

Overview

In this notebook I train a simple mnist model using depthwise convolutions, port it to nengo-dl, and port it to nengo-loihi.

This functionality is not supported on nengo right now so make sure to install the following branches first https://github.com/Michaeljurado24/nengo/tree/Convolution_groups

https://github.com/Michaeljurado24/nengo-loihi/tree/Convolution_groups

Import Everything

```
In [1]: import tensorflow as tf
import tensorflow.keras as keras
import numpy as np
import nengo
import nengo_dl
from nengo_dl.converter import ConvertConv
import nengo_loihi
```

```
2022-02-16 17:07:21.612345: W tensorflow/stream_executor/platform/default/dso_loader.cc:64] Could not load dynamic library 'libcudart.so.11.0'; dlerror: libcudart.so.11.0: cannot open shared object file: No such file or directory; LD_LIBRARY_PATH: /slurm/intel-archi/lib
2022-02-16 17:07:21.612370: I tensorflow/stream_executor/cuda/cudart_stub.cc:29] Ignore above cudart dlerror if you do not have a GPU set up on your machine.
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_dl/version.py:56: UserWarning: This version of NengoDL has not been tested with your Nengo version (3.2.1.dev0). The latest fully supported version is 3.1.0.
  warnings.warn(warnstr)
2022-02-16 17:07:23.506812: W tensorflow/stream_executor/platform/default/dso_loader.cc:64] Could not load dynamic library 'libcudart.so.11.0'; dlerror: libcudart.so.11.0: cannot open shared object file: No such file or directory; LD_LIBRARY_PATH: /slurm/intel-archi/lib
2022-02-16 17:07:24.602469: W tensorflow/stream_executor/platform/default/dso_loader.cc:64] Could not load dynamic library 'libcuda.so.1'; dlerror: libcuda.so.1: cannot open shared object file: No such file or directory; LD_LIBRARY_PATH: /slurm/intel-archi/lib
2022-02-16 17:07:24.602504: W tensorflow/stream_executor/cuda/cuda_driver.cc:269] failed call to cuInit: UNKNOWN ERROR (303)
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/version.py:20: UserWarning: This version of `nengo_loihi` has not been tested with your `nengo` version (3.2.1.dev0). The latest fully supported version is 3.2.0
  warnings.warn(
```

Set SNN parameters

```
In [2]: pres_steps = 200
pres_time = pres_steps * .001
n_test = 50
frequency_target = 250
synapse = .02
```

Train ANN

```
In [4]:
```

```

# Model / data parameters
num_classes = 10
input_shape = (28, 28, 1)

# the data, split between train and test sets
(x_train, y_train), (x_test, y_test) = keras.datasets.mnist.load_data()

# Scale images to the [0, 1] range
x_train = x_train.astype("float32") / 255
x_test = x_test.astype("float32") / 255
# Make sure images have shape (28, 28, 1)
x_train = np.expand_dims(x_train, -1)
x_test = np.expand_dims(x_test, -1)

# convert class vectors to binary class matrices
y_train = keras.utils.to_categorical(y_train, num_classes)
y_test = keras.utils.to_categorical(y_test, num_classes)

# Here we define a model with deptwise convolutions and regular convolutions
model = tf.keras.models.Sequential([
    tf.keras.layers.DepthwiseConv2D(3, input_shape=(28, 28, 1), activation = tf.nn.relu, use_bias = False),
    tf.keras.layers.Conv2D(3, kernel_size = (1, 1), activation = tf.nn.relu, use_bias = False),
    tf.keras.layers.Flatten(),
    tf.keras.layers.Dense(10)
])

batch_size = 128
epochs = 5

model.compile(loss=tf.keras.losses.CategoricalCrossentropy(from_logits = True), optimizer=keras.optimizers.Adam())
model.fit(x_train, y_train, batch_size=batch_size, epochs=epochs, validation_split=0.1)

```

```

Epoch 1/5
422/422 [=====] - 5s 12ms/step - loss: 0.7276 - accuracy: 0.8130
- val_loss: 0.2811 - val_accuracy: 0.9222
Epoch 2/5
422/422 [=====] - 5s 11ms/step - loss: 0.3139 - accuracy: 0.9119
- val_loss: 0.2515 - val_accuracy: 0.9297
Epoch 3/5
422/422 [=====] - 5s 11ms/step - loss: 0.2857 - accuracy: 0.9195
- val_loss: 0.2431 - val_accuracy: 0.9323
Epoch 4/5
422/422 [=====] - 5s 12ms/step - loss: 0.2749 - accuracy: 0.9220
- val_loss: 0.2369 - val_accuracy: 0.9332
Epoch 5/5
422/422 [=====] - 5s 11ms/step - loss: 0.2667 - accuracy: 0.9249
- val_loss: 0.2396 - val_accuracy: 0.9328
Out[4]: <keras.callbacks.History at 0x7f5ce9277610>

```

Prepare a version of the model that can accept spikes directly

```

In [5]: # Convert out model into the correct form
dummy_activation = lambda x: x
input = tf.keras.layers.Input((28, 28, 1))
to_spikes = tf.keras.layers.Activation(dummy_activation) # This will be replaced with new
spiking_input = to_spikes(input) # Will be replaced by SpikingRecifiedLinear in the dl c
x = spiking_input
for layer in model.layers:
    x = layer(x)
spiking_model = tf.keras.Model(inputs = input, outputs = x)

```

Convert model using nengo-dl

In [6]:

```
# Here we register a conversion pipeline for depthwise convolutions. This only works if t
@nengo_dl.Converter.register(tf.keras.layers.DepthwiseConv2D)
class ConvertDepthWise2D(ConvertConv):
    """Base class for converting convolutional layers to Nengo objects."""

    has_weights = True

    def convert(self, node_id):
        kernel = np.moveaxis(self.layer.get_weights()[0], -1, -2) # Convert weights of d
        if self.layer.use_bias:
            raise(Exception("Biases not currently supported"))

        # create nengo object to implement activation function
        output = self.add_nengo_obj(node_id, activation=self.layer.activation)

        # set up a convolutional transform that matches the layer parameters
        if self.layer.data_format == "channels_last":
            n_filters = self.input_shape(node_id)[-1]
        else:
            n_filters = self.input_shape(node_id)[0]

        transform = nengo.Convolution(
            n_filters=n_filters,
            input_shape=self.input_shape(node_id),
            kernel_size=self.layer.kernel_size,
            strides=self.layer.strides,
            padding=self.layer.padding,
            channels_last=self.layer.data_format == "channels_last",
            init=kernel,
            groups = n_filters
        )

        self.add_connection(node_id, output, transform=transform, trainable=False)
        return output

# Convert ANN to SNN
nengo_converter = nengo_dl.Converter(
    spiking_model,
    swap_activations={tf.nn.relu: nengo_loihi.neurons.LoihiSpikingRectifiedLinear() ,
                      dummy_activation: nengo.SpikingRectifiedLinear()}, # off the chip

    synapse=synapse,
    allow_fallback = False,
    inference_only = True
)

# Scale input hz to match a frequency target
x_test = x_test.reshape(x_test.shape[0], -1) * frequency_target
test_images = np.tile(x_test[:, None, :], (1, pres_steps, 1))
test_labels = np.tile(y_test[:, None, None], (1, pres_steps, 1))

# get input/output objects
nengo_input = nengo_converter.inputs[spiking_model.input]
nengo_output = nengo_converter.outputs[spiking_model.output]

# with nengo_dl.Simulator(
#     nengo_converter.net, minibatch_size=5, progress_bar=False
# ) as nengo_sim:
#     data = nengo_sim.predict({nengo_input: test_images})
```

```

# predictions = np.argmax(data[nengo_output][:, -1], axis=-1)
# accuracy = (predictions == np.argmax(test_labels[:n_test, 0, 0], axis = -1)).mean()
# print(f"Nengo-DL Test accuracy: {100 * accuracy:.2f}%")

```

Nengo-loihi Inference

In [7]:

```

with nengo_converter.net as net:
    nengo_input.output = nengo.processes.PresentInput(
        x_test, presentation_time=pres_time
    )

with nengo_converter.net as net:
    nengo_loihi.add_params(net)
    net.config[nengo_converter.layers[to_spikes].ensemble].on_chip = False

with nengo_loihi.Simulator(net, target = "loihi", precompute = False) as loihi_sim:
    loihi_sim.run(n_test * pres_time)

# get output (last timestep of each presentation period)
pres_steps = int(round(pres_time / loihi_sim.dt))
output = loihi_sim.data[nengo_output][pres_steps - 1 :: pres_steps]

# compute the Loihi accuracy
loihi_predictions = np.argmax(output, axis=-1)

correct = 100 * np.mean(loihi_predictions == np.argmax(test_labels[:n_test, 0, 0], ax
print("Nengo Loihi Test accuracy: %.2f%%" % correct)

```

3

INFO:DRV: SLURM is being run in background

/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/simulator.py:159: UserWarning: Model is precomputable. Setting precompute=False may slow execution.

warnings.warn(

INFO:DRV: Connecting to 10.212.98.104:39995

INFO:DRV: Host server up.....Done 0.17s

INFO:DRV: Encoding axons/synapses.....Done 0.07s

INFO:DRV: Compiling Embedded snips....Done 0.36s

INFO:DRV: Compiling Host snips.....Done 0.58s

INFO:DRV: Compiling MPDS Registers....Done 0.50ms

INFO:HST: Args chip=0 cpu=0 /homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nxsdk/driver/compiler/../../../../temp/1645060092.4716344/launcher_chip0_lmt0.bin --chips=1 --remote-relay=0

INFO:HST: Nx...

INFO:DRV: Booting up.....Done 0.77s

INFO:DRV: Encoding probes.....Done 1.23ms

INFO:DRV: Transferring probes.....Done 3.56ms

INFO:DRV: Configuring registers.....Done 0.03s

INFO:DRV: Transferring spikes.....Done 1.24ms

INFO:HST: [Host] Listening for client

INFO:HST: [Host] Connected to client

INFO:HST: chip=0 cpu=0 time 100

INFO:HST: chip=0 cpu=0 time 200

INFO:HST: chip=0 cpu=0 time 300

INFO:HST: chip=0 cpu=0 time 400

INFO:HST: chip=0 cpu=0 time 500

INFO:HST: chip=0 cpu=0 time 600

INFO:HST: chip=0 cpu=0 time 700

INFO:HST: chip=0 cpu=0 time 800

INFO:HST: chip=0 cpu=0 time 900

INFO:HST: chip=0 cpu=0 time 1000

INFO:HST: chip=0 cpu=0 time 1100

```
INFO:HST: chip=0 cpu=0 time 1200
INFO:HST: chip=0 cpu=0 time 1300
INFO:HST: chip=0 cpu=0 time 1400
INFO:HST: chip=0 cpu=0 time 1500
INFO:HST: chip=0 cpu=0 time 1600
INFO:HST: chip=0 cpu=0 time 1700
INFO:HST: chip=0 cpu=0 time 1800
INFO:HST: chip=0 cpu=0 time 1900
INFO:HST: chip=0 cpu=0 time 2000
INFO:HST: chip=0 cpu=0 time 2100
INFO:HST: chip=0 cpu=0 time 2200
INFO:HST: chip=0 cpu=0 time 2300
INFO:HST: chip=0 cpu=0 time 2400
INFO:HST: chip=0 cpu=0 time 2500
INFO:HST: chip=0 cpu=0 time 2600
INFO:HST: chip=0 cpu=0 time 2700
INFO:HST: chip=0 cpu=0 time 2800
INFO:HST: chip=0 cpu=0 time 2900
INFO:HST: chip=0 cpu=0 time 3000
INFO:HST: chip=0 cpu=0 time 3100
INFO:HST: chip=0 cpu=0 time 3200
INFO:HST: chip=0 cpu=0 time 3300
INFO:HST: chip=0 cpu=0 time 3400
INFO:HST: chip=0 cpu=0 time 3500
INFO:HST: chip=0 cpu=0 time 3600
INFO:HST: chip=0 cpu=0 time 3700
INFO:HST: chip=0 cpu=0 time 3800
INFO:HST: chip=0 cpu=0 time 3900
INFO:HST: chip=0 cpu=0 time 4000
INFO:HST: chip=0 cpu=0 time 4100
INFO:HST: chip=0 cpu=0 time 4200
INFO:HST: chip=0 cpu=0 time 4300
INFO:HST: chip=0 cpu=0 time 4400
INFO:HST: chip=0 cpu=0 time 4500
INFO:HST: chip=0 cpu=0 time 4600
INFO:HST: chip=0 cpu=0 time 4700
INFO:HST: chip=0 cpu=0 time 4800
INFO:HST: chip=0 cpu=0 time 4900
INFO:HST: chip=0 cpu=0 time 5000
INFO:HST: chip=0 cpu=0 time 5100
```

```
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardw
are/interface.py:446: UserWarning: Too many spikes (59) sent in one timestep. Increase the
value of `snip_max_spikes_per_step` (currently set to 50). See
```

```
https://www.nengo.ai/nengo-loihi/configuration.html
for details.
```

```
warnings.warn(
```

```
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardw
are/interface.py:446: UserWarning: Too many spikes (61) sent in one timestep. Increase the
value of `snip_max_spikes_per_step` (currently set to 50). See
```

```
https://www.nengo.ai/nengo-loihi/configuration.html
for details.
```

```
warnings.warn(
```

```
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardw
are/interface.py:446: UserWarning: Too many spikes (63) sent in one timestep. Increase the
value of `snip_max_spikes_per_step` (currently set to 50). See
```

```
https://www.nengo.ai/nengo-loihi/configuration.html
for details.
```

```
warnings.warn(
```

```
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardw
are/interface.py:446: UserWarning: Too many spikes (55) sent in one timestep. Increase the
value of `snip_max_spikes_per_step` (currently set to 50). See
```

```
https://www.nengo.ai/nengo-loihi/configuration.html
for details.
```

```
warnings.warn(
```

```
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardw
```


are/interface.py:446: UserWarning: Too many spikes (69) sent in one timestep. Increase the value of `snip_max_spikes_per_step` (currently set to 50). See

<https://www.nengo.ai/nengo-loihi/configuration.html> for details.

```
warnings.warn(
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardware/interface.py:446: UserWarning: Too many spikes (70) sent in one timestep. Increase the value of `snip_max_spikes_per_step` (currently set to 50). See
```

<https://www.nengo.ai/nengo-loihi/configuration.html> for details.

```
warnings.warn(
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardware/interface.py:446: UserWarning: Too many spikes (67) sent in one timestep. Increase the value of `snip_max_spikes_per_step` (currently set to 50). See
```

<https://www.nengo.ai/nengo-loihi/configuration.html> for details.

```
warnings.warn(
/homes/mjurado3/miniconda3/envs/loihi_vishal/lib/python3.9/site-packages/nengo_loihi/hardware/interface.py:446: UserWarning: Too many spikes (62) sent in one timestep. Increase the value of `snip_max_spikes_per_step` (currently set to 50). See
```

<https://www.nengo.ai/nengo-loihi/configuration.html> for details.

```
warnings.warn(
```

```
INFO:HST: chip=0 cpu=0 time 5200
```

```
INFO:HST: chip=0 cpu=0 time 5300
```

```
INFO:HST: chip=0 cpu=0 time 5400
```

```
INFO:HST: chip=0 cpu=0 time 5500
```

```
INFO:HST: chip=0 cpu=0 time 5600
```

```
INFO:HST: chip=0 cpu=0 time 5700
```

```
INFO:HST: chip=0 cpu=0 time 5800
```

```
INFO:HST: chip=0 cpu=0 time 5900
```

```
INFO:HST: chip=0 cpu=0 time 6000
```

```
INFO:HST: chip=0 cpu=0 time 6100
```

```
INFO:HST: chip=0 cpu=0 time 6200
```

```
INFO:HST: chip=0 cpu=0 time 6300
```

```
INFO:HST: chip=0 cpu=0 time 6400
```

```
INFO:HST: chip=0 cpu=0 time 6500
```

```
INFO:HST: chip=0 cpu=0 time 6600
```

```
INFO:HST: chip=0 cpu=0 time 6700
```

```
INFO:HST: chip=0 cpu=0 time 6800
```

```
INFO:HST: chip=0 cpu=0 time 6900
```

```
INFO:HST: chip=0 cpu=0 time 7000
```

```
INFO:HST: chip=0 cpu=0 time 7100
```

```
INFO:HST: chip=0 cpu=0 time 7200
```

```
INFO:HST: chip=0 cpu=0 time 7300
```

```
INFO:HST: chip=0 cpu=0 time 7400
```

```
INFO:HST: chip=0 cpu=0 time 7500
```

```
INFO:HST: chip=0 cpu=0 time 7600
```

```
INFO:HST: chip=0 cpu=0 time 7700
```

```
INFO:HST: chip=0 cpu=0 time 7800
```

```
INFO:HST: chip=0 cpu=0 time 7900
```

```
INFO:HST: chip=0 cpu=0 time 8000
```

```
INFO:HST: chip=0 cpu=0 time 8100
```

```
INFO:HST: chip=0 cpu=0 time 8200
```

```
INFO:HST: chip=0 cpu=0 time 8300
```

```
INFO:HST: chip=0 cpu=0 time 8400
```

```
INFO:HST: chip=0 cpu=0 time 8500
```

```
INFO:HST: chip=0 cpu=0 time 8600
```

```
INFO:HST: chip=0 cpu=0 time 8700
```

```
INFO:HST: chip=0 cpu=0 time 8800
```

```
INFO:HST: chip=0 cpu=0 time 8900
```

```
INFO:HST: chip=0 cpu=0 time 9000
```

```
INFO:HST: chip=0 cpu=0 time 9100
```

```
INFO:HST: chip=0 cpu=0 time 9200
```

```
INFO:HST: chip=0 cpu=0 time 9300
```

```
INFO:HST: chip=0 cpu=0 time 9400
```

```
INFO:HST: chip=0 cpu=0 time 9500
INFO:HST: chip=0 cpu=0 time 9600
INFO:HST: chip=0 cpu=0 time 9700
INFO:HST: chip=0 cpu=0 time 9800
INFO:HST: chip=0 cpu=0 time 9900
INFO:HST: chip=0 cpu=0 time 10000
INFO:HST: chip=0 cpu=0 Waited to exit (nonsense sum -13580)
INFO:DRV:      Executing.....Done 0.45s
INFO:DRV:      Processing timeseries.....Done 1.31ms
Nengo Loihi Test accuracy: 96.00%
INFO:HST: [Host] Received shutdown signal: -1
INFO:HST: [Host] Wrote superhost shutdown signal: 8192 bytes
INFO:HST: [Host] Closing server socket
INFO:HST: chip=0 cpu=0 halted, status=0x0
```

In []: