPU more accessible.**
- Joe's folks …
    - There are two directions we could contribute:
      1. Create the project like https://github.com/alsora/ros2-tensorflow
      We could also add ros2-openvino, ros2-tensorRT…..
      (don't know)
      2. Join or clone the project like https://github.com/microsoft/onnxruntime
      Alex - ONNX seems promising … portable
      and ask if we could develop the ROS2 interface/adapter for it.
    - A clear interface design and definition are required in both direction.
    - contribute to ML performance profiling test
- Actions
    - Joe following up Alberto re whether to bring this forward into current release alsora/ **ros2-tensorflow** 1 (bouncy)
    - Katherine: force Foxy to OpenCV 4.x. **YES**.
    - Update cv-bridge to work with OpenCV 4.2? Find someone to do it (possible Joe's folks, Edward …) IMage pipeline already going to 4.2 per Steve
    - Lewis: Update ros2 ~~wiki~~ to show CV bridge has been ported ROS2, Katherine can merge into https://index.ros.org/doc/ros2/ https://index.ros.org/p/cv_bridge/github-ros-perception-vision_opencv/#eloquent-overview

        - ?: need documentation covering broad brush strokes covering all ROS2 and how these are used in this process
        - (always send meeting minutes via mail and/or discourse, not everyone has access to google doc.)
        - (Joe: add Cam to invite)

March 17th 2019 - 7AM PDT

- Dial-in info and Edge WG discussion thread
- Call attendees
    - Joe Speed
    - Katherine Scott
    - Jeremy Adams
    - Harold Yang
    - Lewis Liu
    - Liyou Zhou
    - Paul Lin
    - Christoph Hellmann Santos
    - Zhen Ju
    - Dhawal Sharda
- Agenda
    - Please list packages using ML here https://discourse.ros.org/t/ros-2-packages-using-ml/13137/7
    - Actions
        - Joe: Check w Steve re Intel OpenVINO for Nav2, any of that making it in
            -
        - Joe: follow up with Alberto re whether to bring this forward into current release alsora/ **ros2-tensorflow** 1 (bouncy)
        - Katherine: confirm latest stable OpenCV 4.x …probably 4.2.
            - CONFIRMED 4.2 this ships on Ubuntu Focal. OpenCV will be stock for both Foxy and Noetic
        - Katherine: force Foxy to OpenCV 4.x, will bring up with Michael Carroll. Michael & William, Ethan (per Jeremy)
        - Harold: will get the OpenCV contact(s), we'll ask OpenCV maintainer to make Ubuntu apt gettable of that latest stable … as alternative version to not break current OpenCV 3.x users
        - Kyle?: if we need Canonical to get involved with updating OpenCV apt
        - ?: Update "CV bridge" to work with that OpenCV 4.x. Ask nicely … or find someone to pay for it
        - Lewis: Update ros2 wiki to show CV bridge has been ported ROS2 https://index.ros.org/p/cv_bridge/github-ros-perception-vision_opencv/#eloquent-overview

        - ?: need documentation covering broad brush strokes covering all ROS2 and how these are used in this process

- Minutes
    - CHS
        - openvino doesn't support reinforcement learning, right?.
        - Using ___ to do reinforcement learning, is lighter / faster than gazebo.
        - Gazebo not best for learning, but what about ignition?
    - Jeremy
        - ROS1 used wiki, ROS2 moved to index. If you're familiar with ROS1 you may not know where to find things in ROS2 documentation
        - OpenCV has improved CUDA, OpenVINO, etc support
        - OpenCV ROS2 support for Eloquent branch https://github.com/ros-perception/vision_opencv/tree/eloquent
    - Katherine
        - General intent for ignition is to refactor into modular libraries instead of monolithic. Probably needs active contribution.
        - Mentioned SageMaker has made many things so much easier.
        - Is more of a knowledge problem than a technology problem
        - OpenCV
    - Joe - generic camera + sagemaker + openvino service created by AWS+Intel+ADLINK, should contribute
    - include packages that use OpenCV? there is much ML contribution to OpenCV. Though much of what I see is in ros2/demos
        - Katherine
            - OpenCV is the way to go for 90% of image primitives
            - Should be trivial shim to include ML libraries & acceleration
            - Issue is folks keeping it secret sause
            - Need "here's a node for a good camera, actually calibrated" put model in there, do bit of touchup
            - Models not generalized well … must collect data in the operating domain e.g. in the factory, from that POV, on the camera that'll be used
            - In near future won't have models that generalize well.
        - Steve
            - Models trained for cameras at eye level … which is not mobile robots POV
    - Lewis
        - "CV bridge" in ROS is senstive to OpenCV version
        - OpenVINO, contributed ~13k lines of code
        - Contributed smart grasp for MoveIt (MoveIt2?)

        - “Object analytics” using OpenCV + OpenVINO to detect 3D localization for objects. Tracking & detection. Applicable for fixed camera & mobile camera
        - Most ML work is being done for latest OpenCV 4.x. But Ubuntu is supporting 3.x?
            - Katherine - need apt gettable latest “good” release of OpenCV, 4.2?

February 27th 2019 6AM PDT Edge AI WG call

- Dial-in info and WG discussion thread
- [Joe] I should have mentioned that today’s call is recorded for those that couldn’t attend and my own note taking. If that’s a problem tell me and I will delete the recording.
- Call Attendees
    a. Joe Speed (host & moderator)
    b. Christoph Hellmann Santos - Fraunhofer IPA using ML in production and possibly agriculture - ROS Industrial (github: ipa-cmh - discourse: c_h_s)
    c. Sumandeep Banerjee- looking for vision solutions in ROS using latest ML, easy, portable
    d. Liyou Zhou (Arm) - observe, how to help, how to get more adoption on Arm platforms
    e. Geoff Biggs (TierIV) - future architectures for autonomous driving
    f. Jeremy Adams (OCI) - DDS, AI, IoT. Longtime ROS user. ROS-Ag! Learning about how this all ties together.
    g. Harold Yang (Intel) - working on Intel ROS contributions. How Intel edge compute/AI can help ROS.
    h. Lewis (Weizhi Liu) - OpenVINO contributor, maintains OpenVINO toolkit for ROS. Object analytics, people analytics. Moveit2
    i. Zhen Ju “Jen” - software engineer, Huawei makes edge AI modules, ASIC for NN

Edge AI WG contributors and members

add yourself and choose your role as contributor or member (attends w/o commitment)

- Joe Speed **@joespeed** ADLINK Technology (WG leader, contributor)
- Steve Macenski **@smac** Samsung (contributor)
- Geoff Biggs **@gbiggs** (will assign) TierIV (contributor)
- Aaron Blasdel **@Aaron_Blasdel** Amazon (member)
- Harold Yang **@hyang5**, Lewis Liu **@LewisLiuPub** Intel (contributors)
- Katherine Scott **@Katherine_Scott** Open Robotics (contributor)
- Christoph Hellmann **@c_h_s** Fraunhofer IPA (contributor)
- Zhen Ju **@crystaldust** (contributor)
- Sumandeep Banerjee **@sumandeepb** Rapyuta Robotics, Tokyo (contributor)

- Bob A @boba (contributor or member?)
- Alex Tyshka @atyshka (member)
- Adrian Bedford @Adrian_Bedford (contributor or member?)
- Jeremy Adams @adamsj@objectcomputing.com github/discourse @adamsj-oci Object Computing (contributor or member?)

Edge AI WG Objectives

High level objects for your review & input from draft by Geoff, Steve, Aaron, Joe:

Objective: make Edge AI easier and ubiquitous in ROS 2

Overview: Edge AI, specifically ML has important applications in ROS 2. For example navigation, perception & picking, inspection, motion planning. We'll work to integrate and support technology for ML. We'll approach this via a process of assessing the current state and working to make it easier for ROS 2 users.

The WG seeks to do the following for ROS 2:

1. perform a survey of ML use in ROS and ROS 2
2. identify gaps and build a roadmap to close them
3. make it easier to use machine learning in ROS 2
4. enable machine learning on embedded processors found in typical mobile robots
5. enable HW acceleration of ML when present e.g. CPU SIMD, GPUs, FPGAs, VPU, NPU, TPU

Actions

1. perform a survey of ML use in ROS and ROS 2 today
2. Need systematic way to gather input - what tool to use? How to structure it? Github … markdown file
    a. Liyou Zhou: github handle: LiyouZhou
        i. 1) model zoo for AI applications - more generalized models, not HW specific
        ii. 2) AI inference engines Apache TVM abstracts HW architectures, deploy to different HW architectures
        iii. Quantized is important for edge
    b. Christoph Hellmann Santos: (github: ipa-cmh - discourse: c_h_s)
    c. Zhen Ju: @crystaldust for both github and discourse
    d. Decision: Github markdown file to start listing such things
    e. (add your github handles above if different than your discourse handles)
    f. ML frameworks? Tensor Flow, Tensor Lite, PyTorch, ? …

    g. Sumandeep - Inference onboard, offboard, onsite, on-cloud?
        i. What do you think?
        ii. Jeremy - (see note below re impedence mismatch between NN & accelerator, e.g. Jeremy having certain NN layers not supported by accelerators e.g. TensorRT)
        iii. Liyou - if we establish model zoo, look at TVM as execution framework, then we can make sure all models in the zoom can be executed / accelerated by TVM on the target HW accelerators ...
3. Each of us knows a part of the ML ecosystem. Contribute that knowledge to create the seed list.
4. After that share with larger community for input
5. How might we find every ROS package that uses ML? Do we need to or instead just focus on the most important ones?
    a. Nav2 in future? Most of Nav2 runs on CPU, has plans to add more abilities. Steve has the future plan.
        i. SLAM packages should use such features
    b. MoveIt2 smart grasp package is AI NN related work
    c. LGSVL simulator
    d. Lewis can find some more
    e. Sumandeep - toolchain, versioning of libraries breaks things. Is big easy of use issue.
    f. [jjadams] Does the edge accelerator work with the network? Newer networks may not be supported but how do you know? It would be nice to summarize supported networks per package..

Next Meeting date & time

1. Monthly or bi-weekly? Every **2 weeks to start …**
2. What day of week and time do you like?
    a. Maybe 1 or 2 or 3 hours later? 7 hours later? Alternate early & late times every 2 weeks.
3. One consistent time or tick-tock between early & late pacific time?

From the discourse thread ...

[**atyshka**] One suggestion for topic of discussion would be better support for GPU builds of software. I recall frustration when trying to use OpenCV compiled with CUDA for ROS, as many packages depend on it and they all would have to be build from source. Similarly, PCL offers limited CUDA functionality but it runs into the same source-compilation issues as OpenCV. If we want Edge AI to be a big part of ROS2, we need better support for GPU acceleration, and CUDA is where a large part of the community is at right now. GPU opens

up so many possibilities for ROS and it's currently underutilized. I might like to see alternate builds of binaries with CUDA support. I understand such an endorsement of proprietary software might be a controversial step, but it would be immensely useful.

[LewisLiuPub] Some robotics-capable CNN detection model. Preference for SSD mobilenet V1/2 or inception. This class of model is typical and will give a good basis for DL solutions that are generally capable of running on the edge on Nvidia hardware, but I'd love to know if I can run it on intel CPUs as well. These features have been supported by this project [16].

- It supports ROS and ROS 2 frameworks.
- It supports Intel OpenVINO and encapsules the utility of openCV.
- It supports Intel CPU / GPU / FPGA. (FPGA is not fully tested although).
- It support SSD mobilenet, YovoV2, master RCNN and other Intel trained models and public pre-trained modles.

I hope our work about ROS/ROS2 OpenVINO Toolkit projects can contribute a bit to this topic.

[sumandeepb Sumandeep Banerjee] great initiative. This in lines of something I have been ideating on in the past few months. I have also jotted down some ideas and approaches to attack this type of task. I am relatively a new comer to ROS but have extensive background in a variety of computer vision / machine learning / deep learning algorithm and product development including academic research, working on mine and other's startups. I definitely want to contribute and be part of this. I have worked on Intel Movidius, Jetson, Tegra GPU platforms, sensors such as RealSense, embedded boards such as RasberryPI, Tinkerboard etc. Portability and reusability of implementations and compute hardware / sensor support for widely varying hardware platforms has been a common challenge I have faced and continue to do so. Anything I can do to help the greater community and help speed up the overall pace of work in this domain is very much my intent. I am currently working on Robotics Perception Problems at Rapyuta Robotics, Tokyo. We are a cloud robotics platform based on ROS. Kindly have a look at my profile https://www.linkedin.com/in/sumandeep-banerjee-1436a17/

[**c_h_s Christoph**] We here at Fraunhofer IPA are also interested in the topic of how to deploy AI modules to robotics systems running ROS. Recently, we have started a research project that touches the area. I think this can be one major differentiator of ROS in the future compared to other industrial robot systems. We should certainly also focus on FPGA and neuromorphic chips as this seems to be an upcoming topic.

[**crystaldust**] Great proposal, is there some designs on architecture yet? If possible, I would like to make contributions, we have our smart device which accelerate the deep neural network, hope the arch will adapt multiple devices.

[**smac Steve Macenski**] Something I'm really interested by, and hopefully something Intel can help with, is an apples-to-apples comparison of a few classes of machine learning models with OpenVino-optimized model, raw CPU, and TensorRT-optimized models on a comparable Nvidia GPU (which, you know is a question of itself. But I have a Jetson Nano, a TX2, and I'm sure I could get access to a Xavier if we wanted to run the gambit).

In particular, starting from a common basis tensorflow model, instructions for how to run, instructions for how to optimize with openvino, instructions on how to optimize with TenorRT (which I can provide). This way you have a full document for "hey I trained this model, and now I want to deploy it on [all the major options]" and as the instructions we can use to reproducibly get test results for different classes of models.

Now what classes of models am I interested in?

- Some robotics-capable CNN detection model. Preference for SSD mobilenet V1/2 or inception. This class of model is typical and will give a good basis for DL solutions that are generally capable of running on the edge on Nvidia hardware, but I'd love to know if I can run int on intel CPUs as well.
- Some lower dimensional detection or classification model. Think: IMU, laser scan, or narrow/shallow CNN. This class of model is interesting to test for simpler problem sets, how much the optimizations really help, and again with an eye for "can I run this on my i5 CPU tractably" without needing a separate GPU unit.
- [other] I'm sure there's another good extreme to test here with other types of machine learning interests.

This is a really valuable investigation but perhaps not *directly* ROS related