

The exact-lab srl and CNR/IOM collaboration on deep learning algorithm and platform applied to SEM nanoimages : the NFFA case study

S. Cozzini, R. Aversa, C. De Nobili, A. Chiusole, G.B Brandino

CNR – IOM / eXact lab srl



Agenda

- Short introduction on eXact-lab and its activities
- The NFFA case study:
 - NFFA project and the data repository
 - The SEM image problem and its classification by means of ML/DL algorithm
- Conclusions



eXact


Stefano Cozzini – eXact lab srl

ENEA, Portici

07/02/2018

about us





Why a company providing scientific computing services?

Industry needs scientific computing to improve productivity and competitiveness

eXact lab is able to drive the needs and enable high performance solutions

Our mission

eXact lab offers solutions, training and on-demand access in the high performance computing market



HIGH PERFORMANCE COMPUTING
FOR YOUR BUSINESS



HIGH PERFORMANCE COMPUTING
FOR RESEARCH & DEVELOPMENT




HIGH PERFORMANCE COMPUTING
FOR THE PUBLIC SECTOR

Our vision

eXact lab wants to ease access to computational resources, for SME and public sector, dismantling the notoriety that still relegates HPC to an elite only playground



6 years activity



The team



eXact lab team

The core team

- 3 Founders: strong background in scientific computing
- 3 Collaborators: scientific software analyst, sys. adm , data scientist

Network of professional contacts

- Consultants/Professionals for specific tasks

Network of academic researchers

- Joint collaboration in R&D projects with local and European scientific partners



Services and products


We offer high performance services

- Infrastructure management
- Computing on-demand
- Coding for Big Data

Infrastructure management

eXact lab designs and manages computing resources, and offer consultancy to optimize your software stack

**Services and consultancy to
boost your resources and software**



ARPA FVG success story

eXact lab

- maintains the HPC resources for ARPA
- runs on its own infrastructure the Weather Model to provide business continuity
- Help research in improve/optimize simulation workflow at all levels

A full spectrum consultancy which notably increased the overall productivity



ARPA FVG

Agenzia Regionale per la Protezione
dell'Ambiente del Friuli Venezia Giulia

San Raffaele Hospital



Personalized medicine - customization of healthcare by use of genetic information

Our team built an high performance, highly available, computing infrastructure for genomic research



OSPEDALE
SAN RAFFAELE

Computing on-demand

eXact lab provides support and resources to tackle pressing computational problems without the burden of installing anything on local machines.

**Cloud computing services
anywhere and anytime**



C3HPC: HPC on-demand



C3HPC is a Cloud supercomputing infrastructure managed by eXact lab

It provides a complete software ecosystem able to tackle any kind of computing needs by means of modern software tools

Examples:

virtual workstation for remote visualization

Deep learning framework in containers

www.c3hpc.com



Virtual workstation on C3HPC

Virtual workstation

- GPU enabled!
- Elastic size
- Remote accessible from any device

Virtual workstation



C3 core infrastructure HPC

- Linux based
- Easy accessible

Vlznode



- Cloud based
- 2 GPU K2



Shared storage infrastructure

Docker container in HPC

- Docker containers as an alternative to Virtual Machines (VM)
 - Close to zero overhead (no virtualization of interfaces), crucial with parallel HPC workloads
 - Ease of communication between containers
 - Fast provisioning
 - No time spent in compiling and packaging scientific libraries on a HPC cluster



Docker on C3HPC

- Integration with the PBS Pro 12.2 queue system:
 - remove Docker containers at the job termination
 - Creation of an hook embedded into PBS Pro + a convention on the names of Docker containers to be removed after the job

```
docker run ... \  
--name example_${PBS_JOBNAME}_${PBS_JOBID}
```

Docker on C3HPC (2)

- 20 scripts in Bash and Python to speed-up the creation and usage of ML models on a HPC platform
- 2 different frameworks for ML and many models studied (CNTK and TensorFlow)

```
$ qsub -q gpuqueue -l nodes=1:ppn=24 \  
-v IMAGES_FOLDER=/images/,OUTPUT_FOLDER=/output/ \  
/scripts/run_fine-tuning_multiGPU_CNTK.sh
```

Coding for Big Data

eXact lab provides state of the art Machine/Deep Learning software stacks on any kind of infrastructure and develops tailored solutions for any Big Data problem.

**Advanced
analytics techniques
for any HPC application**



Foodword suite


Subject: italian food

Semantic analysis on different social media



Foodword suite

Real time monitoring of “food” concept



**A joint case study:
classifying SEM images by
Neural network within
NFFA-EUROPE project**

EU funded project


**it provides the widest range of tools for
research at the nanoscale**

Free transnational access to academia & industry

www.nffa.eu

The consortium

20 partners of which **10 nanofoundries** co-located with Analytical Large Scale facilities



The offer

TA

Transnational Access activities

Multidisciplinary research at the nanoscale performed at nano-laboratories and ALSFs

Integration of theory & numerical analysis with advanced characterization

NA

Networking activities

Interface for different user communities


Industrial exploitation of experimental data

JRA

Joint Research activities

Methods & tools at the frontier in nanoscience research

Improved infrastructures for academic & industrial projects



NFFA Data management

JRA3:


e-Infrastructure for data and information management

A transversal activity devoted to the setup of the first ***Information and Data Repository Platform*** (IDRP) for Nano science

- Definition of new metadata standards for data sharing in nanoscience
- Automatic acquisition of key metadata and create a data repository for future data access


Data infrastructure is complemented **by Data Analysis Services.**

NFFA IDR architecture




Easy data access from all facilities and via the NFFA portal for all NFFA users

NFFA IDRП deployment



Our Issue: SEM images

- One SEM Available at CNR-IOM Trieste with 150,000 images **NOT** classified
- 10 SEM across European partners: the work can be exported to a sizeable part of the community



Sharing images is nice..

- A couple of million nano images can be of some help for some nanoscience..



But before doing that we need to start classifying them...

SEM images classification steps

- STEP1: Classify images (scientific skills)
- STEP2: Train a neural network (deep learning)
- STEP 3: Use the network as classifier (inference)
 - Semi - Automatic tool for SEM users
 - Massive process of all the images

Step I: classify images..

We created and manually annotated ***the first dataset*** of classified SEM images (18,577 images).

Aversa et al., in preparation







Figure 1. Categories chosen for SEM images. The dimensionality of nanoscience objects provided the basis for the choice. Other categories, such as Biological and Tips were added as these were common images found in the SEM database.

Step 2: train the network !



Step2 : the tools/infrastructure...



Step2:The deep Learning network..

- **Models:**


- AlexNet
- Inception-v3/v4
- Densenet

- **Deep learning techniques:**

- Training from scratch
- Transfer learning (Feature extraction, Fine Tuning)

- **Deep learning frameworks:**


- TensorFlow
- Neon/Nervana




Glossary

- **Supervised learning:** labelled examples
- **Transfer learning:** applying knowledge of a trained network to a new domain
- **Check point:** set of parameters saved at a certain point of the training
- **Feature Extraction:** previous layers frozen to the check point + last layer(s) randomly initialized
- **Train from scratch:** all the parameters of all the layers are randomly initialized
- **Fine tuning:** all the parameters are initialized to the last check point and are allowed to vary


Feature extraction: ImageNet Checkpoint




Training from Scratch: alexnet




Training from Scratch: inc-v4



Training from Scratch: inc-v3



Which is the best ?




One epoch: it consists of *one* full training cycle through as much images as the training set size

Architecture	Test Accuracy [%]	Time to 470 eps (hours)	Stable at (hours)
Inc-v3, bs=16	95.3 ± 0.2	76.7	~ 76
Inc-v4, bs=16	94.0 ± 0.1	160.7	> 160

What about batch size ?


batch_size: the number of images in one forward/backward pass.



Architecture	Test Accuracy [%]	Time to 470 eps (hours)	Stable at (hours)
Inc-v3, bs=16	95.3 ± 0.2	76.7	~ 76
Inc-v3, bs=32	96.4 ± 0.2	69.8	~ 70


This is because with larger batches, less back-propagation are needed

Densenet vs Inception..



They seem very similar! But...

Densenet vs Inception




Same accuracy in more than half time!

Technical results

Synchronous Training


One step:

- CPU sends a model clone to each GPU which processes a batch of images and computes gradients.
- When both GPUs finish their computations, they send their result back to the CPU
- CPU averages gradients and updates the model for the next training step.



batch_size = 32, with two GPUs means an **eff_batch = 64**

MultiGPU scalability




Inception-v3 (imgs/sec)	1	2	4
Tesla K20	17.2	33.3	—
Tesla K40	23.0	42.1	78.5
GTX1070	40.9	56.0	66.4
DenseNet-28 (imgs/sec)	1	2	4
Tesla K20	45.5	81.9	—
GTX1070	53.5	100.3	183.7

DenseNet has a good scaling
also on GTX-1070

Neon/Nervana preliminary results


Neon: is a deep learning framework created by **Nervana-Systems** and **Intel**. Optimized on Intel CPUs using Intel Math Kernel Library (MKL)

In its final part, our project was selected by **Intel** as a relevant case study. We then had the opportunity to test **Neon** on Intel hardware (CINECA).



AlexNet: Forward and BackProp

Step 3: data analysis services




Step 3: massive image classification


- Massive Image Processing of all the 150,000 NFFA-EUROPE images
- Process:
 - Use a ML model specifically trained to classify the Images
 - Collect the 5 most probable categories for each image
 - Add the categories as metadata of the image
- Time requested: 150,000 images on 48 cores took 75 hours of computation

Tools: Apache spark cluster by docker integrated within LRMS

- 1 Docker container running the Apache Spark driver, on the first node
- 1 Docker container running an Apache Spark worker on each node, using the maximum amount of resources allocated by PBS Pro




Apache Spark speedup












Sem-classifier.nffa.eu

chunk 1 / image 1 /

Welcome, rossella 



Examples of categories



Insert another entry or click to select from the list

Save  and go back to the chunk list

Next 

Conclusion

- A positive collaboration has been setup between eXact-lab and CNR/IOM
- eXact lab is working to identify some key services for industry
- Big data/ML expertise should be converted in some advanced/customized services

References

- **R. Aversa**, “*Scientific Image Processing within the NFFA-EUROPE Project*”, MHPC thesis, 16-12-2016
- **C. De Nobili**, “*Deep Learning for Nanoscience Scanning Electron Microscope Image Recognition*”, MHPC thesis, 18-12-2017
- **H.M. Modarres, R. Aversa, S. Cozzini, et al.**, “*Neural Network for Nanoscience Scanning Electron Microscope Image Recognition*”, *Scientific Reports* **7**, 13282(2017)
- **R. Aversa, S. Cozzini**, “*The first annotated set of Scanning Electroscope Microscopy images*”, in preparation
- **A. Chiusole** “*Machine Learning As A Service Image recognition technologies on a scientific cluster*” BSc.Thesis University of Trento (2017)
- Web classifier: sem-classifier.nffa.eu

Thanks !