Running AI workloads on IBM Power Systems

Maxime Deloche

Deep Learning Engineer

Jean-Armand Broyelle

Cognitive Systems Lab Technical Leader

Cognitive Systems Lab IBM Garage Montpellier, France



Plan

- Open Cognitive Environment (OpenCE)

- Overview & requirements
- Release cycle

- How to use OpenCE

- How to build
- Conda channels
- Additional repositories

- Containerization

- Advantages
- Build custom development environments

– Useful software

• Horovod, Dask, scikit-optimize...

Open Cognitive Environment (OpenCE) overview



Open Cognitive Environment (OpenCE)

- Software distribution for AI & Deep Learning applications
- Helps building a complete environment for AI development on Power
- Free and based on open-source SW + optimizations for IBM Power Systems

This toolkit includes...

- Deep learning frameworks
- NVIDIA software: CUDA toolkit, CuDNN, NCCL, DALI, TensorRT...
- Data science libraries: numpy, pandas, scikitlearn, Dask, Horovod, ...



OpenCE overview

-<u>https://github.com/open-ce</u>

- Source-to-image project to provide preintegrated recipes and build scripts
- Packages built to run in a conda environment
- Previously "IBM Watson Machine Learning Community Edition" (WMLCE) and before "IBM PowerAI"

Differences with previous "WMLCE":

- More flexible "source-to-image" workflow
- Shift from IBM support to community support
 - Large clusters with active contributors: Oak Ridge National Laboratory "Summit", MIT "Satori", Oregon State University...
- Drop support of a few packages and features: Large Model Support (LMS), RAPIDS, SnapML

Release cycle

- Current version is 1.1.3
- Uncoordinated upstream releases
- Rule of thumb: a new release should include a new release of both Tensorflow and Pytorch
- A release is a set of upstream versions that are guaranteed to work together
- This doesn't mean that a newer version of an upstream won't work!





Software requirements

- Operating System
 - RHEL 7.6 or higher
 - Ubuntu 18.04 or higher
- Python >= 3.6
- NVIDIA GPU Drivers v. 440 (when using GPUs)
- Anaconda installer (https://repo.anaconda.com/archive/Anaconda3-2020.11-Linux-ppc64le.sh)
 - conda >= 3.8.3 and conda-build >= 3.20.5

How to use OpenCE

124 } 125 h3{ font-size: 22gu; color: skelele; font-family: "antservelregiler"; s:...sill bockgeround: wrl(.../imp/million.pmp) m-rupent animer wideb: infine-block; wideb: infine: height: infine; floct: left; mergin: 2px Ppx 0 0; 130 } 131 em.mail{ 136 (20) 133 134 135 136 137 138 139 **)** 139) ~ 140 em.phone[141 bockground: url(../kms/phonet.cs.pmg) morequest casher; 142 display: inline-block; 143 height: 18px; 144 floct: left; 144 floct: left; 145 morenin: Save Rave & A: 146 morenin: Save Rave & A:

OpenCE packaging

Conda is an open-source package manager for multiple languages, not only Python

- Dependencies automatically resolved
- Delivery of packages is continuous
- Conda environments isolate software stacks



Documentation: <u>https://docs.conda.io/projects/conda/en/latest/user-guide/tasks/index.html</u> Anaconda installer: <u>https://repo.anaconda.com/archive/Anaconda3-2020.11-Linux-ppc64le.sh</u>

You need a conda channel. 2 options:

Easy way

- Use an already built, publicly accessible Conda channel of OpenCE: <u>https://opence.mit.edu/</u>
- Built for the IBM Power9 "Satori" cluster at MIT
- Ready to use as conda channel (see Conda section below)
- Useful tips and guides in their documentation: <u>https://mit-satori.github.io/</u>
- Alternative (Oregon State University channel): <u>https://ftp.osuosl.org/pub/open-ce/current/</u>

Alternative

- Build a conda channel of OpenCE
- Cons: you must handle the build and the access to the channel
- Pros: you are autonomous on what packages to build and how often it is updated

Open Cognitive Environment (OpenCE)

Packages	Documentation	
₩ Filters	: Q pytorch	

	Name	Version	Platform	Summary
~	_pytorch_select	1.0	linux-ppc64le	Package used to select the specific PyTorch build variant
~	_pytorch_select	2.0	linux-ppc64le	Package used to select the specific PyTorch build variant
~	pytorch	1.6.0	linux-ppc64le	Meta-package to install GPU-enabled PyTorch variant
~	pytorch-base	1.6.0	linux-ppc64le	PyTorch is an optimized tensor library for deep learning using GPUs and CPUs.
~	pytorch-cpu	1.6.0	linux-ppc64le	Meta-package to install CPU-only PyTorch variant
~	pytorch	1.7.1	linux-ppc64le	Meta-package to install GPU-enabled PyTorch variant
~	pytorch-base	1.7.1	linux-ppc64le	PyTorch is an optimized tensor library for deep learning using GPUs and CPUs.
~	pytorch-cpu	1.7.1	linux-ppc64le	Meta-package to install CPU-only PyTorch variant
~	pytorch-lightning	0.9.0	noarch	PyTorch Lightning is the lightweight PyTorch wrapper for ML researchers. Scale your models. Write less boilerplate.
~	pytorch-lightning	1.1.0	noarch	PyTorch Lightning is the lightweight PyTorch wrapper for ML researchers. Scale your models. Write less boilerplate.
~	pytorch-lightning-bolts	0.2.5	noarch	Pretrained SOTA Deep Learning models, callbacks and more for research and production with PyTorch Lightning and PyTorch.

Open-CE

Open Cognitive







How to build OpenCE (1/2)

- Install requirements (including CUDA if GPUs)
- Git clone the "open-ce" repo
- Pick the environment you want to build from the "open-ce-environments" repo
- Run the "open-ce build env..." command
 - \$ git clone https://github.com/open-ce/open-ce.git
 - \$./open-ce/open_ce/open-ce build env \

https://raw.githubusercontent.com/open-ce/\

open-ce-environments/main/envs/opence-env.yaml

– Full "OpenCE" build requires 10 to 15 hours ; and a lot less for smaller YAML envs.

Ľ	bazel-env.yaml
Ľ	conda_build_config.yaml
ß	dali-env.yaml
ß	horovod-env.yaml
Ľ	lightgbm-env.yaml
Ľ	opence-env.yaml
Ľ	pytorch-env.yaml
ß	pytorch-extras-env.yaml
Ľ	spacy-env.yaml
ß	tensorboard-env.yaml
Ľ	tensorflow-env.yaml
ß	tensorrt-env.yaml
ß	transformers-env.yaml
ß	xgboost-env.yaml

How to build OpenCE (2/2)

- Previous build produces a "./condabuild" dir
- \$ conda install -c ./condabuild \
 numpy tensorflow ...
- You can possibly move that directory to a web server and make it available to your users (the same way the MIT OpenCE conda channel works)

- Possibility to build single feedstocks
 - \$ open-ce build feedstock ...
- Possibility to run tests after a build (defined for each package in their feedstock repo)
 - \$ open-ce test env ...
- Possibility to build inside a container (clean environment for the build and makes builds more "system independent")
 - \$ open-ce build env \
 --container_build ...

Configuration and use

or replace with your own Conda channel

- Add OpenCE to the conda channels (or modify directly "~/.condarc" file):
 - \$ conda config --prepend channels http://opence.mit.edu/
- List installed packages
 - \$ conda list
- Search a package
 - \$ conda search numpy
- Install a package (optionally specifying version)
 - \$ conda install numpy=1.17.4
- Uninstall a package
 - \$ conda remove numpy

GPUs and CUDA

- GPU drivers (v. 440 or higher) need to be installed (<u>https://www.nvidia.com/Download/index.aspx</u>)
- CUDA runtime package is available in the Anaconda default channel: "cudatoolkit" (version 10.2 or 11.0)
 - Package link: https://anaconda.org/anaconda/cudatoolkit
- It includes everything needed by GPU-accelerated apps:
 - GPU-accelerated libraries, CUDA runtime, compiler, profiling tools...

Manage Conda environments

- Easy to create, install and switch between software stacks!
- List environments
 - \$ conda env list
- Create an environment
 - \$ conda activate my_env
- Switch between environments
 - \$ conda create -n my_env python=3.7

- Conda environments allow to have multiple software stacks:
 - Multiple Python versions
 - Different frameworks or versions

(my_env) mdeloche-14-8sq4l @pwrai ~/documents >

Active environment is displayed at the beginning of your prompt.

(opence) mdeloche-14-8sq4l @pwrai ~/documents > conda info

active environment	: opence
active env location	: /opt/anaconda/envs/opence
shell level	
user config file	: /home/pwrai/.condarc
populated config files	: /home/pwrai/.condarc
conda version	: 4.7.12
conda-build version	: not installed
python version	: 3.7.7.final.0
virtual packages	:cuda=10.2
bas <u>e environment</u>	:/opt/anaconda (writable)
channel URLs	: https://opence.mit.edu/linux-ppc64le
	https://opence.mit.edu/noarch
	https://repo.anaconda.com/pkgs/main/linux-ppc64le
	https://repo.anaconda.com/pkgs/main/noarch
	https://repo.anaconda.com/pkgs/r/linux-ppc64le
	https://repo.anaconda.com/pkgs/r/noarch
package cache	: /opt/anaconda/pkgs
	/home/pwrai/.conda/pkgs
envs directories	: /opt/anaconda/envs
	/home/pwrai/.conda/envs
platform	: linux-ppc64le
user-agent	: conda/4.7.12 requests/2.23.0 CPython/3.7.7 Linux/4.14.0–115.18.1.el7a.ppc64le ubuntu/18.04.4 glibc/2.27
UID:GID	: 2051:2051
netrc file	: None
offline mode	: False

Additional repositories

- Supplementary channel (useful packages that are not part of OpenCE):
 - <u>https://anaconda.org/powerai</u>
- "old" WMLCE main channel:
 - <u>https://public.dhe.ibm.com/ibmdl/export/pub/software/server/ibm-ai/conda/</u>
- "old" WMLCE early-access channel:
 - <u>https://public.dhe.ibm.com/ibmdl/export/pub/software/server/ibm-ai/conda-early-access/</u>

If your package is still not available?

- Try pip:

- Install pip with conda, and search for your package
- If pip finds the package:
 - If there is a binary built for ppc64le, it will install it and dependencies
 - If not, it will try to install it using its "setup.py" file (see setuptools)

Warning	This conflicts if the same package is installed in Conda: only use it when you can't find it in
	the various Conda repositories!

- If pip compilation fails, try to recompile the package manually

Hint Build a wheel file (self-contained pip package) that you can then save and re-use on any ppc64le system (and possibly add it to your built Conda channel).

Air-gapped systems

- Local channel needed, not only with packages from OpenCE but also standard packages from main repos
- Solution: think ahead about your stack and create a "custom" conda channel
- 1. On an internet-facing system, clean the cache
- \$ conda clean --all
- 2. Build an environment will all packages you need ; and locate the cache
- \$ conda create -n my_env PACKAGES
- \$ conda info | grep "package cache"

- Copy all ".tar.bz2" and ".conda" files from the cache to "MY_CHANNEL/noarch/" on the air-gapped system
- 4. Install Anaconda or your air-gapped system and build a local conda channel from it (this creates metadata files needed by conda)
 - \$ conda index MY_CHANNEL/
- 5. Install using that directory, or add it to your conda configured channels

\$ conda config --prepend channels \
 file://path/to/MY_CHANNEL/

Containerization



Container images

- Advantages of working in containers:
 - Isolation of users on your system
 - Resources limits (GPUs in particular)
 - Custom software stack for each container
- Can run on a container orchestrator (Kubernetes/Red Hat Openshift) or standalone using Docker commands

- No (known) OpenCE available images yet
- Previous "WMLCE" IBM images: <u>https://hub.docker.com/r/ibmcom/powerai</u>

 The image we built for our customers (that you can freely modify and reuse): <u>https://gitlab.com/PSLC/wmlce-mop</u>

Your own minimal image

FROM ubuntu:20.04

ENV PATH /opt/conda/bin:\$PATH

RUN apt-get update -y && apt-get install -y wget

RUN wget --quiet \
 https://repo.anaconda.com/archive/Anaconda3-2020.11-Linux-ppc64le.sh && \
 sh ./Anaconda3-2020.11-Linux-ppc64le.sh -b -p /opt/conda && \
 conda config --prepend https://opence.mit.edu/ && \
 conda create -n opence python=3.8 tensorflow pytorch

Add additional libs, conda environments, GUI, SSH service...

JupyterLab (1/2)

💭 File Edit View Run Kernel Tabs Settings Help

	+ 🗈	<u>*</u>	G		
_	🖿 / 🚥 / scripts / tf_cnn_benchr	marks /			
۱.	Name	•	Last Modified		
	models		4 minutes ago	^	
_	platforms		4 minutes ago		
-	🖿 test_data		4 minutes ago		
_	all_reduce_benchmark_test.py		4 minutes ago		
ř	🍦 all_reduce_benchmark.py		4 minutes ago		
	Nallreduce_test.py		4 minutes ago		
2	neduce.py		4 minutes ago		
	ᅌ batch_allreduce.py		4 minutes ago		
ן	🝦 benchmark_cnn_distributed_te	st	4 minutes ago		
	👌 benchmark_cnn_distributed_te	st	4 minutes ago		
ŀ	🍦 benchmark_cnn_test.py		4 minutes ago		
	🍦 benchmark_cnn.py		4 minutes ago		
	🝦 cnn_util_test.py		4 minutes ago		
	👌 cnn_util.py		4 minutes ago		
	🔁 coco_metric.py		4 minutes ago		
	net constants.py		4 minutes ago		
	🕏 convnet_builder.py		4 minutes ago		
	🔁 datasets.py		4 minutes ago		
	🔁 flags.py		4 minutes ago		
	leading_indicators_test.py		4 minutes ago		
	🕏 mlperf_test.py		4 minutes ago		
	👌 mlperf.py		4 minutes ago		
	🔁 preprocessing.py		4 minutes ago		
	README.md		4 minutes ago		
	🍦 run_tests.py		4 minutes ago		
	ssd_constants.py		4 minutes ago		

1 s_ 1 🥶

NVIDIA-SMI 440.	64.00 Driver	Version: 440.64.00	CUDA Versio	n: 10.2		
 GPU Name Fan Temp Perf	Persistence-M Pwr:Usage/Cap	+ Bus-Id Disp.A Memory-Usage	+ Volatile GPU-Util	Uncorr. ECC Compute M.		
 0 Tesla V100 N/A 26C P0	-SXM2 Off 36W / 103W	00000004:04:00.0 Off 11MiB / 16160MiB	+====== 0%	0 Default		
Processes: GPU PID No running pro (opence) pwrai@md	Type Proces cesses found eloche-opence-1	s name -z2d5f:/wmlce/data\$ con	da env list	GPU Memory Usage 		
# conda environme # base opence theia	opt/an/ /opt/an/ /opt/an/	aconda aconda/envs/opence aconda/envs/theia				
(opence) pwrai@md _pytorch_select oytorch bytorch-base torchtext torchvision-base (opence) pwrai@md	leloche-opence-1 2.0 1.7.1 1.7.1 0.8.1 0.8.2 leloche-opence-1	-z2d5f:/wmlce/data\$ con cuda10.2 hca541ab cuda10.2_py38 py38 cuda10.2_py38 cuda10.2_py38 -z2d5f:/wmlce/data\$ []	da list g _1 https _1 https _8 https _4 https _5 https	rep torch ://opence.mit ://opence.mit ://opence.mit ://opence.mit	.edu .edu .edu .edu .edu	
torchtext torchvision-base (opence) pwrai@md	0.8.1 0.8.2 eloche-opence-1	 cuda10.2_py38 -z2d5f:/wmlce/data\$ []	_4 https _5 https	://opence.mit ://opence.mit	.edu .edu	

Untitled.ipynb Х + % 🗇 🖞 Python 3 O ■ C → Code \sim • [1]: import torch [5]: if torch.cuda.is_available(): for i in range(torch.cuda.device_count()): print(torch.cuda.get_device_name(i)) print(torch.cuda.get_device_properties(i)) Tesla V100-SXM2-16GB _CudaDeviceProperties(name='Tesla V100-SXM2-16GB', major=7, minor=0, tota 1_memory=16128MB, multi_processor_count=80)

JupyterLab (2/2)

Source: https://github.com/jupyterlab/jupyterlab

Doc: https://jupyterlab.readthedocs.io/en/stable/getting_started/installation.html

RUN conda run -n opence conda install -y nodejs jupyterlab

Theia (1/2)

File Edit Selection View Go Run Terminal Help

L.J.	EXPLORER: DATA Ö I	ð ••	tf_cnn_benchmarks.p		⊳	e allreduce.p			
	 ipynb_checkpoints 	unner	36 flags.def 37 for name 38 abs1_fl 39 abs1_flag 40 abs1_flag 41 'mrin 42 'Prin 43 'MLPe 44 'http: 45 'PYTH 46 'PYTH 47 def main(49 # Comma 50 # '-di prevent 51 51 # this positio 52 53 assert 54 if len(55 raise 56 57	<pre>ine_flags() in flags.param_specs.keys(): ags.declare_key_flag(name) s.DEFINE_boolean(erf_compliance_logging', False, t logs required to be compliant with MLPerf. If set, must cl of training repo <u>https://github.com/mlperf/training</u> and add s://github.com/mlperf/training/tree/master/compliance to the DNPATH') positional_arguments): nd-line arguments like 'distortions False' are equivalent stortions=True False', where False is a positional argument. From silently running with distortions, we do not allow nal ents. Ren(positional_arguments) >= 1 positional_arguments) > 1: ValueError('Received unknown positional arguments: %s'</pre>	Part of An Angel and Angel		<pre>rate valueError('all_reduce_spec (%s) contains non-integer range %s' %</pre>		
	constants.py		59 with ml	<pre>berf.mlperf_logger(absl_flags.FLAGS.ml_perf_compliance_loggi</pre>	ıg,		ef build_all_reduce_device_prefixes(job_name, num_tasks): """Build list of device prefix names for all_reduce.	Marcaran and a state of the second se	
က်နှ	 datasets py flags.py leading_indicators_test.py mlperf_test.py mlperf py preprocessing.py README.md run_tests.py ssd_constants.py ssd_dataloader.py test_util.py 		● Problems >_pw (opence) pwrai@mde Cloning into 'bencl remote: Enumeratin, remote: Counting of remote: Compressing remote: Total 4944 Receiving objects: Resolving deltas: (opence) pwrai@mde ICENSE README.md (opence) pwrai@mde	<pre>ai@mdeloche-opence-1-z2d5f:/wmlce/data × oche-opence-1-z2d5f:/wmlce/data\$ git clone https://github marks' ; objects: 1007, done. ; objects: 100% (100/100), done. ; oche-opence-1-z2d5f:/wmlce/data\$ ls benchmarks/ perface0 scripts</pre>	.com/tensorflow	/benchmarks		•	

Theia (2/2)

Doc: https://theia-ide.org/docs/

```
COPY theia-package.json /tmp/theia-install/package.json
RUN cd /tmp/theia-install && \
        conda create -y -n theia python=2 nodejs && \
        conda run -n theia npm install -g yarn && \
        conda run -n theia yarn && \
        conda run -n theia yarn && \
        conda run -n theia yarn theia build
CMD ["conda", "run", "-n", "theia", \
        _____"yarn", "theia", "start", "--hostname", "127.0.0.1"]
```

Containerization - Final thoughts (1/2)

- Those are 2 examples among many web-based IDEs. You can take advantage of them even when not running in containers.
- Containers allow to keep the host environment clean and stable, and still let the cluster users customize their work environment and tailor them to their development needs.
- When building containers, you can use CI/CD pipelines (Github Actions, Gitlab CI/CD, ...) to simplify the development workflow: by automating the build, test and push to a registry steps.

Containerization - Final thoughts (2/2)

- JupyterHub (<u>https://github.com/jupyterhub/jupyterhub</u>) is multi-user hub of Jupyter Notebooks
 - Spawns and manages multiple instances of single-user environments
 - You can bring in your own Docker image! (see <u>https://jupyterhub-dockerspawner.readthedocs.io/en/latest/</u>)



Useful software

ion a return a className getAttribute "className" appendChild(a).id=u, in.getElementsByName[[in.getElementsByName(te("id")===b}}):(delete d.find.ID,d.filter.ID=function(a){var b=a.re undefined"!=typeof b.getElementsByTagName?b.getElementsByTagName(a): ..getElementsByClassName&&function(a,b){return"undefined"!=typeof b.g allowcapture=''><option selected=''></option></select>", a. querySelected=''></option></select>", a. querySelected=''></option></selected=''></option></selected='' h||q.push("~="),a.querySelectorAll(":checked").length||q.push(":checked ame=d]").length&&q.push("name"+L+"*[*^\$|!~]?="),a.querySelectorAll(":en wia(function(a){c.disconnectedMatch=s.call(a,"div"),s.call(a,"[s!= ntElement:a.d=b&&b.parentNode;return a===d||!(!d||1!==d.nodeType||!(c.cc ompareDocumentPosition; return d?d:(d=(a.ownerDocument[]a)===(b 0 4 function(a b) if(a===b)return l= 0 0 var c ?ka -v71-0 toLowerCase attribute stchild 计输出法 机动物 医马勒氏 医白细胞 C. March SIGN

and and a sector of the sector

HOROVOD

- Distributed deep learning training framework
- Interfaces with Tensorflow, Keras, Pytorch and MXNet
- Based for underlying communications on:
 - Use MPI (Message Passing Interface) to communicate on the control plane
 - Use Nvidia NCCL to communicate on the data plane (MPI-compatible library with operations optimized for GPUs, especially "allreduce")
- Latest version 0.21.0 included in OpenCE, as well as NCCL and OpenMPI

Horovod (2/3)

```
# initialize Horovod
horovod.init()
```

```
# pin GPU to be used to process local rank (one GPU per process)
config.gpu_options.visible_device_list = str(horovod.local_rank())
```

```
# define dataset and model
train dataset = ...
```

```
model = Sequential().add(...)
```

```
# create a distributed optimizer
opt = horovod.DistributedOptimizer(keras.optimizers.Adam())
```

train the model # Callback broadcasts initial variable states from rank 0 to all other processes. # This is necessary to ensure consistent initialization of all workers. model.fit(train_dataset, callbacks=[horovod.callbacks.BroadcastGlobalVariablesCallback(0)])

Horovod (3/3)

Horovod defines a "horovodrun" command that is a wrapper to "mpirun" calls.

- Run on single host with 4 GPUs:
- \$ horovodrun -np 4 -H localhost:4 python train.py
- Run on 3 hosts with 4 GPUs each:
 - \$ horovodrun -np 12 -H server1:4,server2:4,server3:4 python train.py

 Benchmarks shows a scaling of ~90% on 128 4-GPU servers

Dask (1/2)

- Parallel computing library for Python scale from a laptop to a cluster with the same familiar API
- 2 components:
 - Distributed task scheduler -
 - Dask Collections optimized for distributed environments and larger-than-memory data
- Great doc: <u>https://docs.dask.org/en/latest/</u>

Dask (2/2)

- Familiar programming interfaces:

```
import dask.dataframe as dd # pandas-like
```

```
df = dd.read_csv("*.csv")
df.groupby(df.user_id).value.mean().compute()
```

 Lazy evaluation: build a graph of operations that is triggered by calling "compute()"

- Dask-ML extends scikit-learn with the same API
 - Can also be used with Pytorch, Keras, and Tensorflow using packages that bring a Scikitlearn like API to these frameworks (Skorch, SciKeras...)
- See <u>https://ml.dask.org/index.html</u>

Scikit-optimize (1/3)

- Hyper-parameter optimization library
- Based on scikit-learn but not limited to it, it can optimize any function with parameters
- Embeds Bayesian Optimization algorithm
 - Black box function (no gradients)
 - Expensive to evaluate
 - Observations are noisy
- Typically a neural network training
- To be installed with 'pip'

```
import skopt
import matplotlib.pyplot as plt
import numpy as np
DIMENSIONS = [
    skopt.space.Real(0.0, 1.0, name="x")
]
def f(x, noise_level=0.1):
    return np.sin(x[0]*5) * \
```

```
(1-np.tanh(x[0]**2)) \
+ np.random.randn() * noise_level
```

```
opt = skopt.Optimizer(DIMENSIONS)
```

```
for i in range(20):
    x = opt.ask()
    y = f(x)
    result = opt.tell(x, y)
print(result)
```

Scikit-optimize (2/3)

x* = 0.8252, f(x*) = -0.3961

×

0.0

1.25 -

1.00

0.75

0.50

0.25

0.00

-0.25

-0.50 -

0.0 0.2 0.4 0.6 0.8 1.0

μωθ0
 Observation

Scikit-optimize (3/3)

Thank you! Questions?

