FREQUENTLY ASKED QUESTIONS

For AI Jump Start Applicants

A. WHAT KINDS OF BUSINESSES CAN APPLY?

B. WHAT IS THE TIMELINE FOR THE PROGRAM, INCLUDING DEADLINES FOR APPLICATION ?

C. WHAT WILL HAPPEN AT THE AI JUMPSTART WORKSHOP ?

D. WORKING WITH AI JUMPSTART FACULTY EXPERTS

E. ESTIMATED DURATION OF PROJECT ENGAGEMENT WITH FACULTY AND OVERALL PROJECT

F. RESOURCES FOR DEVELOPING AND IMPLEMENTING NEW BUSINESS OUTCOMES AT YOUR COMPANY

G. COMPUTING RESOURCES AVAILABLE TO AI JUMP START COMPANIES

H. WHAT IS COLOSSEUM ?
   a. SUMMARY DESCRIPTION OF FUNCTIONALITY
   b. SUMMARY TECHNICAL DETAIL
   c. ACCESSIBILITY AND FEATURES

I. WHO CAN I CONTACT IF I HAVE QUESTIONS ?

==========

A. WHAT KINDS OF BUSINESSES CAN APPLY ?

Any product manufacturer or service company currently operating in Massachusetts that is seeking to introduce, upgrade, or advance its use of data-based computing solutions to improve its competitiveness and profitability.

B. WHAT IS THE TIMELINE FOR THE PROGRAM, INCLUDING DEADLINES FOR APPLICATION ?

Initial on-line submissions due on Monday 10 May 2021.

Applicants will be notified within 30 days or sooner if selected to attend the AI Jump Start Workshop scheduled for Friday 11 June 2021.

Following the workshop, supplemental applications can be submitted anytime up to Friday 9 July.

Companies which submit supplemental applications will be notified of a funding/no funding decision by Monday 16 August. Depending on final review of the submissions, the program intends to fund the participation of 20-25 companies in this first cohort of AI Jump Start Program members.

Final approval notices will include the name and contact information of assigned faculty experts for companies to contact to initiate their projects.

C. WHAT WILL HAPPEN AT THE AI JUMPSTART WORKSHOP ?

Detailed information and instructions will be provided only to companies selected to participate in the Workshop. Faculty presentations on the objective and management of the program will be designed to help companies prepare and submit their applications

D. WORKING WITH AI JUMPSTART FACULTY EXPERTS

The project has enlisted ~ 35 experts among the faculties of Northeastern, Tufts, and Boston Universities.

Examples of the technical and developmental expertise include (but are not limited to) the following:

- ✓ Machine Learning
- ✓ Deep Learning
- ✓ Computer Vision and Image Processing
- ✓ Natural Language Processing/text analysis
- ✓ Signals, Sensing and Spatio-Temporal Analysis
- ✓ Data Mining
- ✓ Data Integration and Cleaning
- ✓ Data Management
- ✓ High Performance Computing and Acceleration
- ✓ Edge Computing
- ✓ Distributed Data Processing and optimization
- ✓ Embedded Systems
- ✓ Cloud Computing
- ✓ AI in Health
- ✓ Imaging Systems
- ✓ Robotics, control, autonomous systems
- ✓ Wireless Systems
- ✓ Internet of Things

Initial budgeting is based on a discounted hourly consulting rate of $100 per hour indicating that, on average, $4,500 can be made available to participating companies to pay for the initial consultation with assigned faculty experts. This estimated amount will pay for ~45 hours of consulting and development time for a company's project. It is anticipated that this initial consulting period may lead to further collaboration between companies and consultants and that they will define those terms as needed.

E. ESTIMATED DURATION OF PROJECT ENGAGEMENT WITH FACULTY AND OVERALL PROJECT

The duration of the faculty expert engagement and overall project will vary somewhat from company to company. A useful, preliminary estimate is @ 45 hours to be allocated for a period up to one month. This initial funding will "jump start" the relationship between the company and the faculty consultant to pave the way for further collaboration to be discussed and defined by the company and the consultants."

F. RESOURCES FOR DEVELOPING AND IMPLEMENTING NEW BUSINESS OUTCOMES AT YOUR COMPANY

Today's AI/ML applications require an ever-increasing amount of computer horsepower to solve AI/ML tasks. The AI Jump Start Program is seeking to address needs of

Massachusetts' industry by enabling researchers at Northeastern, BU and Tufts to demonstrate the power of AI/ML. We seek to accomplish this by helping businesses such as yours operate in the Commonwealth more efficiently, effectively and more profitably by thriving in the era of AI/ML computing.

G. COMPUTING RESOURCES AVAILABLE TO AI JUMP START COMPANIES

The computing equipment and associated software tools assembled for AI Jump Start are being combined with customized solutions expertise from the faculty at Northeastern, Tufts and Boston University who are affiliated with or participating in this program. By learning about your product and operating challenges, we will help assess, design and implement computing solutions using AI/ML methodologies. This pilot program, which is undertaken in partnership with the John Adams Innovation Institute, a division of the Commonwealth's MassTech Collaborative, is helping to extend the long tradition of university-based researchers at Northeastern, BU and Tufts supporting Massachusetts companies by engaging in consulting engagements with these businesses.

The Northeastern AI/ML Cluster was acquired through a grant from the John Adams Innovation Institute/ Massachusetts Technology Collaborative. The cluster

provides a distributed computing solution, with resource installed in the Massachusetts Green High Performance Center and the Northeastern University Burlington Innovation Campus. This project collaborates with a number of research centers and laboratories/institutes at Northeastern including the ALERT DHS Center of Excellence, the WIOT Institute, NUCAR and SPIRAL.

H. COMPUTING RESOURCES AVAILABLE TO AI JUMP START COMPANIES

The computer cluster at Northeastern University available to participants of the AI Jump Start Program combines the latest high-performance computing technology to address next-generation challenges in the applications of artificial intelligence/machine learning (AI/ML). It includes state of the art computing capabilities and cutting edge AMD and NVIDIA platforms for companies to learn about and apply to their individual business needs.

The AI/ML cluster includes 7 state-of-the-art NVIDIA DGX-A100's, each capable of delivering 5 petaflops for deep learning applications. The AI/ML cluster also offers 14 nodes of AMD CPU/GPU heterogeneous computing, with 8 MI-50 GPUs, each rated at 26.5 TFLOPs for FP16, and 13.3 TFLOPS for FP32. The system includes a large memory node (3TB of DDR4 RDRAM).

In addition, the system is closely tied to the university's Colosseum Cluster which provides computational support for cutting-edge research in 5G and 6G technologies.

I. WHAT IS COLOSSEUM AND WHAT ROLE DOES IT HAVE IN THIS PROGRAM ?

Colosseum is the world's largest RF emulator, designed to support research and development of large-scale, next generation radio network technologies in a repeatable and highly configurable RF environment.

Created by an investment of tens of millions of dollars by the U.S. Department of Defense Advanced Research Projects Agency (DARPA) – and now owned and operated by Northeastern's Institute for the Wireless Internet of Things at the Burlington campus – Colosseum emulates wireless signals, with granularity at the RF signal level, within a custom-built, massive RF computational facility.

Colosseum can create virtual worlds for testing devices using RF, creating an artificial environment as if radios are operating in an open field, urban area, shopping mall, or a desert, simulating signals traversing space and reflecting off multiple objects and obstacles as they travel from transmitters to receivers by generating more than 52 terabytes of data per second.

SUMMARY DESCRIPTION OF COLOSSEUM FUNCTIONALITY

Colosseum provides the users with preconfigured and ready to use LXC containers for basic testing. Users can customize these containers based on their needs to develop and implement their own radio codes. These customized container images can be uploaded to the Colosseum servers, and the user may choose to load the default or the customized containers onto the SRN's during their reserved session. This allows for completely customizable RF testing with a variety of repeatable scenarios that enables the users to conduct virtually any type of research on waveforms and different layers of protocol stack as well as networked applications, among others.

As an RF testbed, Colosseum enables full stack wireless networking research from applications down to the physical layer at scale - with up to 256 interacting radios. It can be utilized to:

- Emulate multiple operational environments including a 1 sq. km open field, a dense urban city, a suburban shopping mall, a desert, or anything in-between.
- Emulate in real-time multipath and fading effects with high-fidelity ray-tracing.
- Support high-fidelity and large-scale research on waveforms, protocols at all layers, networked applications, jamming and security, MIMO, and beamforming.
- Provide full-stack repeatable environment (from RF to application layer).

- Carry out large scale testing, up to 256 radio nodes with 256x256 configurable channels.
- Support cellular networks (4G/5G and beyond), IoT, cognitive radio, ad hoc networks, edge computing, and cloud RAN research.
- Implement Machine Learning algorithms in different wireless communications techniques such as spectrum sharing, dynamic spectrum access, extraction of signal intelligence and optimized routing, by providing built-in powerful computational resources.

COLOSSEUM SUMMARY TECHNICAL DETAIL

- Colosseum combines 128 Standard Radio Nodes (SRNs) with a Massive digital Channel Emulator (MCHEM) backed by an extensive FPGA routing fabric.

    - Each of the SRNs provides a platform for Software Defined Radio as well as Machine Learning applications with two key hardware components, namely, a Dell R730 Server which comes with an NVIDIA K40M GPU, and an Ettus Research USRP X310 Software-defined Radio that is equipped with a XILINX Kintex 7 FPGA.

  The MCHEM facilitates real-world wireless RF channel emulation between the SRNs and can emulate fading, multipath, etc., for up to 256 x 256 independently customizable channels. This allows for large scale RF testing with up to 256 independent radio nodes each with powerful computational capabilities.

- Colosseum can be used to create realistic city-scale 5G/6G scenarios. As an example, a Rome emulation scenario has been created by programming into Colosseum the GPS coordinated of real-world cellular towers deployed in Rome, Italy, derived from an open access database of cellular base stations. The scenario involves 50 nodes (with up to 10 cellular base stations) serving pedestrian users deployed at different distances from the base stations (e.g., close, medium and far distance), and moving at different speeds (e.g., static nodes, low and moderate speed).

- Accessible as a cloud-based platform, Colosseum also provides other resources to create a real-time high-fidelity radio scenario such as traffic generation, timing, and GPS synthesis.

- Allows 256 x 256 100 MHz RF channel emulations, with USRP X310 radio nodes in the loop interacting in over 65,000+ channels at the same time.

For additional information, about Colosseum please see:

https://colosseum.net or
https://www.northeastern.edu/wiot

COLOSSEUM ACCESSIBILITY AND FEATURES

- Colosseum operates 24/7/365.
- Can be accessible remotely to AI Jump Start users.

- Users can reserve Colosseum resources through a simple web interface.
- Emulation jobs can be done either manually in an interactive session during the scheduled time
    - Alternatively, the user may choose to create batch jobs that run automatically per user's instructions.

For additional information about Colosseum, please see: colosseum.net

J. WHO CAN I CONTACT IF I HAVE QUESTIONS ABOUT AI JUMP START ?

AIJumpStart@MassTech.org