FaultLine

On-device Training and Inference for Power Signature Analysis and Fault Detection

ELECOMP Capstone Design Project 2021-2022

This project is a continuation of 2020-2021 Project FaultLine

Acumentrics is continuing their support of the Program for the 6th consecutive year. Links are available at the end of this proposal.

Sponsoring Company:

Acumentrics Inc.
10 Walpole Park South
Walpole, Massachusetts 02081
https://www.acumentrics.com/

Company Overview:

Acumentrics, Inc., headquartered in Walpole, Massachusetts, has been a trusted market leader in RUPS™ (rugged AC and DC uninterruptible power supplies) for harsh and combat environments as well as autonomous power and heavy-duty industrial applications, since 1994. Acumentrics products provide clean power conditioning and battery backup when reliability is mission critical. Acumentrics is a preferred supplier of US-made power electronics to many of the world's largest prime defense contractors.

The modern military relies on computers and other sophisticated electronic equipment and relies on Acumentrics' products to keep that equipment online in harsh environments. Electrical variance, surges, spikes, sags, and interruptions can cause communication breakdown and data loss, especially during the rigors of active duty. With new autonomous power systems, these products can range from rack-mounted units to carry-on luggage, backpacks, and even handheld devices.

Technical Director:

Brenden Smerbeck (URI CoE '17)
ELECOMP Capstone Graduate 2017
Software Engineer
bsmerbeck@acumentrics.com

Project Motivation:

Predictive maintenance is a growing field of interest across all industries. The goal of predictive maintenance (PdM) is to reduce the likelihood of catastrophic failure by detecting variances when compared to equipment's normal operation. Successful implementations of these solutions operate by analyzing data collected by an array of sensors attached to the device. However, these require modification to the product and are neither easily deployable nor maintainable.

All electronic devices require energy to operate – which can be derived to a voltage and current value. In these electronics, power consumption changes over time; consuming a varying amount of power depending on the intended action. By analyzing the consumption of power, a system can not only uniquely identify a specific device amongst others, but also learn how that device operates. Through learning this "power signature", a system could identify when a device is behaving abnormally and notify a user before such abnormalities become catastrophic.

Acumentrics' products have a commonality – clean, reliable, and rugged power. Regardless of condition, our products operate at the highest level. These systems are, however, not immune to failure. And in mission-critical operations, these failures are not tolerable. Additionally, our systems are deployed in secure environments which do not permit data collection for analysis. Therefore, most cases of failure analysis occur when a system is returned to the company; after a fault has occurred.

Year 1: 2019-2020 Project FaultLine

The goal of FaultLine in its first year was to establish baseline functionality. The students were successful in designing and implementing a Fault Detection Unit (FDU) - capable of analyzing the power signature of a connected appliance and detecting fault states in real time. The prototype consisted of a custom printed circuit board (PCB) that connected both to a single-phase appliance and a Raspberry Pi. To "detect" faulty devices, the team had to design a machine learning (ML) model capable of inferring what it means for an appliance to be functioning normally. Once trained, the ML model was loaded onto the Raspberry Pi where it could sense power data from an appliance and infer its state in real-time.

In addition to the single-phase FDU, the team also created a prototype polyphase FDU. As Acumentrics' products support both single and polyphase power environments, this was an absolute necessity.

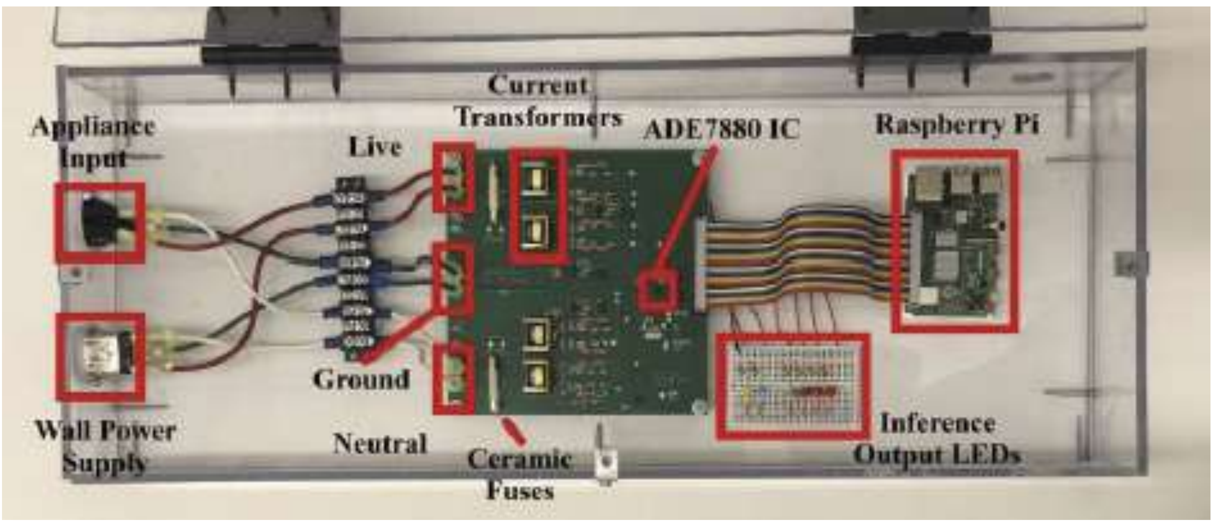

Year 2: 2020-2021

With a successful proof-of-concept developed, FaultLine's second year of development was focused on <u>integration into an existing product.</u> The Acumentrics SmartPDU is a commercially available product that was first designed by ELECOMP Capstone students during the 2017-2018 academic year. This year's team was responsible for taking the previous year's work, merging it with the SmartPDU, and then expanding on that original work – allowing the new fault detection system to run simultaneously on eight different outlets.

The Acumentrics SmartPDU

Year 3: 2021-2022

Machine learning is an incredible data science that has numerous advantages across countless industries. However, machine learning is limited by the amount of time and processing power required to train a ML model. The greater the complexity of the model or the more data required to process, the longer the training process will take. The motivation for this year's project is to push the limits of modern single-board computers (SBC) to determine if on-device training is possible and feasible. That is, training the ML model on an embedded device, rather than training on a powerful machine externally and transferring the model onto an embedded device. Fortunately, the SmartPDU was designed with this goal in mind.

Anticipated Best Outcome:

The Anticipated Best Outcome (ABO) is to successfully implement on-device model training using the embedded computing platform within the SmartPDU. The prototype system shall be able to infer the state of connected appliances and detect abnormalities based on variances in the power signature of the appliance (i.e., the FaultLine system).

Project Details:

Project FaultLine:

To better understand the technologies used in this project, as well as the work completed thus far, please visit the following links to previous project presentations and documentation. You can also visit the homepage for the projects using links at the end of the proposal to view each team's symposium presentation as was as their project posters.

- **2019-2020**
  - Original Project Proposal
  - Summit Presentation
- **2020-2021**
  - Original Project Proposal
  - Summit Presentation

On-device Training

On-device training is a process where machine learning models can be trained onboard an embedded system. These embedded systems are limited by storage size, operating power, and overall processing power. As such, training must be optimized to be technically feasible.

Traditionally, training of machine learning models would occur on more powerful systems. Once trained, these models could be loaded onto an embedded system where it would simply need to input data and view the results. This process is referred to as **model embedding**. While it allows for offline usage, the model is generally static and unchanging. A better method is **model serving**: in which the model is trained and hosted externally and accessed as it's needed. This model can also be stored in a data format that is readable programmatically; allowing it to be integrated into software that wishes to utilize it.

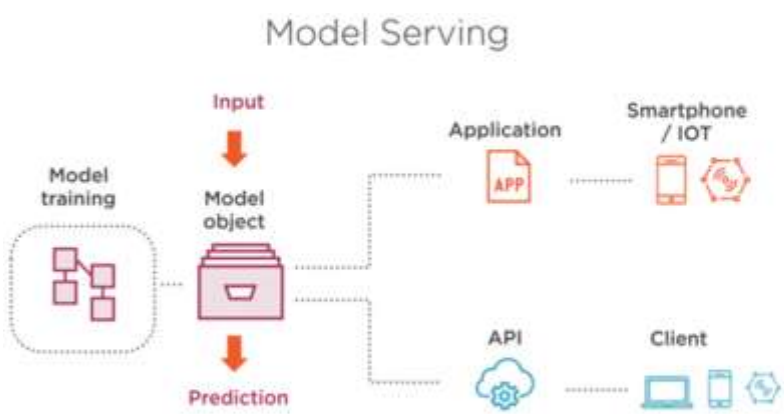

While both methods allow for a higher-complexity model to be trained and used, they have drawbacks. Embedded models are static and are not expected to be changed. Served models rely on some form of networked communication between the client and the delivery service/server. On-device training aims to allow embedded systems to train models locally. This type of training allows for data collectedly locally to directly influence the models structure. Additionally, these models are unique to the device, and don't require a network over which to communicate.

One such implementation, used by Apple and Google, is better defined as model "personalization". In this implementation, a generalized model is trained on a large data set externally. This model is then transferred to a device where "personalized training" can occur. This training is better thought of as "fine-tuning" – taking a general idea and creating a more detailed picture. This process is referred to as **transfer learning**.

Transfer learning describes a pattern in which a model is initially trained to create a "general" knowledge. This model can be trained on a significantly larger dataset external to the device. Once ***transferred*** to the target system, it is then "personalized" or "customized" by local data. One such example could be a smartphone keyboard recording taps and determining which key is pressed. While a generalized model would be capable of acceptable accuracy, using this model as the base "knowledge" and then learning based on an individual's screen taps would result in a far more accurate keyboard prediction. In fact, Apple uses machine learning with its own keyboard already, predicting your next word using autocompletion.

Computing Platform

The SmartPDU was designed with the intention of becoming a platform on which new technologies could be developed without substantial bottleneck. For that reason, the NVIDIA Jetson Nano was selected as the computing platform.

Jetson Nano Developer Kit (left), Jetson Nano Module (right)

The Jetson Nano is a small-form-factor single-board-computer touting 4GB of 64bit LPDDR memory, a quad-core ARM Cortex processor, 16GB eMMC 5.1 memory, and a GPU with 128 NVIDIA CUDA cores. NVIDIA designed the Jetson Nano to deliver the power of modern AI to the IoT ecosystem. We selected the Nano for that reason – to provide enough power to perform power signature analysis and fault detection within the SmartPDU.

Putting it All Together

With an adequate computer platform selected and previous year's projects success, what remains is an implementation of a transfer learning model architecture with respect to predictive maintenance. The SmartPDU would then be able to learn about the behavior of a connected electrical appliance using only the power of the Jetson Nano.

ML/AI Tasks:

- Familiarize selves with ML technologies and concepts including:
    - TensorFlow/TensorFlow Lite
    - Keras API
    - Long Term Short Memory (LTSM)
    - Recurrent Neural Networks (RNN)
- Familiarize selves with previous team's model and implementation
    - Model design
    - Data Collection
    - Data Processing
    - Training
    - Inference
- Familiarize selves with Jetson Nano platform and Jetson SDK
- Design a general model to serve as the knowledge base for transfer
- Design new model(s) for personalization/customization of connected devices

Computer Tasks:

- Familiarize selves with SmartPDU computing system, and underlying system code
    - Jetson Nano
    - IC communication
    - Data logging
    - Relay control
    - SmartPDU1 Software
- Perform data collection / training of general model
- Perform data collection / training of on-device model
- Demonstrate real-time inference of connected devices on the SmartPDU using transfer learning model

Composition of Team:

2-3 Computer Engineers (CPE)

Preference will be given to students who have previously studied in /worked with Machine Learning or those who have taken Google's Crash Course in Machine Learning (link). All students interested in the project should review the previous years' projects as they provide more detail in power signature analysis as it pertains to the project.

Skills Required:

Computer Engineering Skills Required:

- Proficiency in Object Oriented Programming (OOP)
- Experience in Python development
- Calculus, Linear Algebra (MTH 362, MTH 242 + MTH 243)
- Linear Systems (ELE 313)

Computer Engineering Skills Requested:

- Experience with TensorFlow
- Experience with TensorFlow Lite, Keras
- Experience with Recurrent Neural Networks
- Experience using a SoC/SBC (System on Chip, Single-board Computer) (e.g. Raspberry Pi, Beaglebone)

Anticipated Best Outcome's Impact on Company's Business, and Economic Impact

As Acumentrics' products are built to last in normally inoperable environments, integrity is an absolute requirement. As the company extends its knowledge of power systems to autonomous power, the need for data analytics and understanding only grows. To have a deeper understanding of electrical devices and their power signatures, Acumentrics can not only improve its own products but also improve the longevity of devices connected to those products.

Should the project succeed, there is a growing market for predictive maintenance across all industries – not just the military industry in which Acumentrics primarily exists. Therefore, the economic impact is too large to accurately measure. For existing customers, the project would allow Acumentrics to better understand its devices points of failure and continue to improve the ruggedness and longevity of those devices.

Broader Implications of the Best Outcome on the Company's Industry:

Power systems are a necessity often overlooked by the typical consumer. Yet power systems can be found behind every startup tech company and telecommunications firm, every first responder, every family, every military operation, and – frankly – every modern business. As technology advances, so too does the energy requirement of the world. In order to provide clean and reliable energy in every environment, companies within the industry of power systems must evolve as well. Today, machine learning and artificial intelligences have opened a gateway to medical innovations, seemingly impossible designs, and otherwise limitless opportunity. Yet these opportunities are equally matched by a demand for enhancement and improvement at every stage. It is because of this ever-growing demand that the innovations of the next generation are possible, and why projects like this are necessary.

Previous Projects

2020-2021: **FaultLine – SmartPDU1 Fault Detection Unit Integration**
2020-2021: N-Plus – Optimization of Parallel UPS Operation
2019-2020: **FaultLine – Power Signature Analysis for Fault Detection and Predictive Maintenance**
2019-2020: Volta – Automatic Variable Load Testing of High Voltage DC Output Boards
2018-2019: AcuBMS – Battery Management System for Rechargeable Lithium-based Batteries
2017-2018: AcuPDU – Network Managed Power Distribution Unit for Military Application
2016-2017: AESA – Acumentrics Easy Simple Network Management Protocol (SNMP) Application