Intel® Unnati
Data-Centric Labs in Emerging Technologies

Give Your Students the Intel Edge.

Deep Learning

Today, there is a wide, and growing, skill gap between technical graduates and IT industry expectations. To propel India's digital economy transformation, it is imperative that the higher education system in the country bridges this gap by developing new curricula and offering courses in emerging technologies. The National Education Policy 2020[1] recognises this, and stresses the need for greater industry-academic linkages, and for higher education institutions to focus on research and innovation.

With the **Intel® Unnati Program**, you can keep pace with fast changing industry needs and expectations. It will help you:

Equip your students with industry relevant data-centric skills

In this age of data explosion, there is immense opportunity. Give your students the edge by equipping them with data-centric skills that will help them glean better insights and develop high-value solutions.

Unleash your students' creative potential

We, in India, have an incredible opportunity to unleash the creative potential of the largest student population in the world by training them in the right skills to drive India's digital transformation.

Build a strong reputation

With an Intel co-branded lab, you can be recognised as an institute that is committed to train your students in the latest technology to prepare them for industry, and focus on faculty development.

Build capability for the long term

Establish and maintain your leadership with the help of our **System Integrator Associates**. From Intel's recommendations for end-to-end technology labs set up and course content to training, customisations of your lab set up, or your maintenance and support requests, you can rely on them for all your needs.

Winning with the Intel® Unnati Community

With an Intel® Unnati Lab, you—and your faculty and students—become part of the **Intel® Unnati Community**, and get exclusive benefits:

- **Intel® Unnati Grand Challenge** and **Intel® Unnati Rapid Challenge**, where your students solve industry relevant, high impact problems through technology, with cash awards and the opportunity to be evaluated for internships at Intel
- **Intel® Unnati Ignite** workshops that offer hands-on experience with Intel technologies
- **Intel® Unnati Catalyst** sponsorships for select tech contests and events that you conduct, that encourage the use of new technologies
- **Intel® Unnati Research Launchpad**, which offers grants to your faculty members for original research in new and emerging technologies

[1]Ministry of Human Resource Development, Government of India, National Education Policy 2020, *https://static.pib.gov.in/WriteReadData/userfiles/NEP_Final_English_0.pdf*

Intel® Unnati
Deep Learning
with Habana® Gaudi®

Deep Learning: Transforming the way the world works

Deep Learning (DL) is at the very core of the way businesses work today—indeed, a World Economic Forum article[1] terms it "the future of business". It touches everything from automation of processes and tasks for greater efficiency and reliability to predictive analyses for customer retention.

If you are looking to equip your students and faculty on the latest industry-relevant technologies in the DL space, the **Intel® Unnati Online Lab featuring Deep Learning with Habana® Gaudi®, available on Amazon Web Services (AWS)**, is very relevant for you. It is ideally suited for beginner to intermediate level engineering and science students and faculty.

[1]Davos Agenda 2022, *How deep learning can improve productivity and boost business*, January12, 2022. *https://www.weforum.org/agenda/2022/01/deep-learning-business-productivity-revenue/*

Train More, Spend Less

Habana® Gaudi® processors are poised for rapid industry adoption as they offer 40% better price performance than current GPU-based EC2 instances for Machine Learning (ML) workloads. According to AWS, they offer the "best price performance for training Deep Learning models in the cloud".

Habana® Gaudi® is an AI processor designed from the ground up to accelerate DL training workloads. It is supported by the SynapseAI® Software Suite optimised for Habana® Gaudi® and DL workloads, for DL model development, and to ease migration of existing GPU-based models to the Habana® Gaudi® platform hardware. Among its key features:

- It offers up to 40% better price performance than latest GPU-based instances (*see performance benchmarks below*).
- The Habana® SynapseAI® software platform integrates TensorFlow* and PyTorch*.
- It is easy to port existing TF models to Habana® Gaudi® with minimal code changes
- While this lab is based on the Habana® Gaudi® v1 processor, Habana® Gaudi® v2 processors are out-performing GPUs in terms of training throughput by nearly a factor of 2, making a strong case for an alternative to GPUs for DL model training (*see performance benchmarks on the opposite page*).

Cost Savings with DL1 Model Training

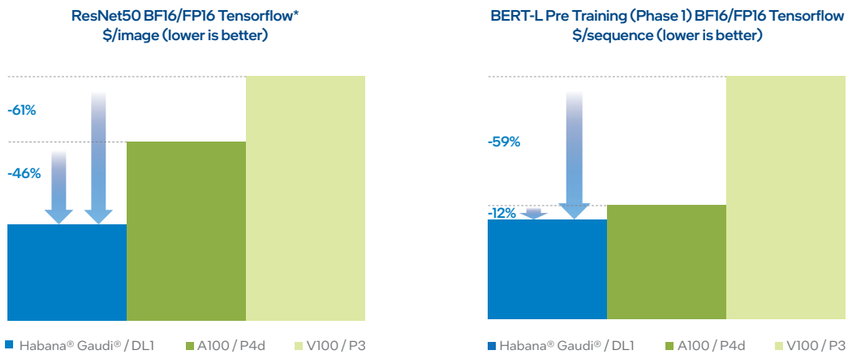

Cost savings based on Amazon EC2 On-Demand pricing for P3, P4d and DL1 instances respectively. Performance data was collected and measured using the following resources:

Habana BERT model: *https://github.com/HabanaAI/Model-References/tree/master/TensorFlow/nlp/bert*

Habana ResNet50 Model: *https://github.com/HabanaAI/Model-References/tree/master/TensorFlow/computer_vision/Resnets/resnet_keras*

Habana Test Container: *https://vault.habana.ai/ui/native/gaudi-docker/1.0.1/ubuntu18.04/habanalabs/tensorflow-installer-tf-cpu-2.5.1/1.0.1-81*

A100 / V100 Benchmark Sources: *https://ngc.nvidia.com/catalog/resources/nvidia:bert_for_tensorflow/performance, https://ngc.nvidia.com/catalog/resources/nvidia:resnet_50_v1_5_for_tensorflow/performance*

Results published on DGXA100 and DGX-1

The AWS Habana® Gaudi® Deep Learning Training Lab

AWS EC2 DL1 Instance (*https://aws.amazon.com/blogs/compute/amazon-ec2-dl1-instances-deep-dive/*)

Instance Size	vCPU	Instance Memory (GiB)	Habana® Gaudi® Accelerators	Network Bandwidth (Gbps)	Total Accelerator Interconnect (Gbs)	Local Instance Storage	EBS Bandwidth (Gbps)
d1.24xlarge	96	768	8	4 x 100 Gbps	700	4 x 1TB NVMe	19

- Available on AWS EC2 DL1 instances
- Content: ~35 hours in total, with 15-20 hours of hands-on exercises
- Includes exercises in Computer Vision, Natural Language Processing and Distributed training
- Exercises are in Jupyter* notebook format
- Includes quizzes

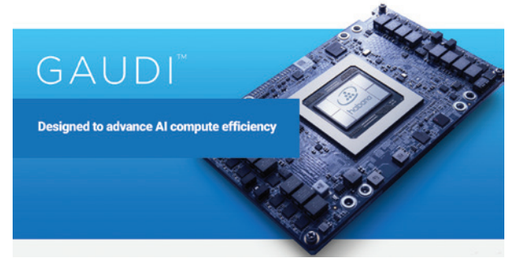

Features of DL1 Instances Powered by Habana® Gaudi®

AWS Custom 2nd Gen Xeon Scalable Processors

Up to 8 Habana Gaudi accelerators with 32GB HBM per processor

400Gbps Networking & 4TB of NVMe Storage

All-to-all 100Gbps interconnect

SynapseAI SDK integrated with TensorFlow and PyTorch

Support for developing custom kernels

Habana® Gaudi® v2 Processor: Outstanding Training Throughput

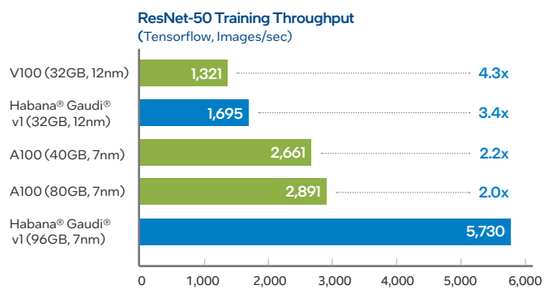

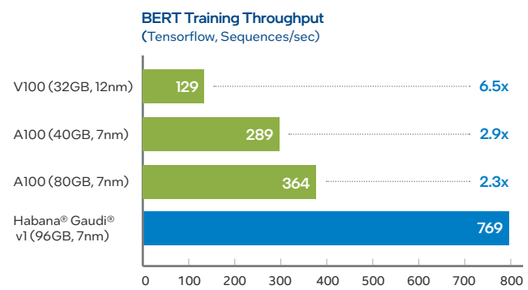

To know more about how your institution can benefit from the Intel® Unnati Program, please contact:

Intel_Unnati_DL_Jan2023