Semi-Supervised On-Device Neural Network Adaptation for Remote and Portable Laser-Induced Breakdown Spectroscopy

On-Device Intelligence Workshop

April 9, 2021

Kshitij Bhardwaj and Maya Gokhale

LLNL-PRES-XXXXXX

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Laser-Induced Breakdown Spectroscopy (LIBS)

- Fast elemental analysis technique: determine chemical composition
- Several applications:
  - Industrial analysis of metals, geological research, forensic analysis
  - Most notably: space exploration. Used in *Mars Perseverance rover*
- <u>Rise of portable LIBS: both remote and handheld devices</u>
  - Require fast on-device data processing in real time
  - ThermoFisher's battery-operated LIBS: perform alloy identification in 10 secs

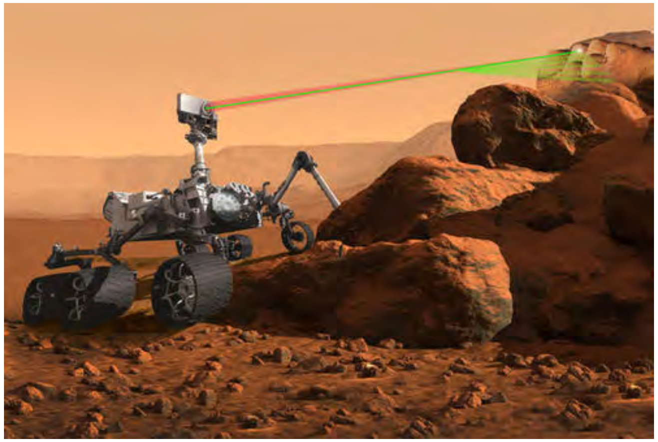

Supercam LIBS instrument on Perseverance [NASA, Mars 2020]

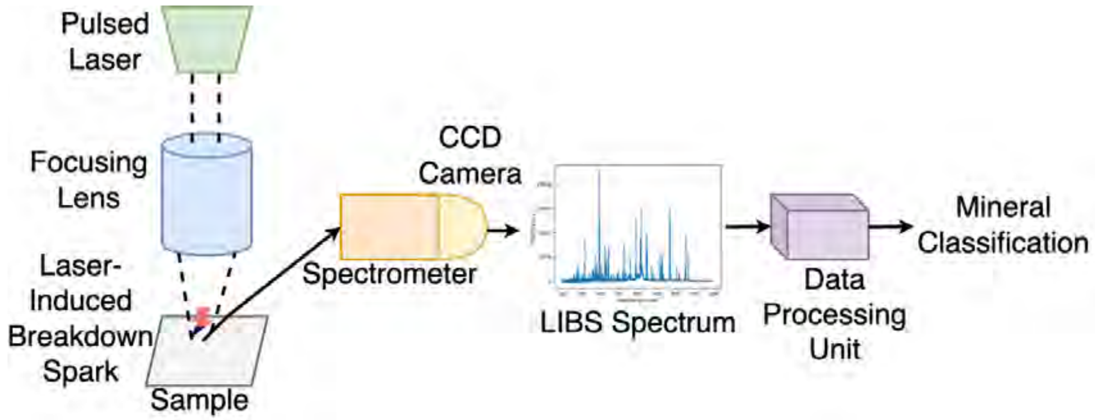

An example LIBS apparatus [Sensei et al., 2021]

Machine Learning in LIBS Systems

- Rise in the use of ML for LIBS [Chen et al., 2020] [Vrabel et al., 2020]
  - Fast, effective in handling high-dimensional and complex LIBS spectra
  - E.g., random forests, support vector machines, and neural networks

- Challenges with ML models for battery-operated/portable and remote LIBS:
  - Model should be lightweight while achieving high accuracy
    - Reduced memory and power consumption with fast predictions
  - Model must be able to efficiently handle *domain shift* without supervision
    - *Especially challenging for remote LIBS:* limited access to the device and the ML model needs to self-adapt without new labeled data
    - Domain shift due to encountering spectra with different distributions than training
    - Dynamic environmental or sensor noise

Example of Domain Shift in LIBS Data

- Use a well-known LIBS dataset [Kepes et al., 2020]
  - 138 soil mineral samples, each belonging to one of 12 mineral classes
  - There are 500 spectra for each soil sample
  - Each spectrum has 40002 wavelength points
- Example of domain shift: 2 spectra
  - Belong to the same mineral class
  - Considerably different probability distributions
  - Training on samples similar to 1 but testing on 2 may lead to incorrect results

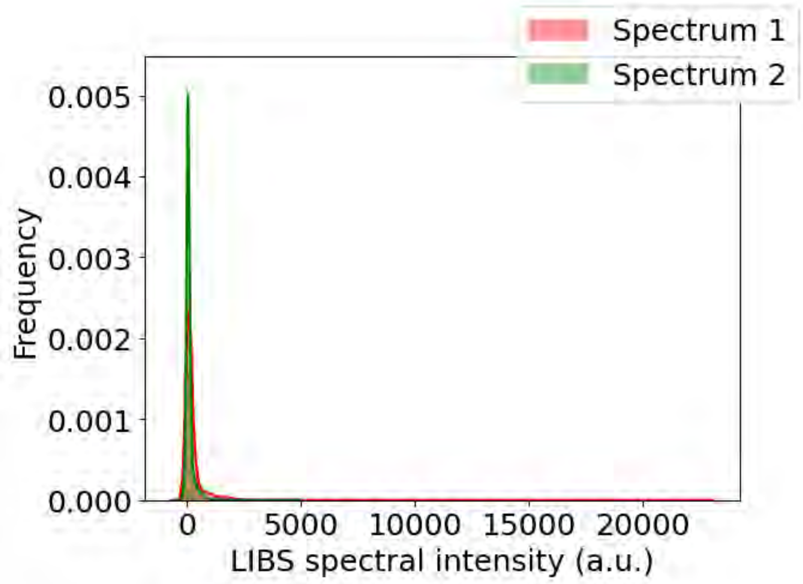

Contributions

- A new lightweight multi-layer perceptron (MLP) model for LIBS
  - MLP-LIBS has 2 hidden layers with 64 neurons each
  - Achieves an average accuracy of 88.2% for LIBS dataset
    - On par with other models [Vrabel et al., 2020] that do not handle domain shift
- Extend MLP-LIBS to handle domain shift: MLP-LIBS-ADAPT
  - Model adaptation (or retraining) after deployment in semi-supervised manner
    - No labels needed for new, possibly domain-shifted, inputs
- For a data streaming case study
  - MLP-LIBS-ADAPT shows up to 2.1% better accuracy than MLP-LIBS
  - Characterized inference and retraining time on Google pixel2
- Propose a new heterogeneous accelerator architecture for MLP-LIBS-ADAPT

MLP Model for LIBS

- Data preprocessing
  - — Dataset: 40002 wavelength points, 12 mineral classes [Kepes et al., 2020]
  - — Normalization to scale values to [0,1]
  - — Dimensionality reduction using UMAP [McInnes et al., 2018]
    - For each spectral sample: from 40002 wavelength points to 100 features

- Lightweight MLP-LIBS model:

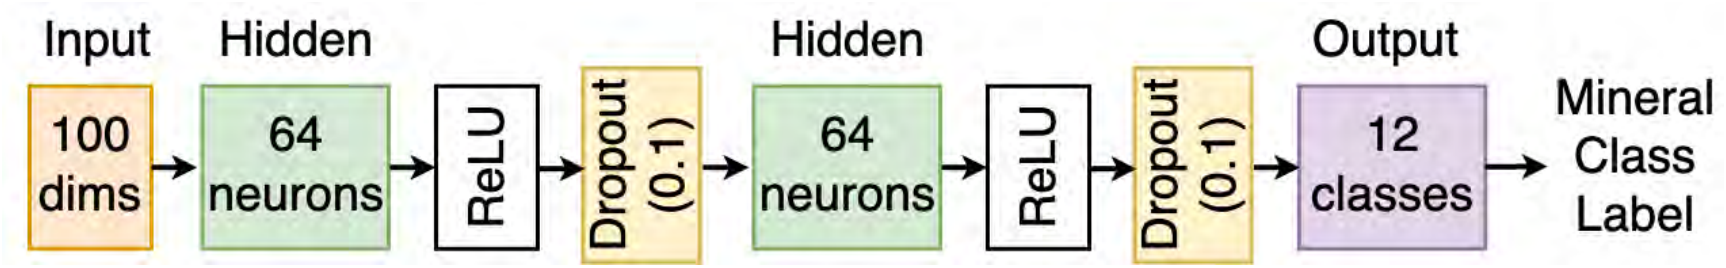

Semi-supervised MLP Adaptation for LIBS

- Extend MLP-LIBS using adaptation technique based on standard backpropagation [Ganin & Lempitsky, 2015]
  - Add a gradient reversal layer
  - Add a domain classifier: classify input spectrum as
    - From source domain (labeled data on which the model was initially trained on)
    - From target domain (new, possibly domain-shifted, unlabeled data)
- Lightweight MLP-LIBS-ADAPT model:

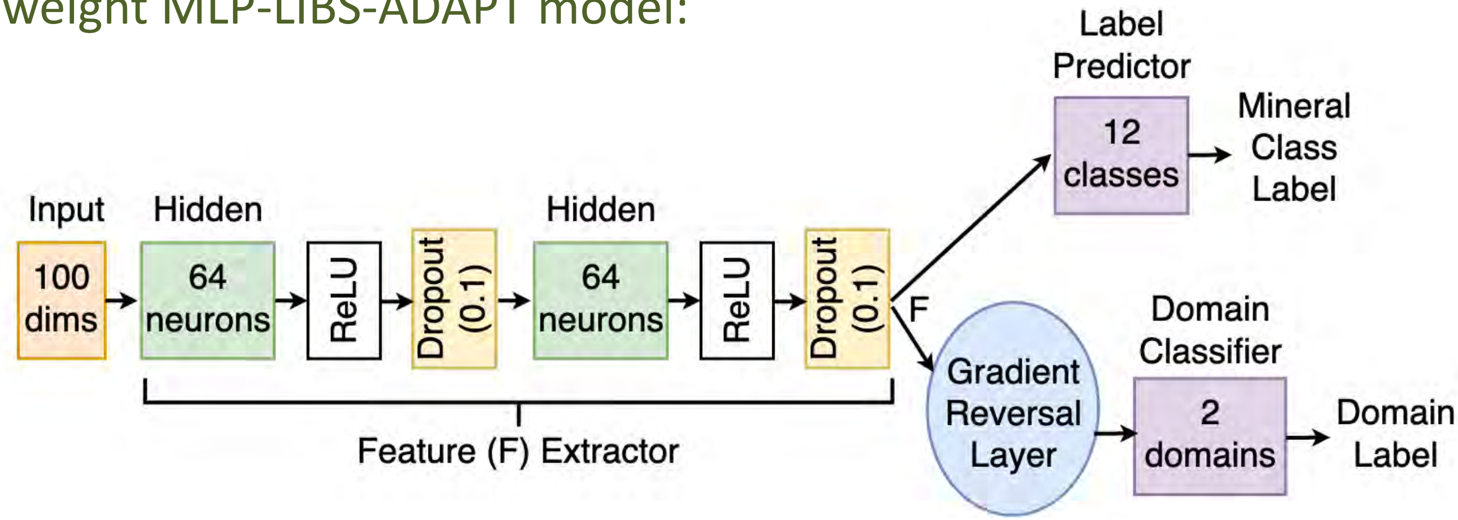

On-Device Adaptation of MLP-LIBS-ADAPT

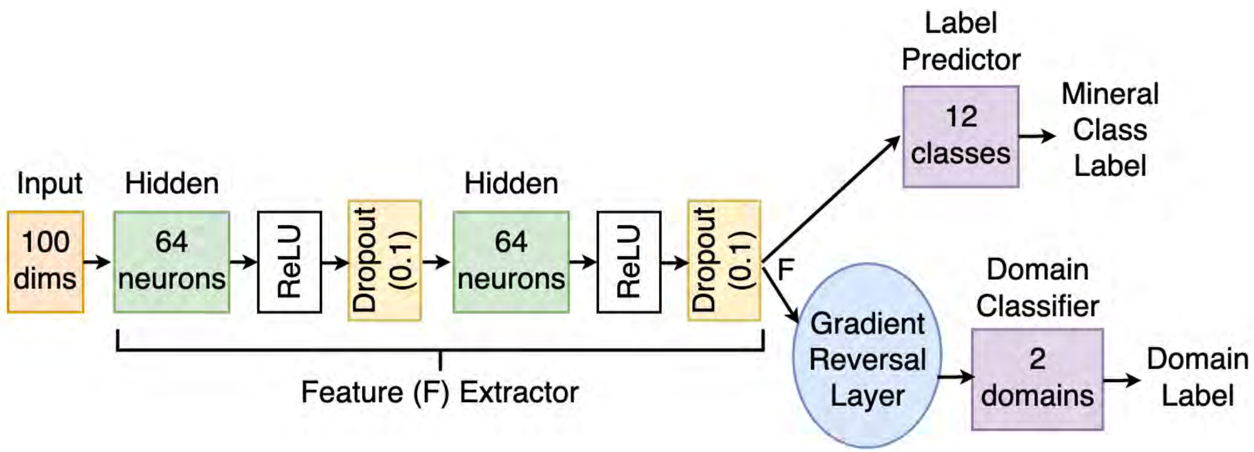

- Offline: initial training using labeled LIBS data (source domain)
- After deployment, retrain at user-defined instances using 2 datasets
  - *__Labeled long term memory (LLTM):__* initial offline labeled LIBS spectra (source)
  - *__Unlabeled short term memory (USTM):__* recently seen new unlabeled data (target)
- Basic idea: *during on-device retraining, learn to make classification decisions without being hindered by shift in the two domains*
  - Gradient reversal layer multiplies gradient with a negative constant

Model feature (F) extractor learns features that:

- Minimize label prediction loss (features are discriminative) and maximizes domain prediction loss (features are domain-invariant)

Experimental Setup

- MLP-LIBS and MLP-LIBS-ADAPT are initially trained/validated using LIBS dataset
- Amount of initial training data varied: 20000, 18000, and 4000 spectra
- Separate 25000 spectra in the LIBS dataset used to simulate data streaming for testing
  - Prequential evaluation (interleaved test-then-train) used for MLP-LIBS-ADAPT
    - Accuracy reported every 2500 spectra (data chunks), followed by retraining using the most recent USTM and pre-stored LLTM
  - No retraining for MLP-LIBS: only report accuracy for each data chunk
- For retraining MLP-LIBS-ADAPT, LLTM and USTM also varied:
  - LLTM same as the initial offline training data or some fraction of it
  - USTM: recently seen 50 or 100 unlabeled spectra

MLP-LIBS-ADAPT vs. MLP-LIBS: Sensitivity to LLTM and USTM

- Vary initial offline training data (i), LLTM (L), and USTM (U): <u>LLTM same as initial data</u>

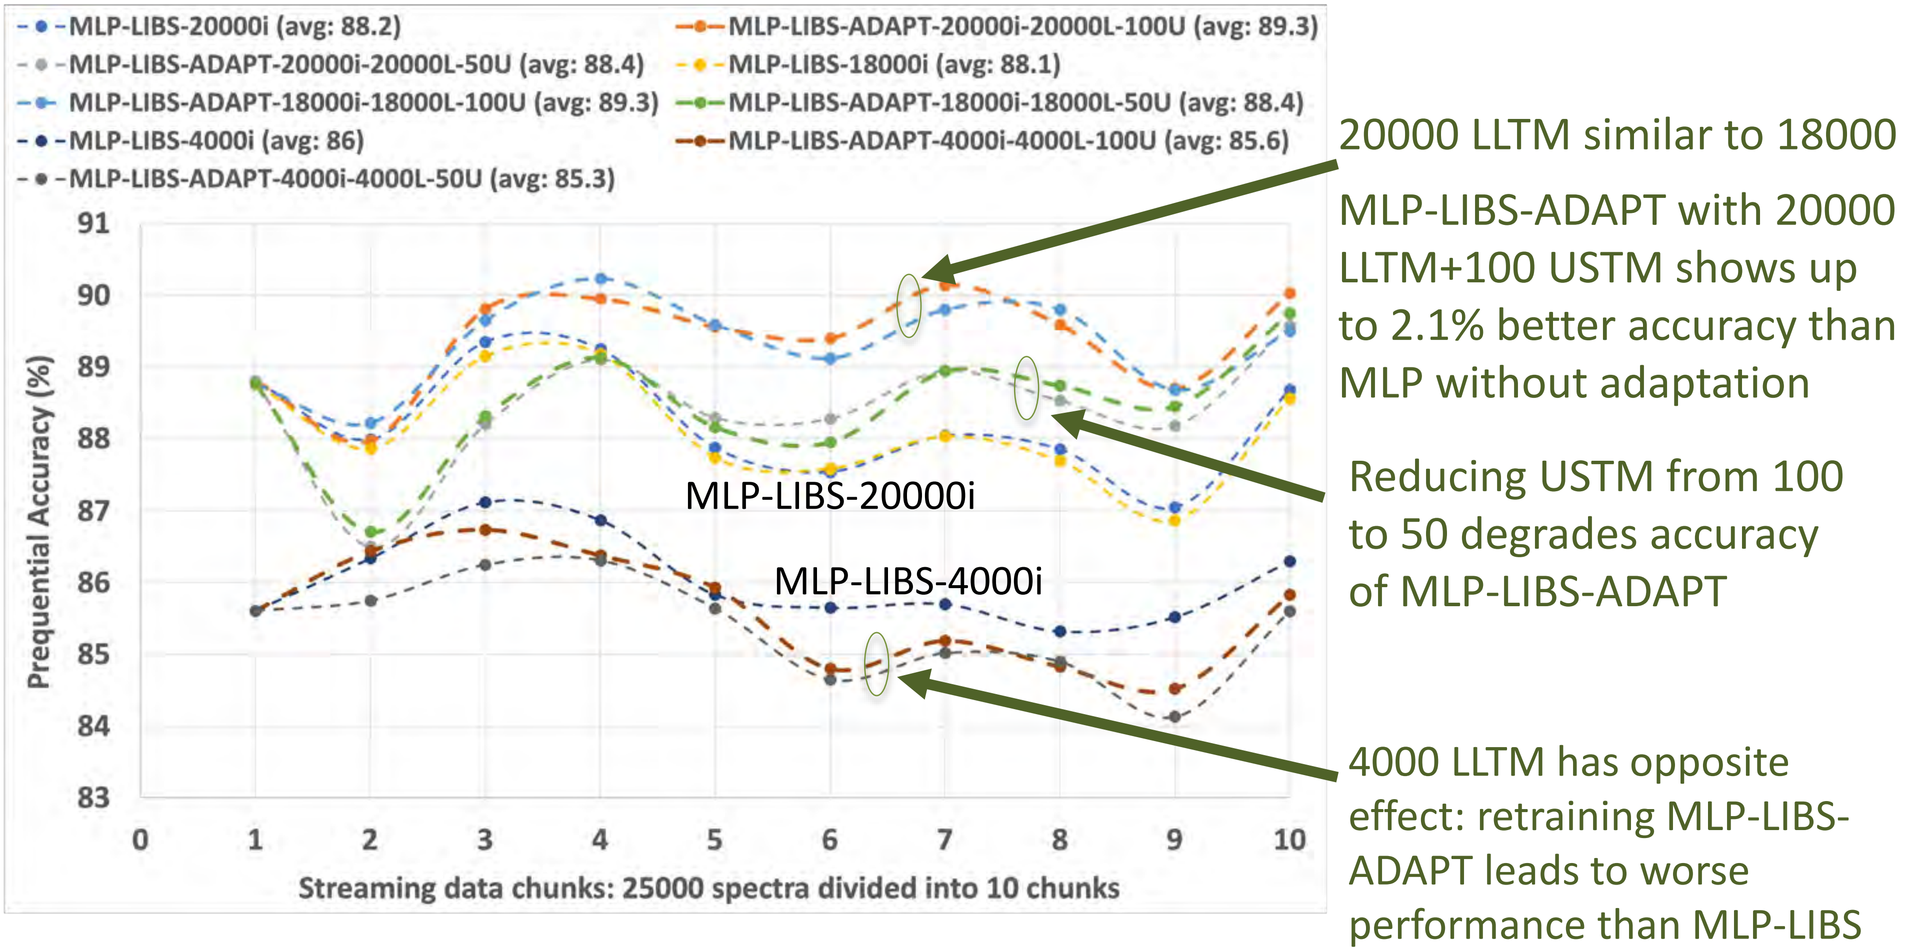

Effect of Using LLTM as Fraction of Initial Data

- Using different fractions of initial offline training data (i) as LLTM (L)
- Initial data fixed to be 18000, LLTM varied as 18000, 14400, 9000, and 4500
- <u>Accuracy of MLP-LIBS-ADAPT suffers if LLTM is not equal to initial training data</u>

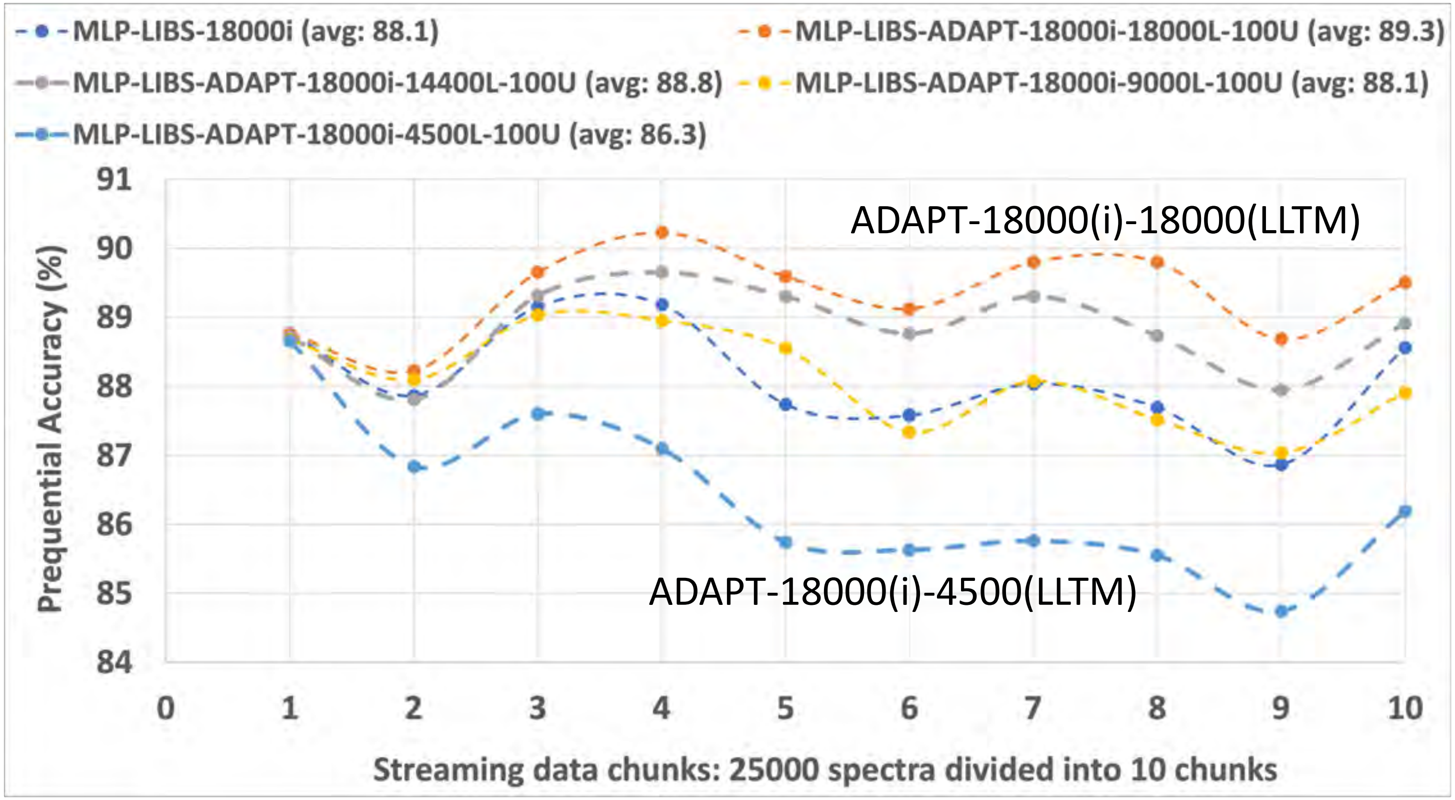

Model Performance on Google Pixel2

- CPU only runs: Qualcomm Kryo 280 processor (similar to Arm A73/A53)
- Ran multiple configurations of MLP-LIBS-ADAPT and selected the best one
  - In terms of accuracy, memory requirement, and MLP processing times
- Best configuration selected: LLTM as 18000 and USTM as 100
  - 89.3% average accuracy during data streaming
  - Avg. inference time: 0.097s, avg. retraining time: 599s
- LLTM as 20000: same avg. accuracy but requires more memory
  - Higher avg. retraining time: 682.2s
- LLTM as 4000: very small memory with avg. retraining time: 126.5s
  - Shows significant accuracy loss (avg. accuracy: 85.6%)

Heterogeneous Accelerator Design for LIBS

- Heterogeneity both in terms of memory and compute is required
- High-density non-volatile memories (NVMs) to store read-only LLTM data, SRAM scratchpad for USTM
- Specialized MLP accelerator to perform fast on-device training and inference

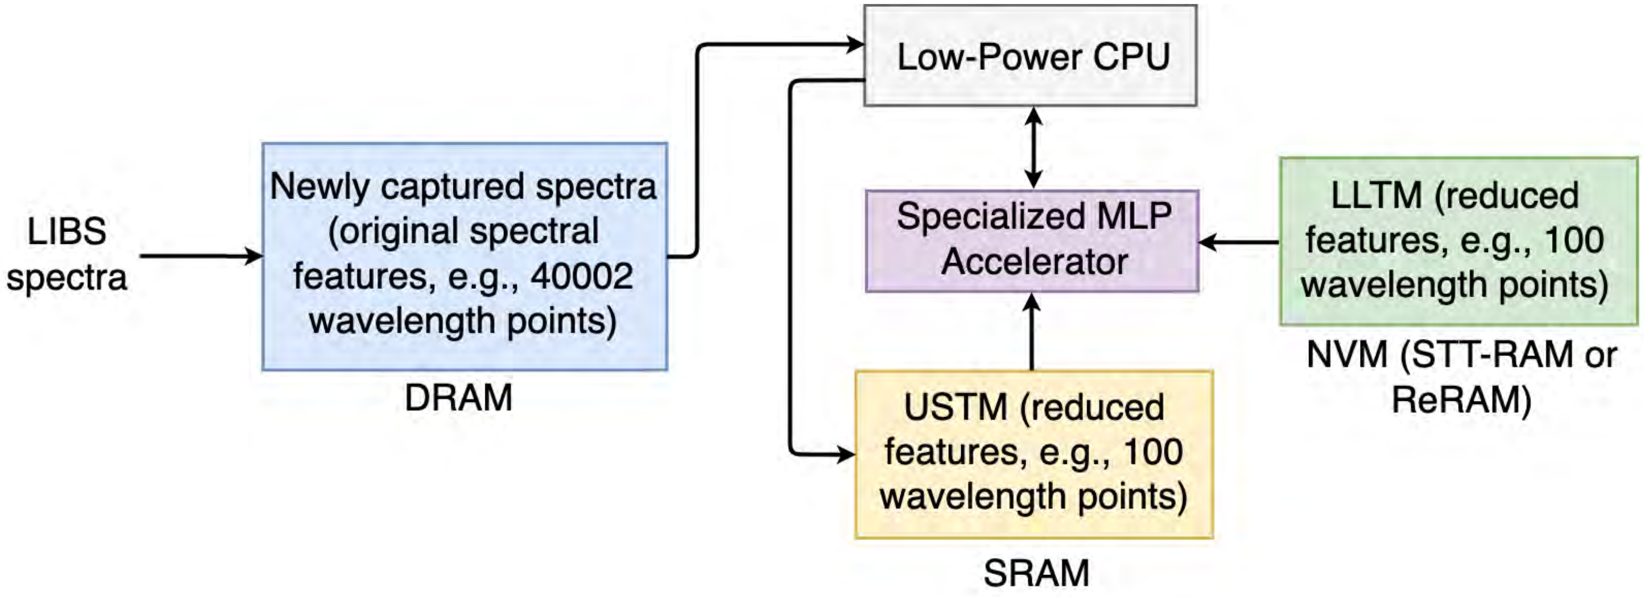

Conclusion

- A new, lightweight MLP-LIBS-ADAPT for portable and remote LIBS systems
  - Adapt to domain shift in a semi-supervised manner
- Achieves an average accuracy of 89.3% during data streaming
  - Up to 2.1% better than an MLP model without support for adaptation
- For effective model adaptation
  - Labeled long term memory: equal to initial offline training data
  - Unlabeled short term memory: sufficient and not too small
- Characterize training/inference time on Google Pixel2
  - Select best configuration in terms of accuracy, memory, and model processing time
- Proposed a heterogeneous accelerator to efficiently run MLP-LIBS-ADAPT

Lawrence Livermore National Laboratory

Backup

Example LIBS Spectrum

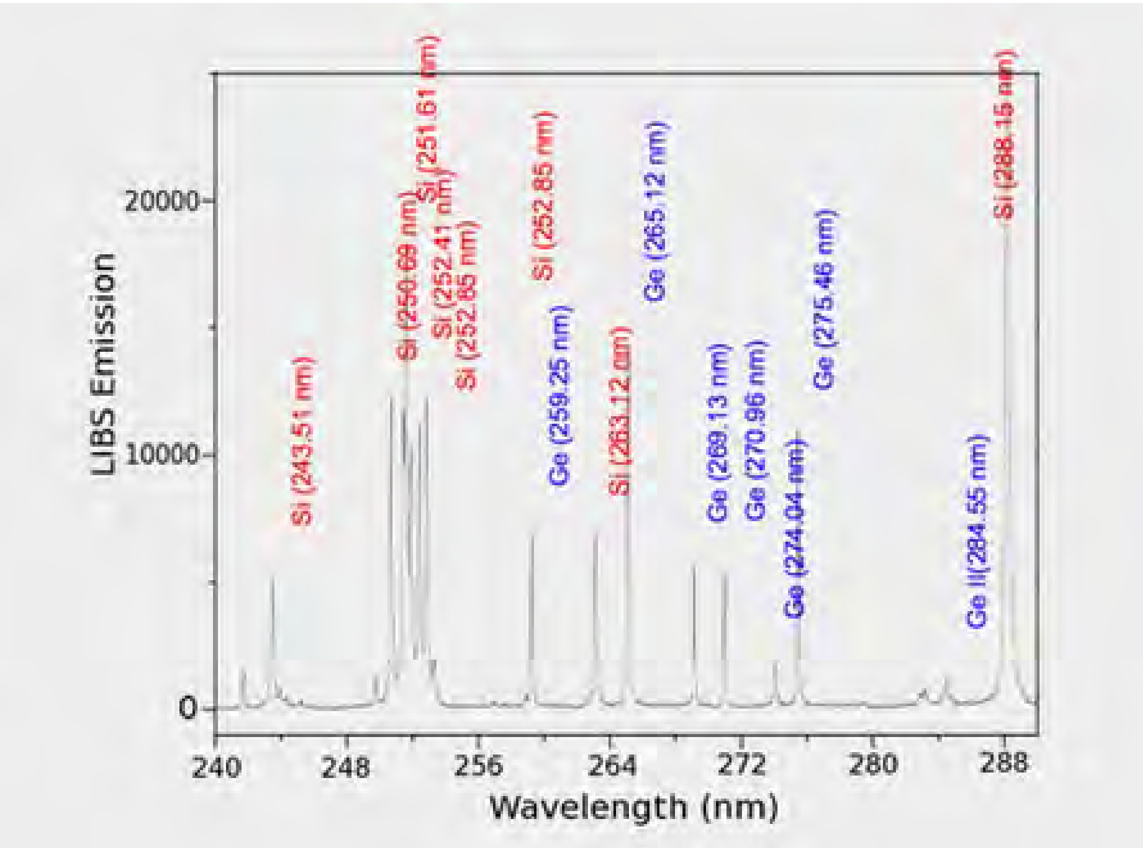

Each element in the periodic table is associated with unique LIBS spectral peaks. By identifying different peaks for the analyzed samples, its chemical composition can be rapidly determined. Often, information on LIBS peak intensities can be used to quantify the concentration of trace and major elements in the sample.

Source: appliedspectra.com