Applications of Machine Learning and Artificial Intelligence in Modern Networks

Special track in the 16th International Conference on Digital Telecommunications, ICDT 2021, April 18 – 22, 2021 in Porto, Portugal

Stan McClellan
Ingram School of Engineering
Texas State University
San Marcos TX, USA
stan.mcclellan@txstate.edu

***Abstract*—This paper summarizes six presentations in the special track "AMAIN: Applications of Machine Learning and Artificial Intelligence in Modern Networks". The research deals with the following key issues :**

- **Optimization of data handling via hardware offload such as FPGA for large, sparse datasets.**
- **Architectures and techniques for loosely coupled, distributed systems (e.g. Internet of Things).**
- **The use of Machine Learning (ML) in powerline communications and the corresponding use of complex-valued training data in ML algorithms.**

The contributions in this track (a) address existing research questions that have fundamental, practical utilization in many current applications, and (b) frame new research questions in important areas which need further scrutiny.

Keywords: Machine Learnibng (ML), Artificial Intelligence (AI), Power Line Communications (PLC), Secure Shell (SSH), Internet of Things (IoT), Field Programmable Gate Array (FPGA), Sparse Matrix, LU Decomposition.

I. INTRODUCTION

The convergence of Machine Learning and Artificial Intelligence (ML/AI) with modern information networks (such as 5G wireless telecommunications) is enabling a new class of always-connected, multi-dimensional, and previously impossible applications. The data available from distributed sensors, probes, and embedded computing devices makes practical applications from situations that have heretofore been impossible to imagine. From fleets of LED-equipped drones and choreographed "fireworks" to network analysis and optimization, to real-time attitude and position of first responders, applications derived from ubiquitous network connectivity and enabled by ML/AI are becoming incredibly important. Similarly, comparisons of data acquisition & formatting as well as training / testing processes are critical to application successes.

This special track focuses on the technologies, purposes, and architecture of such applications, with a particular interest in multi-disciplinary applications and practical relevance of the problems, enablers, comparisons, and outcomes.

II. SUBMISSIONS

A. *Big Data and Computation*

The first two papers in the track by Murthy and Aslan, "Optimization of Sparse Matrix Arithmetic Operations and Performance Improvement using FPGA" [1] and "Optimized Architecture for Sparse LU Decomposition on Matrices with Random Sparsity Patterns" [2] deal with the optimization of FPGA resources in the presence of sparse matrices.

Large datasets, which are popular in modern applications, and which are critical in ML/AI algorithms, are often stored in sparsely populated matrices. Additionally, large, sparse matrices often appear in scientific or engineering applications when solving partial differential equations. In such cases, low temporal locality of the data in the matrix leads to inefficiencies in access, storage, and manipulation.

Specialized algorithms and data structures are required to take advantage of the sparse structure of such matrices and improve the efficiency of the processing algorithms as well as optimize the allocation of FPGA resources. Algorithms are presented which effectively perform this complex, multidimensional optimization by minimizing critical FPGA resources (e.g. gate count, area, computational time, latency, etc.) while improving processing performance.

The designs presented are architected towards simple, scalable implementation with minimal input and output parameters. Furthermore, the architectures are generic and can be implemented irrespective of the application domain.

B. *IOT Architectural Considerations*

The second two papers in the track, "IoT Applications with Common Distributed Architecture for Data Acquisition" by Thapa, Lokesh, and Seets [3] and "Remote Filesystem Event Notification and Processing for Distributed Systems" by Lokesh, Thapa, and McClellan [5] explore commonalities in application architecture and efficient event notification techniques for implementing loosely-coupled, highly-distributed applications such are common in the "Internet of Things" (IoT).

In [3], the authors argue that multiple applications benefit from a common architecture denoted as "Coordinated IoT For Data Acquisition" (CIDAQ). The authors note that hundreds of IoT start-ups as well as several Global 500 companies offer applications, services, and tools which align with the CIDAQ architecture. Examples cited include technology-intensive training scenarios for first responders which involve motion capture, novel IoT-based monitoring of patients in convalescent facilities, very unique and interesting applications of tracking endangered species using heavily distributed embedded systems, and monitoring the power grid for anomalies as well as leveraging ultra-low-frequency communications techniques.

In [5], the authors note that monitoring events in a loosely coupled architecture can be difficult, but is often important in distributed applications. Additionally, they propose the use of a filesystem-based event monitoring system which is not available outside of localized, operating-system-specific options. For example, local filesystem events can be monitored by several conventional tools on multiple operating systems, such as inotify, Direvent, iWatch, Kqueue, FSEvents, etc.

In contrast, the paper presents a simple, scalable, and efficient technique using multiplexed Secure Shell (SSH) and redirected filesystem events. The approach, which is compatible with Internet-reachable and firewalled systems, enables highly secure, remote file system monitoring with minimum overhead. Metrics are provided which demonstrate the effectiveness of the multiplexed SSH technique.

C. *Machine Learning*

The final two papers in the special track, "Supervised Machine Learning in Digital Power Line Communications" by Thapa, McClellan, and Valles [4] and "An Evaluation of Neural Network Performance Using Complex-Valued Input Data" by Thapa and McClellan [6] describe evaluations of processing for ML algorithms using complex-valued data and outcomes of application-level ML algorithms in the demodulation of fundamental communication signals in power line communications.

In both [4] and [6], the authors argue that complex-valued data is common in many applications, including biomedical imaging, seismic sensing, signal processing, and communications. In particular, [4] describes the ultra-low-frequency power line communications environment as suffering from difficult out-of-band interference, a highly reactive channel, and complex spectral allocation issues which may benefit from appropriate ML treatments. Clearly, the use of complex-valued information streams is typical in communications systems, so the performance evaluation of ML algorithms using various formats of complex-valued data in [6] is of interest here, particularly since mainstream ML algorithms are optimized for real-valued input data.

More specifically, in [4] the authors compare the performance of various ML algorithms such as Support Vector Machine (SVM), Logistic Regression, Decision Tree, and Artificial Neural Networks (ANN) in demodulating amplitude-shifted, phase-shifted, and frequency-shifted data streams. They conclude that ML algorithms can efficiently extract information carried by power line communications signals while minimizing the impact of noise components which are difficult to model effectively. In particular, strong out-of-band interference can be compensated automatically via proper choice of ML algorithms and appropriate training, and complex harmonic structures present in the power line channel can be effectively minimized during the demodulation process if the correct ML algorithm structure is employed, and if the format of the (complex-valued) input data is selected to match application constraints.

In [6], different approaches to pre-processing complex-valued training data are compared and contrasted in conventional ML and neural network scenarios. These approaches include "stacking" real and imaginary components framewise, and using either the rectangular or polar forms of complex data sets as training inputs for conventional ML algorithms. The authors conclude that the "stacking" approach tends to perform better in many cases, but tradeoffs in efficiency, size of dataset, training time, and other application-specific variables hold key considerations.

III. Conclusion

The proliferation of large datasets in contemporary applications has led to a number of complex issues in computer and algorithm architecture. From speed and efficiency, to formatting and architectural structure of the application, engineers and scientists need to be aware of bottlenecks and complexity at every stage of the application.

In particular, the convergence of ML/AI with modern information networks is enabling a new class of always-connected, multi-dimensional, and previously impossible applications typically called the "Internet of Things" (IoT). The data available from these distributed sensors, and embedded computing devices, whether they are monitoring endangered species or making sense of communication signals, results in novel applications as well as poorly phrased or understood technical barriers. Pertinent comparisons of data acquisition & formatting, training / testing processes, and architectural tradeoffs are critical to application successes and implementation efficiency.

This special track focused on specific technologies, common architectures, misunderstood issues, and interesting outcomes of modern, data-intensive applications and algorithms. The practical relevance of the presented approaches and solutions may provide some utility to practitioners and researchers in related fields.

Following are specific observations and conclusions:

A. *Big Data and Computation*

- Large datasets are prevalent in many contemporary applications. Hardware-assist is a common approach to offloading general-purpose CPUs in such cases (e.g. graphics processors and toolkits have been used successfully in many data-intensive applications)

- The use of FPGA-driven offload engines may be appropriate for many datasets, and further research is needed to increase logic resources with a comparable increase in I/O bandwidth and on-chip memory, esp. for applications where sparse matrices are common.

B. IOT Architectural Considerations

- Many IoT-based applications share a common underlying system architecture which is essentially a heavily distributed, network-connected data acquisition system. Many unusual applications (e.g. endangered species monitoring) can benefit from such structures.
- Efficient monitoring of filesystem events in such loosely coupled, distributed systems may be a more effective "publish/subscribe" bus or "message bus" than often-used and relatively heavy REST-driven APIs, and may provide better security (e.g. via multiplexed SSH).

C. Machine Learning

- Complex-valued data is used in many applications, and is a critical component of many prevalent ones, and other unusual ones (e.g. low-frequency power line communications). Unfortunately, most important ML algorithms and Neural Network constructs are not compatible with complex-valued input data.
- Pre-processing complex-valued data in various ways (e.g. rectangular, polar, etc.) may produce different outcomes or efficiencies when using ML/AI algorithms which are not "aware" of complex numbers, or have not been constructed to handle complex-valued inputs

Acknowledgment

My sincere thanks to the organizers for considering this special track in the conference. Additionally, tremendous thanks to the reviewers who provided invaluable feedback on each paper. Finally, the authors of the submitted papers deserve special appreciation for the interesting submissions, the diversity of concepts contained in the papers, and their willingness to publish and present them.

References

1. D. K. Murthy and S. Aslan, "Optimization of Sparse Matrix Arithmetic Operations and Performance Improvement using FPGA," in Special Track: Applications of Machine Learning and Artificial Intelligence in Modern Networks, along with ICDT 2021, IARIA XPS Press, 2021.
2. D. K. Murthy and S. Aslan, "Optimized Architecture for Sparse LU Decomposition on Matrices with Random Sparsity Patterns," in Special Track: Applications of Machine Learning and Artificial Intelligence in Modern Networks, along with ICDT 2021, IARIA XPS Press, 2021.
3. K. Thapa, V. Lokesh, and K. Seets, "IoT Applications with Common Distributed Architecture," in Special Track: Applications of Machine Learning and Artificial Intelligence in Modern Networks, along with ICDT 2021, IARIA XPS Press, 2021.
4. K. Thapa, S. McClellan, and D. Valles, "Supervised Machine Learning in Digital Power Line Communications," in Special Track: Applications of Machine Learning and Artificial Intelligence in Modern Networks, along with ICDT 2021, IARIA XPS Press, 2021.
5. V. Lokesh, K. Thapa, and S. McClellan, "Remote Filesystem Event Notification and Processing for Distributed Systems," in Special Track: Applications of Machine Learning and Artificial Intelligence in Modern Networks, along with ICDT 2021, IARIA XPS Press, 2021.
6. K. Thapa and S. McClellan, "An Evaluation of Neural Network Performance Using Complex-Valued Input Data," in Special Track: Applications of Machine Learning and Artificial Intelligence in Modern Networks, along with ICDT 2021, IARIA XPS Press, 2021.