SRF Cavity Fault Classification Using Machine Learning at Jefferson Lab

Chris Tennant | Jefferson Lab

December 3, 2021

FRIB/MSU Seminars Series

About the Talk

- this talk is not about
  - ✓ advancing the state-of-the-art in machine learning
  - ✓ implementing cutting edge algorithms
- this talk is about
  - ✓ a conceptually simple machine learning classification problem
  - ✓ the development of a deployed system – from data collection to implementation
  - ✓ valuable lessons learned along the way
  - ✓ the challenges of working with real-world data

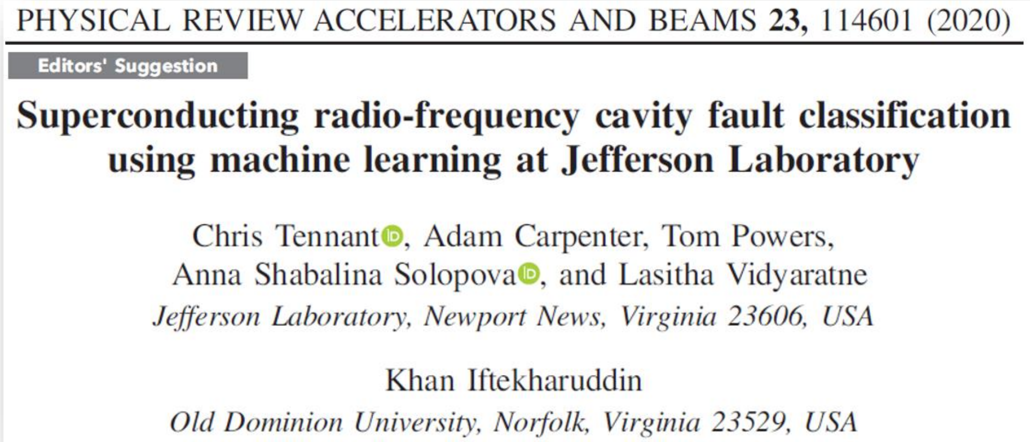

PHYSICAL REVIEW ACCELERATORS AND BEAMS **23**, 114601 (2020)

Editors' Suggestion

Superconducting radio-frequency cavity fault classification using machine learning at Jefferson Laboratory

Chris Tennant, Adam Carpenter, Tom Powers,
Anna Shabalina Solopova, and Lasitha Vidyaratne
Jefferson Laboratory, Newport News, Virginia 23606, USA

Khan Iftekharuddin
Old Dominion University, Norfolk, Virginia 23529, USA

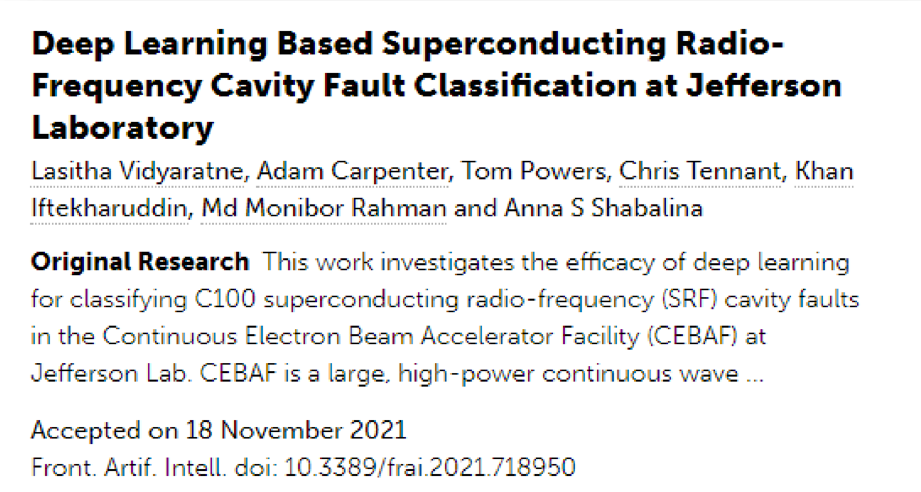

Deep Learning Based Superconducting Radio-Frequency Cavity Fault Classification at Jefferson Laboratory

Lasitha Vidyaratne, Adam Carpenter, Tom Powers, Chris Tennant, Khan Iftekharuddin, Md Monibor Rahman and Anna S Shabalina

Original Research This work investigates the efficacy of deep learning for classifying C100 superconducting radio-frequency (SRF) cavity faults in the Continuous Electron Beam Accelerator Facility (CEBAF) at Jefferson Lab. CEBAF is a large, high-power continuous wave ...

Accepted on 18 November 2021
Front. Artif. Intell. doi: 10.3389/frai.2021.718950

Acknowledgements

Adam Carpenter
Tom Powers
Monibor Rahman
Lasitha Vidyaratne
and many others!

Outline

- ***ML Landscape***
- ***Definitions***
- ***Fault Isolation and Identification***
  - ✓ Supervised Learning: Machine Learning
  - ✓ Supervised Learning: Deep Learning
  - ✓ Unsupervised Learning
- ***Fault Prediction***
- ***Other Work***
- ***Summary***

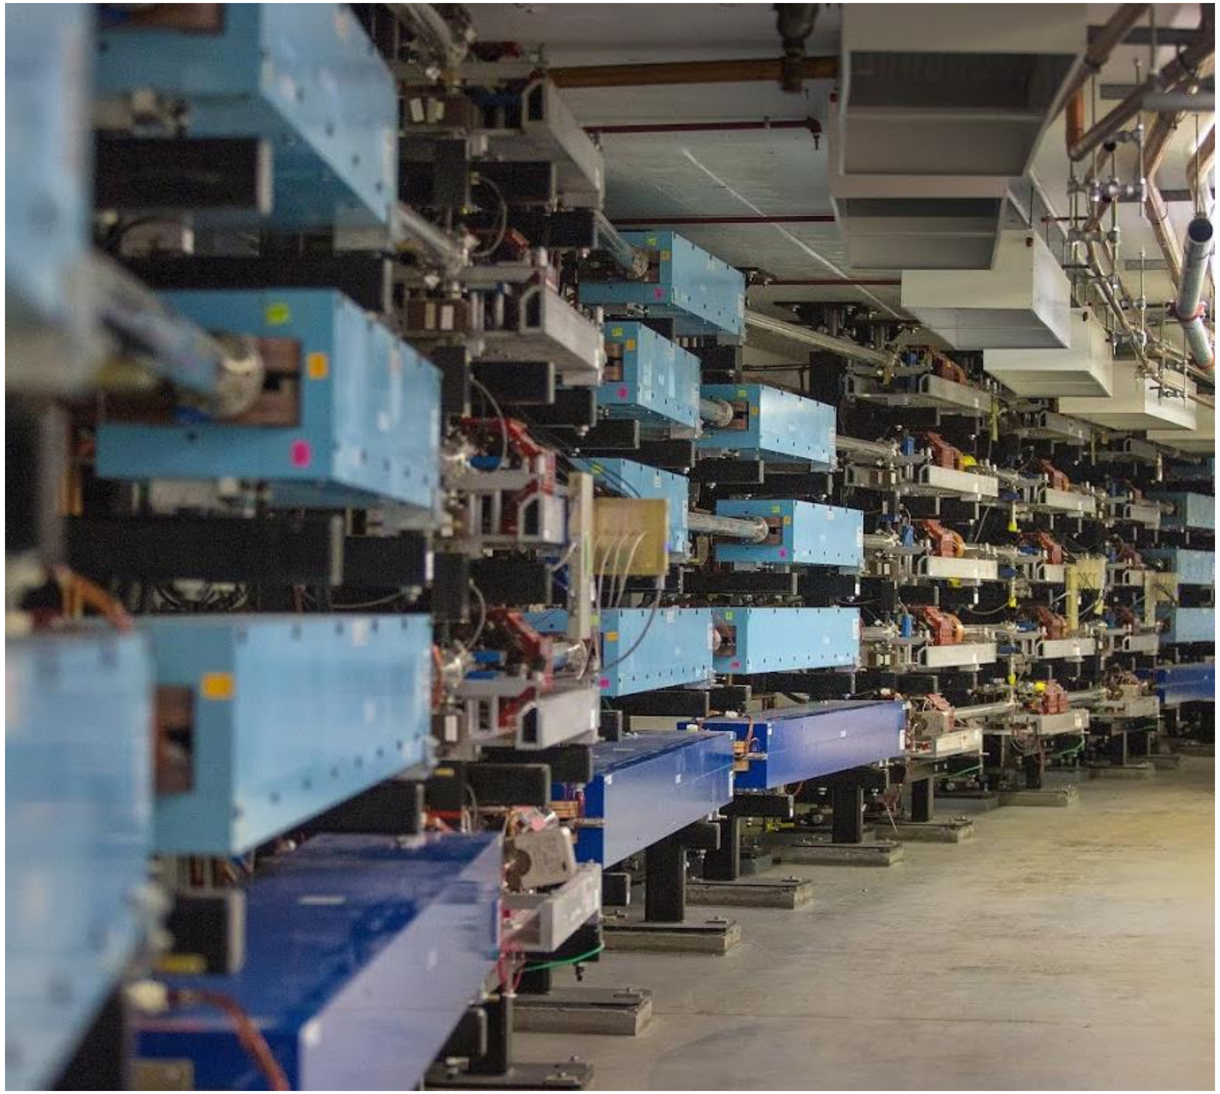

AI/ML in Accelerator Physics

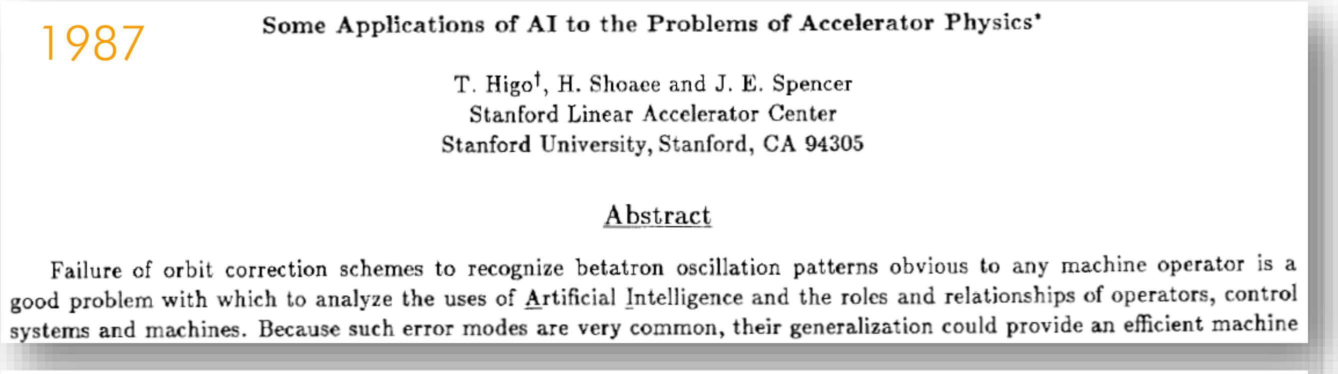

1987

Some Applications of AI to the Problems of Accelerator Physics*

T. Higo[†], H. Shoaee and J. E. Spencer
Stanford Linear Accelerator Center
Stanford University, Stanford, CA 94305

Abstract

Failure of orbit correction schemes to recognize betatron oscillation patterns obvious to any machine operator is a good problem with which to analyze the uses of Artificial Intelligence and the roles and relationships of operators, control systems and machines. Because such error modes are very common, their generalization could provide an efficient machine

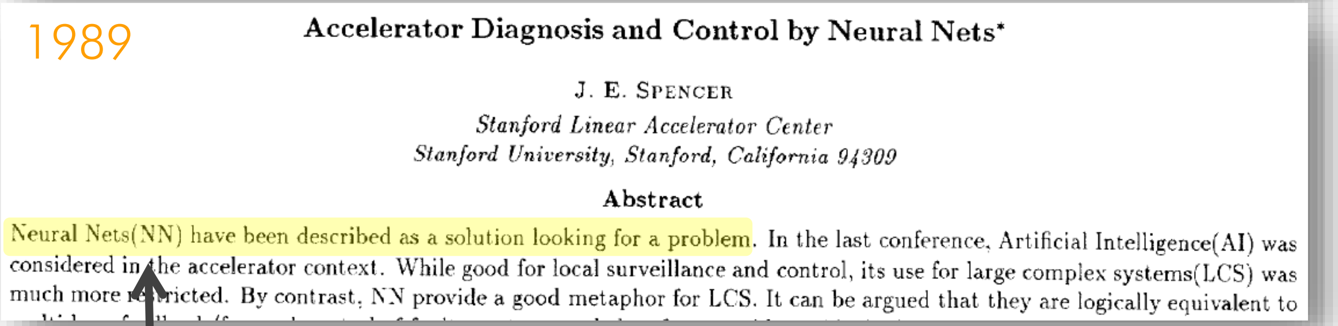

1989

Accelerator Diagnosis and Control by Neural Nets*

J. E. Spencer

Stanford Linear Accelerator Center
Stanford University, Stanford, California 94309

Abstract

Neural Nets(NN) have been described as a solution looking for a problem. In the last conference, Artificial Intelligence(AI) was considered in the accelerator context. While good for local surveillance and control, its use for large complex systems(LCS) was much more restricted. By contrast, NN provide a good metaphor for LCS. It can be argued that they are logically equivalent to

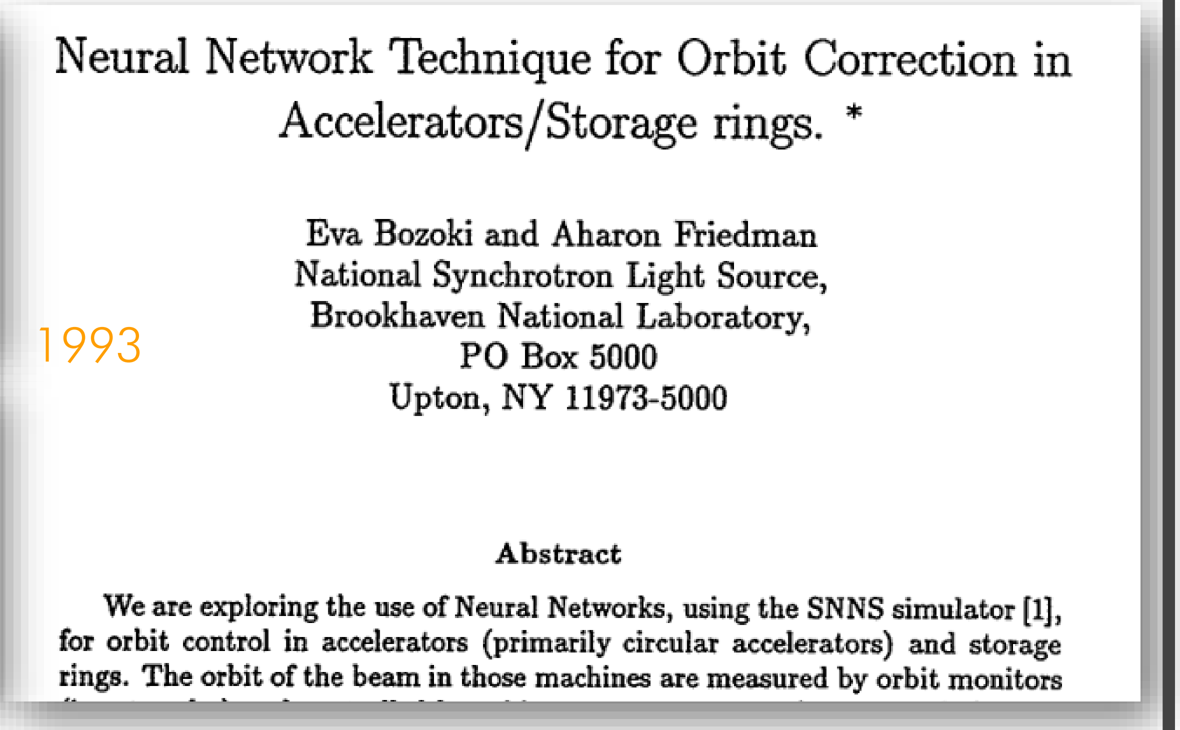

Neural Network Technique for Orbit Correction in Accelerators/Storage rings. *

Eva Bozoki and Aharon Friedman
National Synchrotron Light Source,
Brookhaven National Laboratory,
PO Box 5000
Upton, NY 11973-5000

1993

Abstract

We are exploring the use of Neural Networks, using the SNNS simulator [1], for orbit control in accelerators (primarily circular accelerators) and storage rings. The orbit of the beam in those machines are measured by orbit monitors

"Neural Nets (NN) have been described as a solution looking for a problem."

Rise of Peer-Reviewed ML Publications

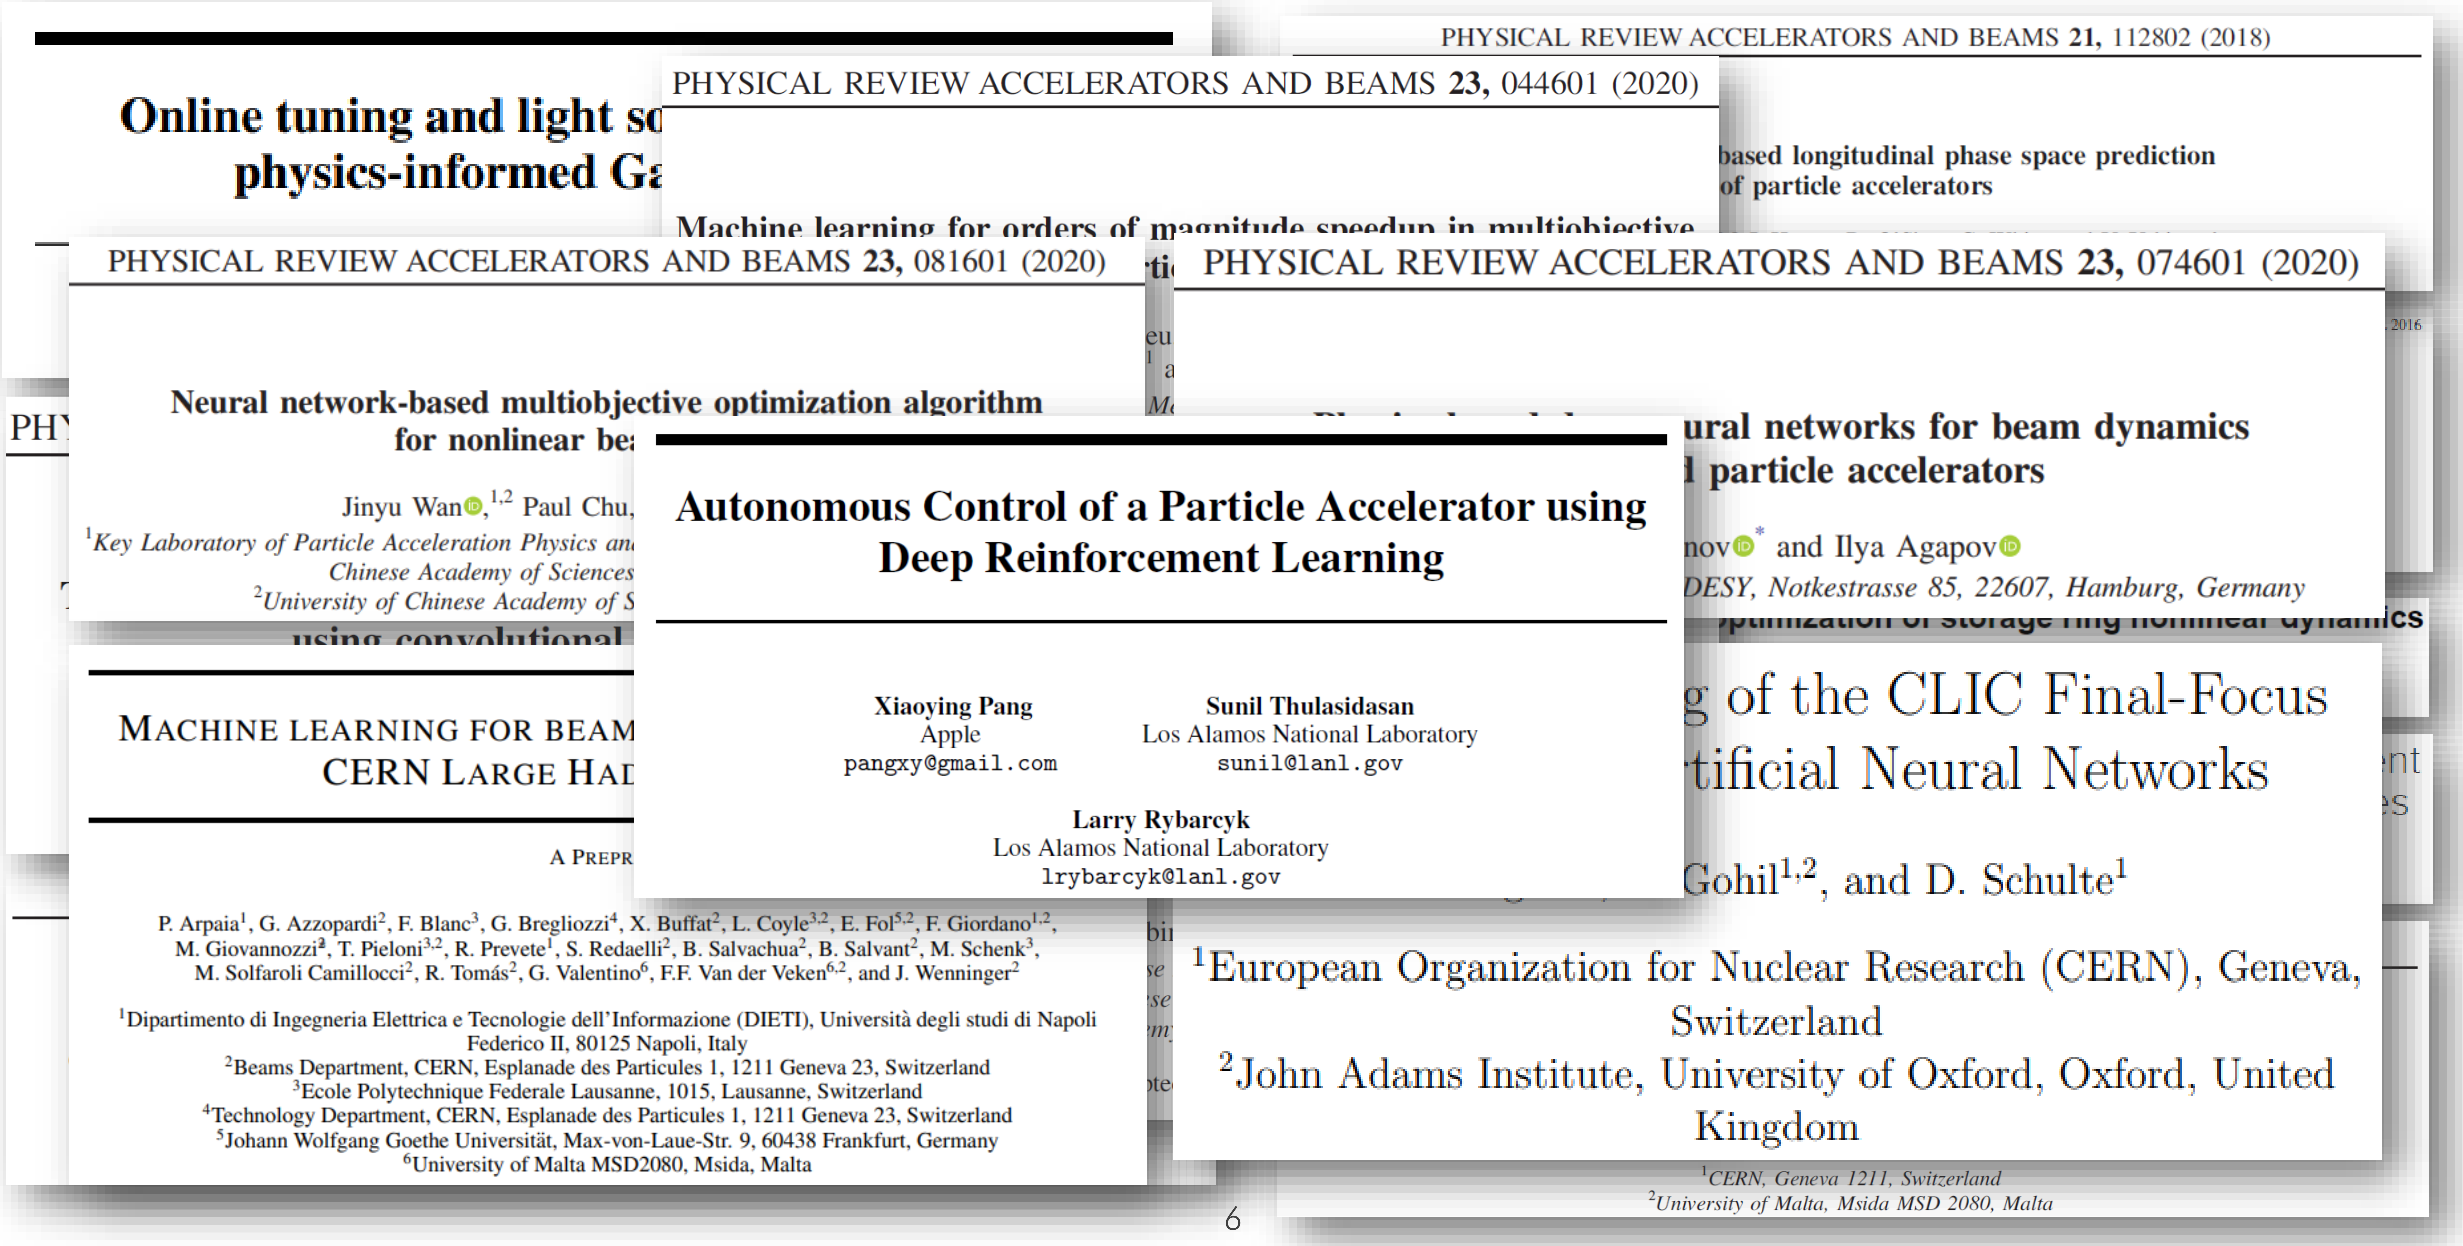

Online tuning and light so
physics-informed Ga

PHYSICAL REVIEW ACCELERATORS AND BEAMS **21,** 112802 (2018)

based longitudinal phase space prediction
of particle accelerators

PHYSICAL REVIEW ACCELERATORS AND BEAMS **23,** 044601 (2020)

Machine learning for orders of magnitude speedup in multiobjective

PHYSICAL REVIEW ACCELERATORS AND BEAMS **23,** 081601 (2020)

Neural network-based multiobjective optimization algorithm
for nonlinear bea

Jinyu Wan[1,2] Paul Chu,
[1]*Key Laboratory of Particle Acceleration Physics an*
Chinese Academy of Sciences
[2]*University of Chinese Academy of S*

using convolutional

PHYSICAL REVIEW ACCELERATORS AND BEAMS **23,** 074601 (2020)

ural networks for beam dynamics
l particle accelerators

nov* and Ilya Agapov
DESY, Notkestrasse 85, 22607, Hamburg, Germany

optimization of storage ring nonlinear dynamics

Autonomous Control of a Particle Accelerator using Deep Reinforcement Learning

Xiaoying Pang
Apple
`pangxy@gmail.com`

Sunil Thulasidasan
Los Alamos National Laboratory
`sunil@lanl.gov`

Larry Rybarcyk
Los Alamos National Laboratory
`lrybarcyk@lanl.gov`

MACHINE LEARNING FOR BEAM
CERN LARGE HAD

A PREPR

P. Arpaia[1], G. Azzopardi[2], F. Blanc[3], G. Bregliozzi[4], X. Buffat[2], L. Coyle[3,2], E. Fol[5,2], F. Giordano[1,2], M. Giovannozzi[2], T. Pieloni[3,2], R. Prevete[1], S. Redaelli[2], B. Salvachua[2], B. Salvant[2], M. Schenk[3], M. Solfaroli Camillocci[2], R. Tomás[2], G. Valentino[6], F.F. Van der Veken[6,2], and J. Wenninger[2]

[1]Dipartimento di Ingegneria Elettrica e Tecnologie dell'Informazione (DIETI), Università degli studi di Napoli Federico II, 80125 Napoli, Italy
[2]Beams Department, CERN, Esplanade des Particules 1, 1211 Geneva 23, Switzerland
[3]Ecole Polytechnique Federale Lausanne, 1015, Lausanne, Switzerland
[4]Technology Department, CERN, Esplanade des Particules 1, 1211 Geneva 23, Switzerland
[5]Johann Wolfgang Goethe Universität, Max-von-Laue-Str. 9, 60438 Frankfurt, Germany
[6]University of Malta MSD2080, Msida, Malta

g of the CLIC Final-Focus
rtificial Neural Networks

Gohil[1,2], and D. Schulte[1]

[1]European Organization for Nuclear Research (CERN), Geneva, Switzerland
[2]John Adams Institute, University of Oxford, Oxford, United Kingdom

[1]*CERN, Geneva 1211, Switzerland*
[2]*University of Malta, Msida MSD 2080, Malta*

Machine Learning 101

machine learning: *"The field of study that gives computers the ability to learn without being explicitly programmed"* A. Samuel

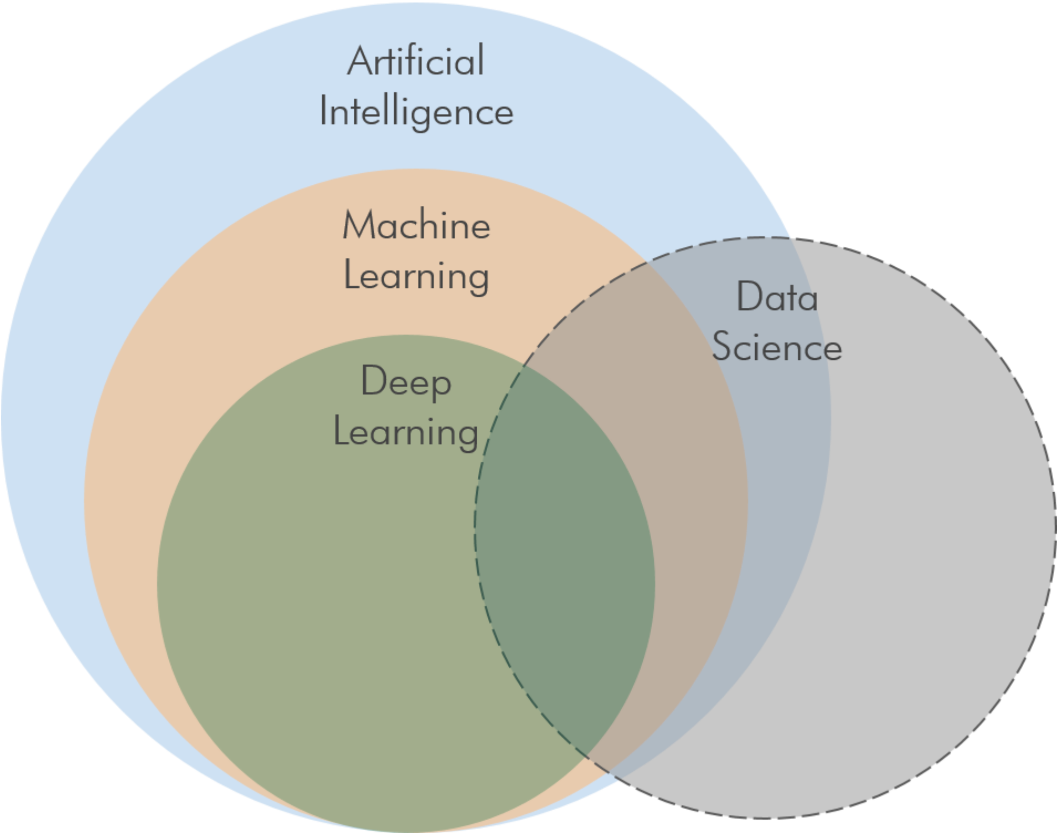

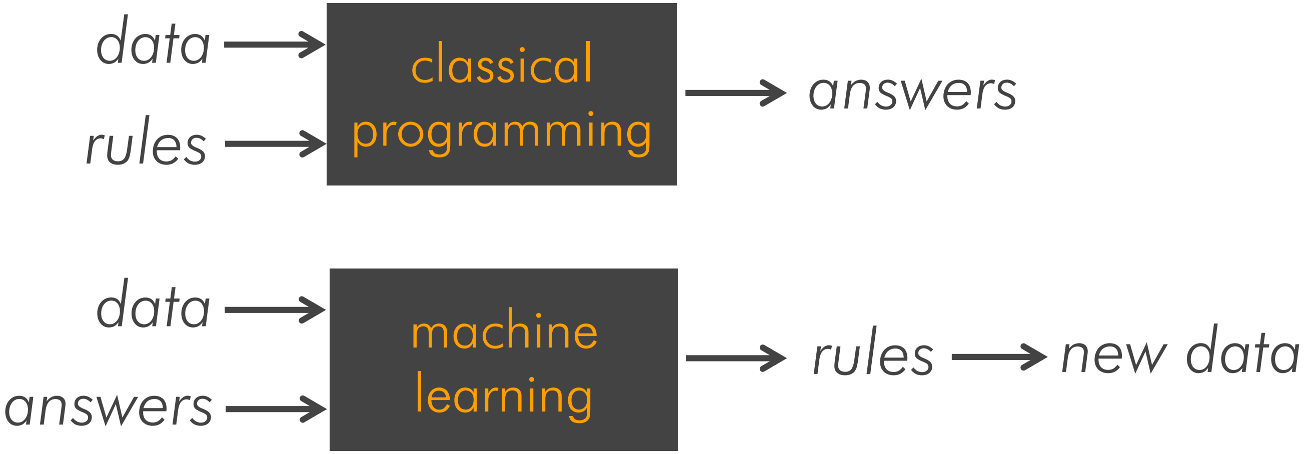

data → **classical programming** ← *rules*
→ *answers*

data → **machine learning** ← *answers*
→ *rules* → *new data*

"a machine-learning system is trained rather than explicitly programmed" F. Chollet

ML and Particle Accelerators

- particle accelerators represent the most complex scientific instruments designed, built, and operated
- there is clear motivation to maximize scientific output per operating dollar

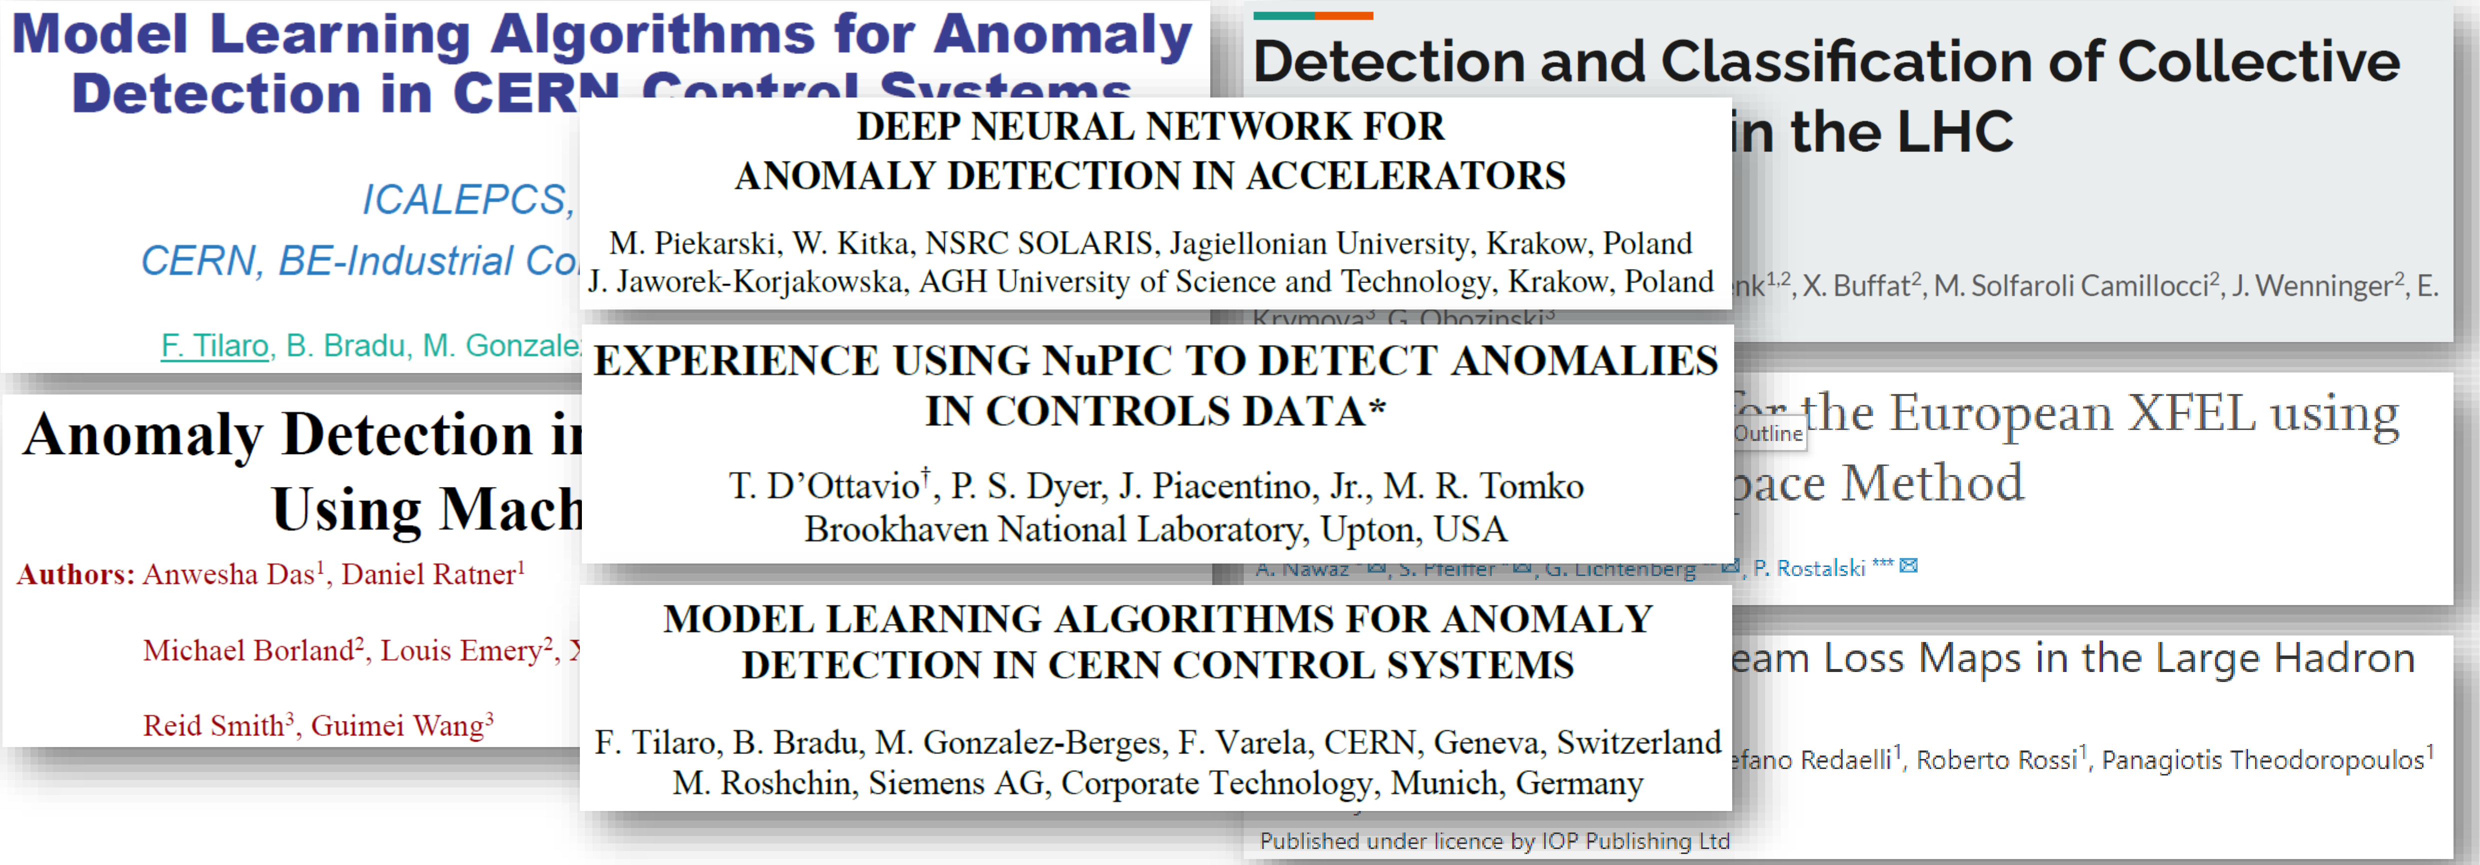

Model Learning Algorithms for Anomaly Detection in CERN Control Systems

ICALEPCS,

CERN, BE-Industrial Co

F. Tilaro, B. Bradu, M. Gonzale

Anomaly Detection i

Using Mach

Authors: Anwesha Das[1], Daniel Ratner[1]

Michael Borland[2], Louis Emery[2],

Reid Smith[3], Guimei Wang[3]

DEEP NEURAL NETWORK FOR ANOMALY DETECTION IN ACCELERATORS

M. Piekarski, W. Kitka, NSRC SOLARIS, Jagiellonian University, Krakow, Poland
J. Jaworek-Korjakowska, AGH University of Science and Technology, Krakow, Poland

EXPERIENCE USING NuPIC TO DETECT ANOMALIES IN CONTROLS DATA*

T. D'Ottavio†, P. S. Dyer, J. Piacentino, Jr., M. R. Tomko
Brookhaven National Laboratory, Upton, USA

MODEL LEARNING ALGORITHMS FOR ANOMALY DETECTION IN CERN CONTROL SYSTEMS

F. Tilaro, B. Bradu, M. Gonzalez-Berges, F. Varela, CERN, Geneva, Switzerland
M. Roshchin, Siemens AG, Corporate Technology, Munich, Germany

Detection and Classification of Collective

in the LHC

nk[1,2], X. Buffat[2], M. Solfaroli Camillocci[2], J. Wenninger[2], E. Krymova[3], G. Obozinski[3]

the European XFEL using

pace Method

Outline

A. Nawaz, S. Pfeiffer, G. Lichtenberg, P. Rostalski ***

eam Loss Maps in the Large Hadron

fano Redaelli[1], Roberto Rossi[1], Panagiotis Theodoropoulos[1]

Published under licence by IOP Publishing Ltd

Definitions

- ***Fault***: an unpermitted deviation of at least one characteristic property or parameter of the system from acceptable, usual or standard conditions

- ***Fault Detection***: monitoring measured variables to determine if a fault has occurred (if a fault has occurred, it may be important to determine the time at which the fault occurred)

- ***Fault Isolation***: determining the location of a fault once it is known that a fault has occurred

- ***Fault Identification***: determining the type of fault

- ***Fault Prediction***: providing advanced warning of an impeding fault

(R. Isermann, *"Fault Diagnosis Systems"*)

Detection vs (Isolation, Identification, Prediction)

- machine protection systems, personal safety systems, alarms, and other engineered systems are able to detect many types of faults
- in these instances, detection is not necessary → it’s (painfully) obvious when a fault has occurred

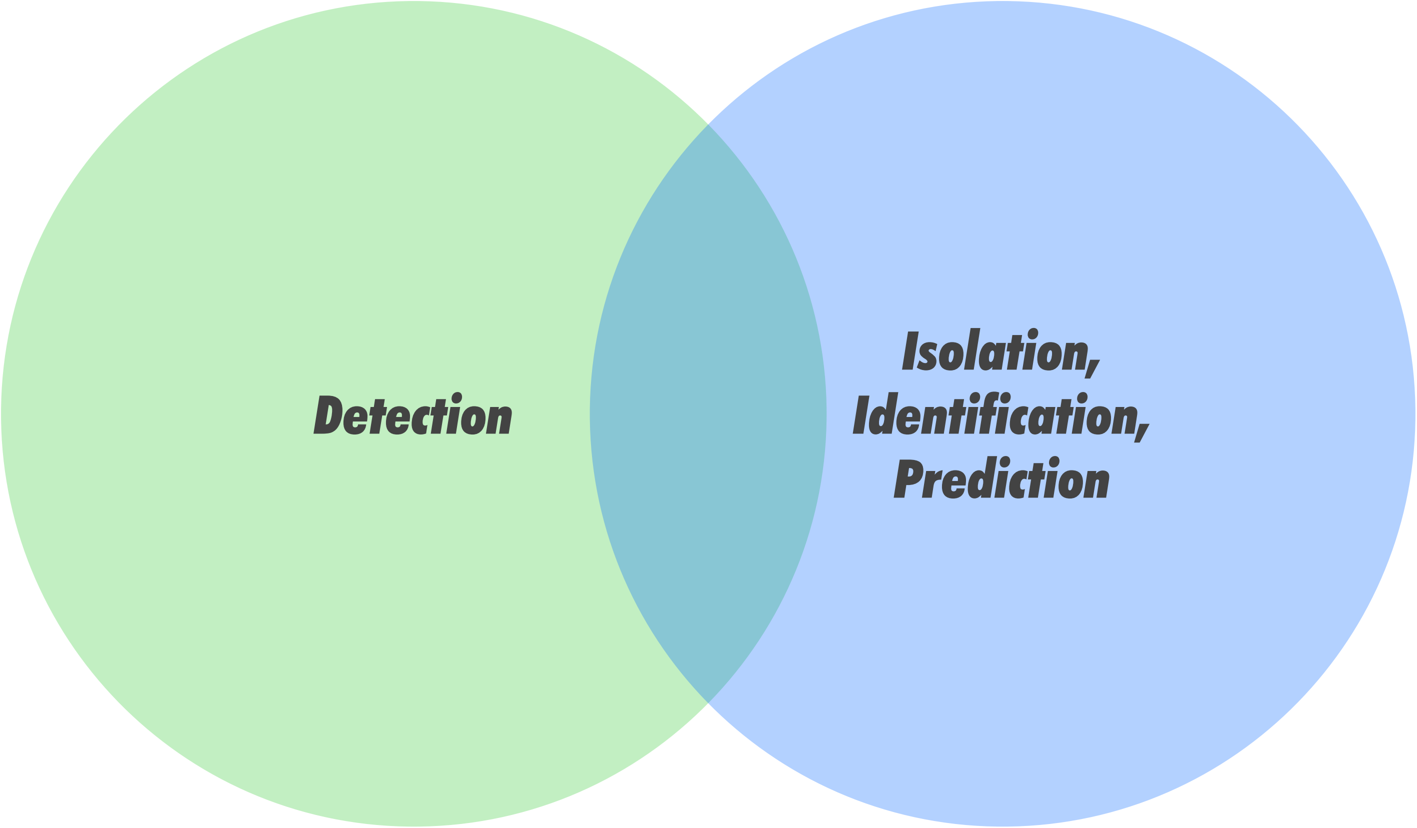

Continuous Electron Beam Accelerator Facility

- CEBAF is a CW recirculating linac utilizing 418 SRF cavities to accelerate electrons up to 12 GeV through 5-passes
- it is a nuclear physics user-facility capable of servicing 4 experimental halls simultaneously
- the heart of the machine is the SRF cavities

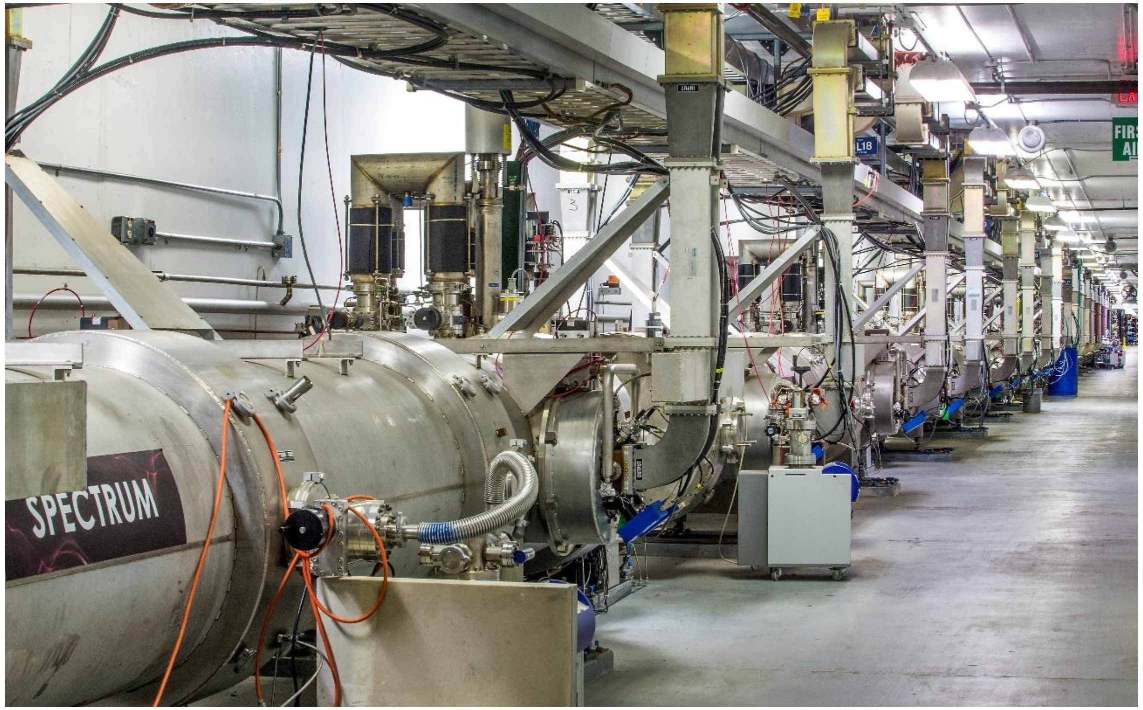

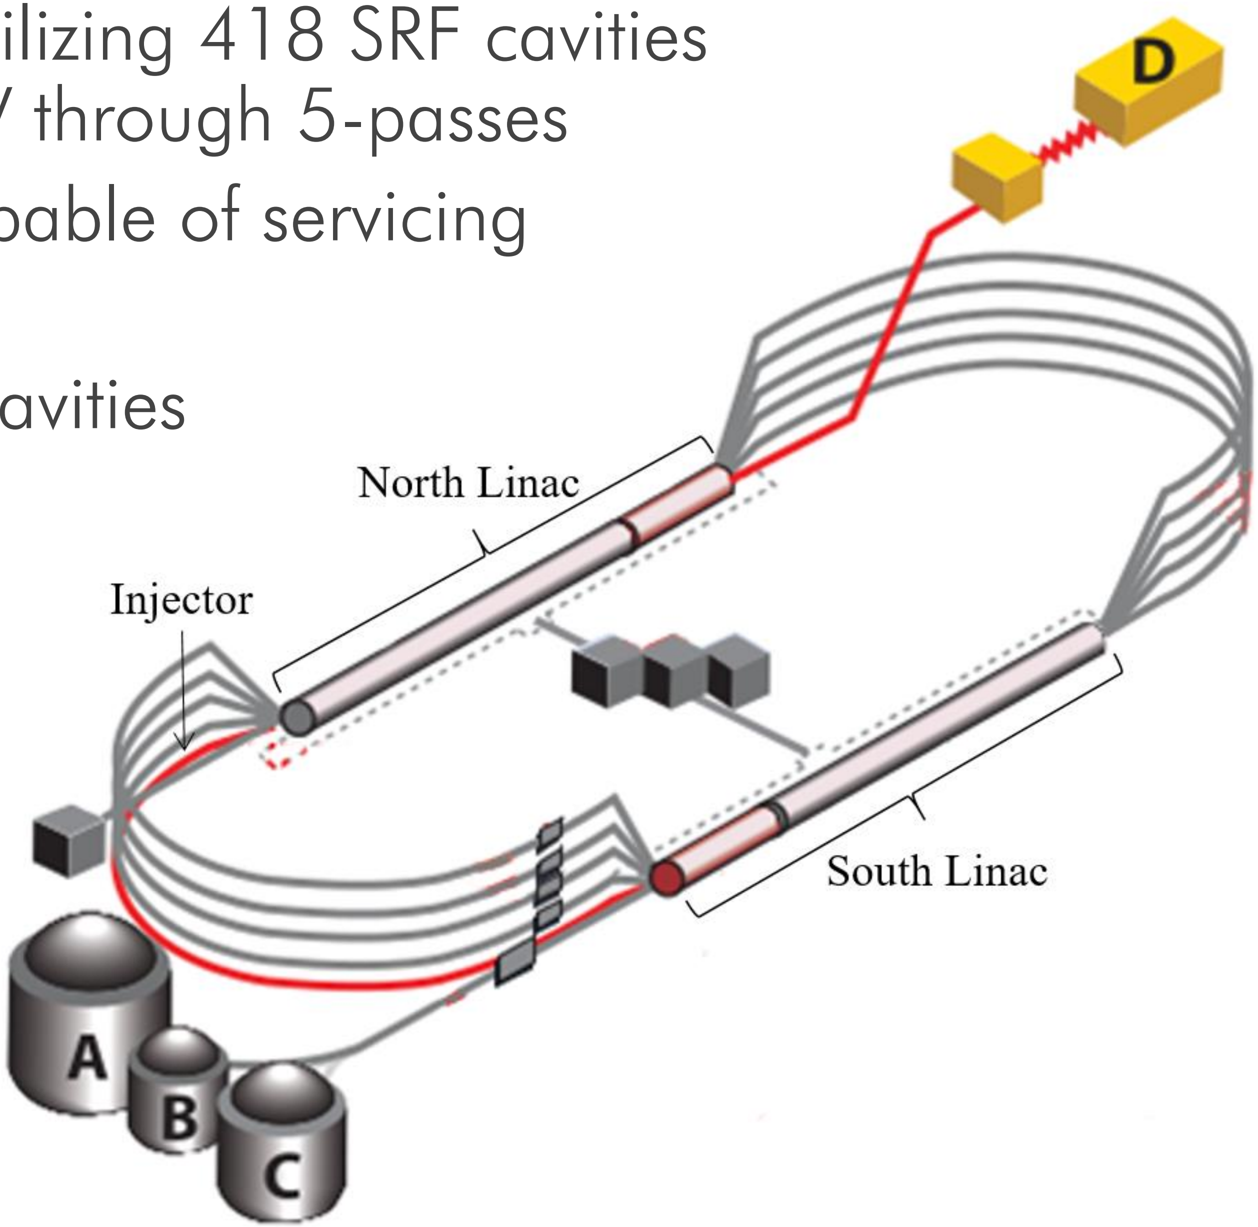

CEBAF Down Time Manager

- CEBAF short machine downtime trips (< 5 min.) in 2019

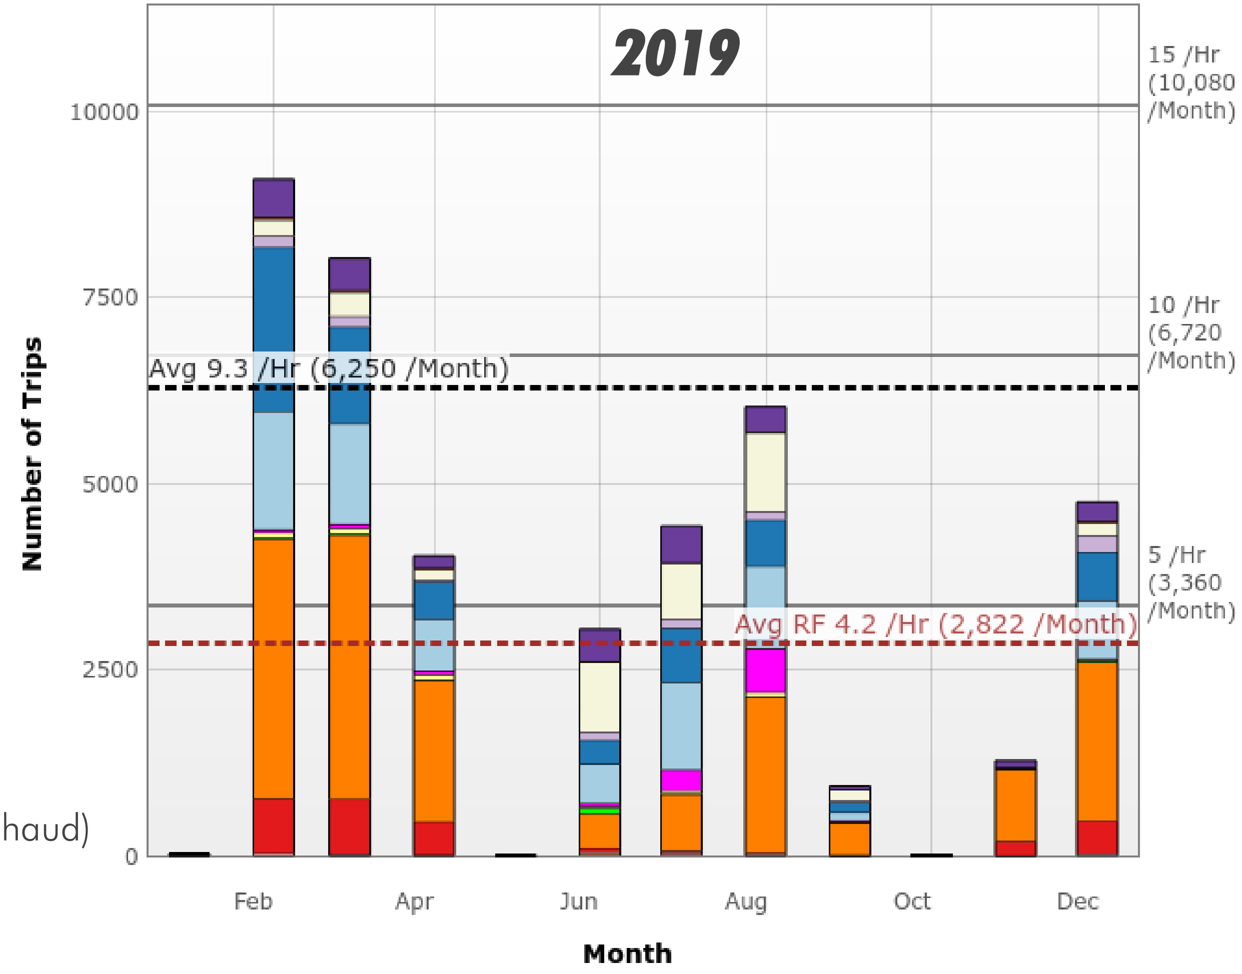

(courtesy R. Michaud)

		Trips /Hr	Lost Hrs	Mins /Trip
🟪	Dump (Insert.)	0.6	12.7	0.3
⬜	Dump (Station.)	0.0	0.0	0.3
🟫	Gun/Laser	0.0	1.5	0.5
⬜	Hall	0.8	17.7	0.3
🟪	MPS (BCM/BLA)	0.2	4.9	0.3
🟦	MPS (BLM)	1.4	35.9	0.3
🟦	MPS (Multi/Other)	1.6	39.7	0.3
🟥	Magnets	0.0	0.1	1.2
🟪	Multiple/Other	0.2	5.6	0.3
🟨	Unknown/Missing	0.1	2.3	0.4
🟩	Vacuum	0.0	1.9	0.8
🟧	RF (C25/C50)	3.5	109.2	0.4
🟥	RF (C100)	0.6	75.2	1.7
🟥	RF (Multi/Other)	0.0	3.0	1.0
🟧	RF (Separator)	0.0	1.1	1.3
	Total:	**9.3**	**310.7**	**0.4**

- Max Trip Duration: 5 Minutes
- Rate from Program (4492.47 hrs)
- SAD Trips excluded

Defining the Problem

we have the ability to record high-fidelity data from 12 cryomodules

1 cryomodule = collection of 8 cavities

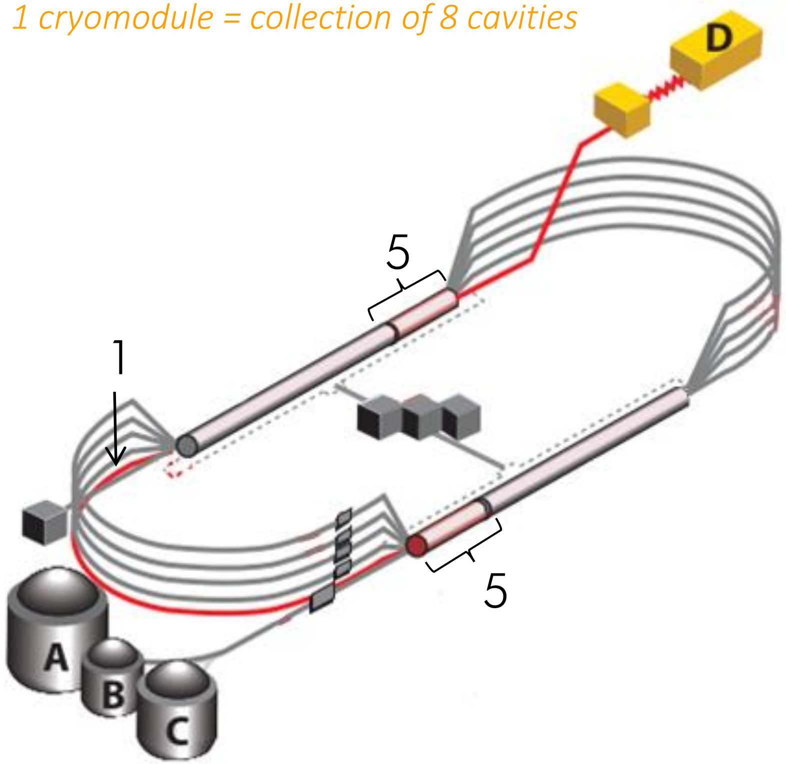

FAULT ISOLATION

Which of the 8 cavities faulted first?

17 signals/cavity × 8 cavities = 136 signals

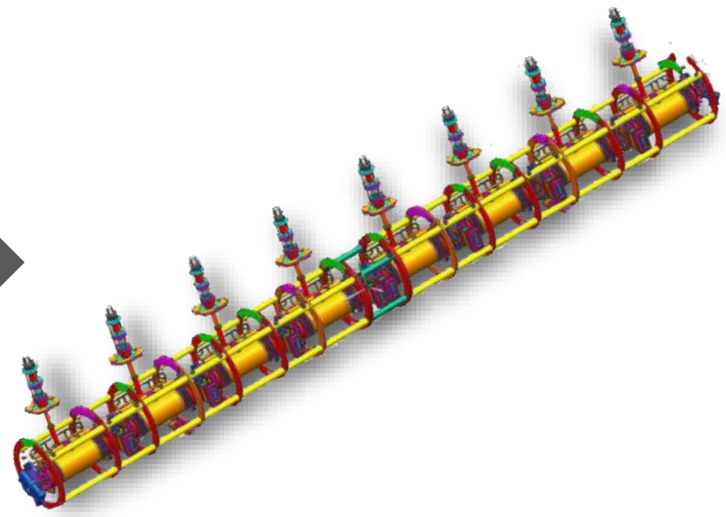

FAULT IDENTIFICATION

What kind of trip was it?

17 signals

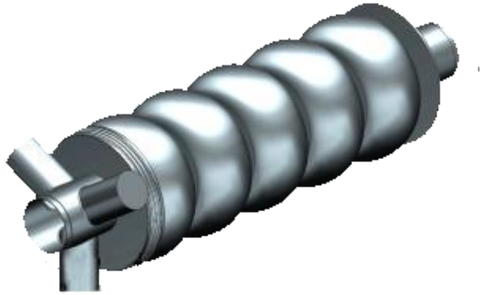

train a model to correctly classify the <u>cavity</u> and <u>type</u> of RF fault given waveform data

machine learning

multi-class classification

time-series data

Brute Force Data Analysis

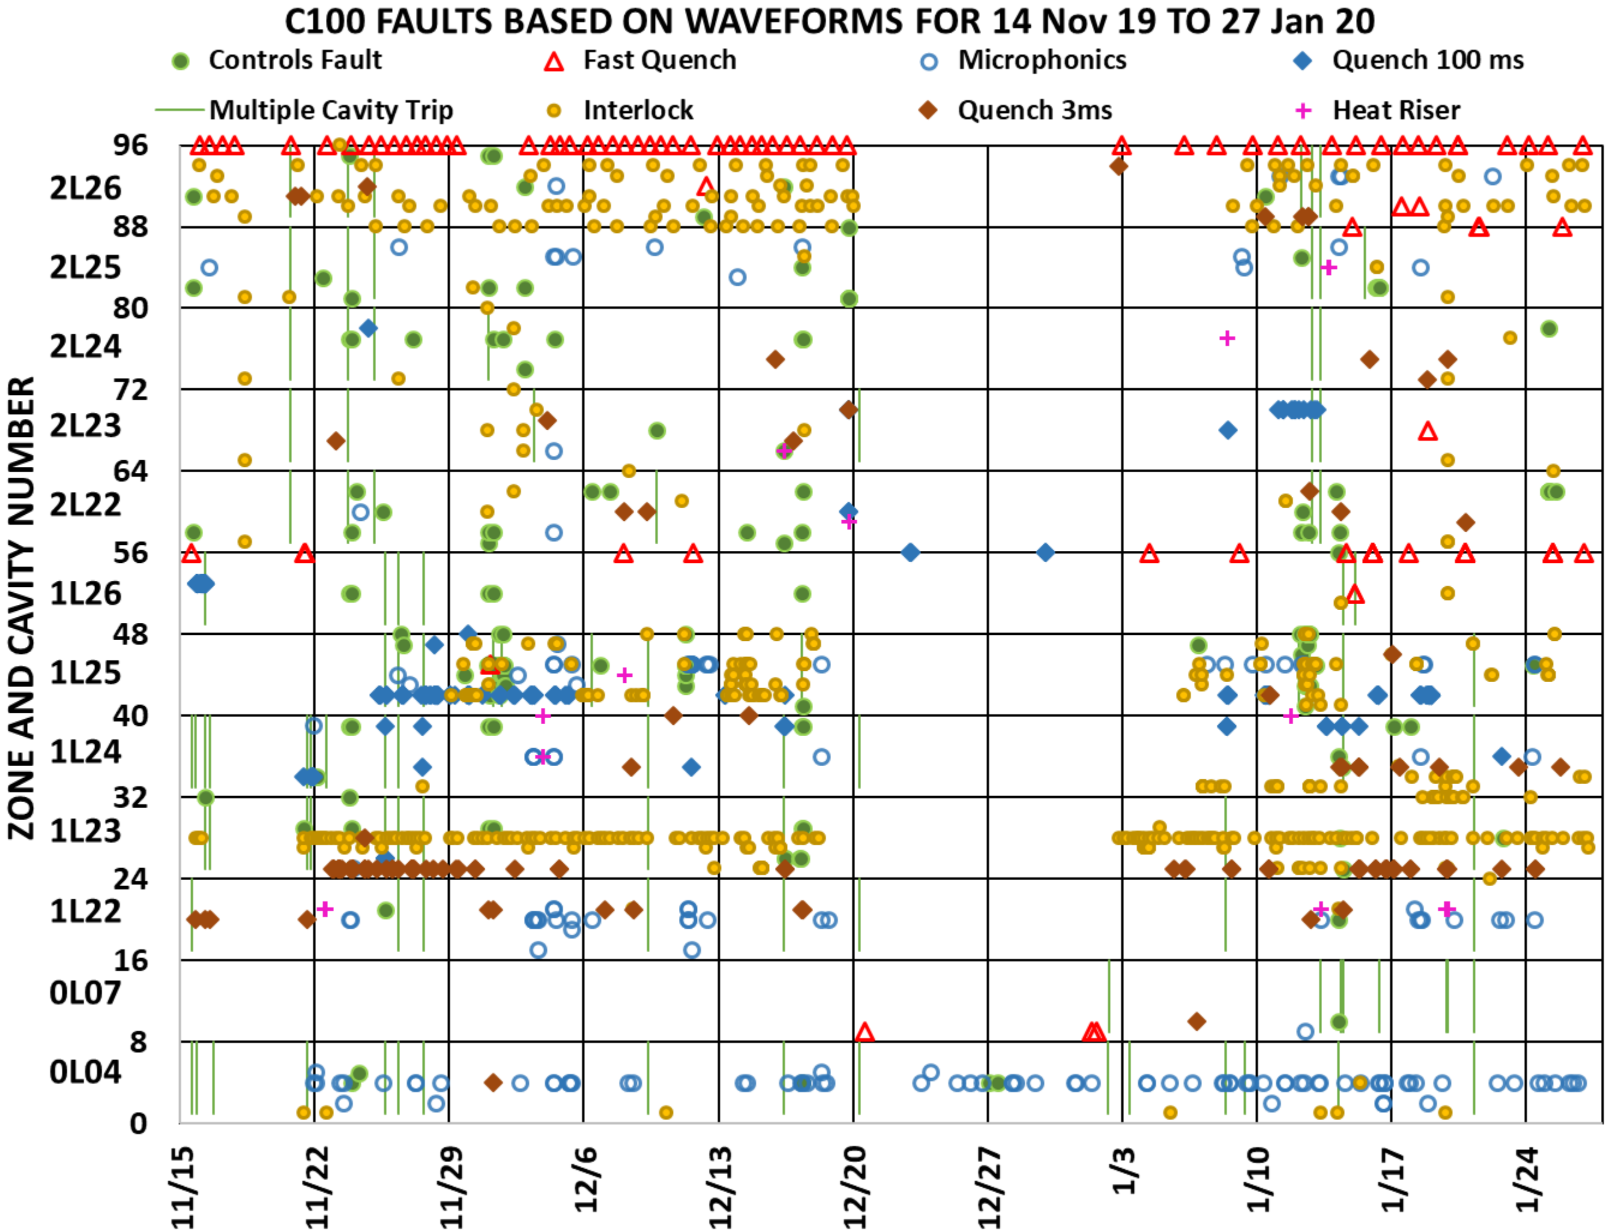

(courtesy T. Powers)

Jefferson Lab

Data Acquisition System

- waveform harvester was developed to capture RF time-series signals after a fault and write them to file for later analysis
  - ✓ each of the 17 harvested waveform signals is 8,192 points long
  - ✓ trigger set such that 94% of the recorded data precedes the fault and 6% after
  - ✓ pre-fault data provides valuable information about the root cause of the trip

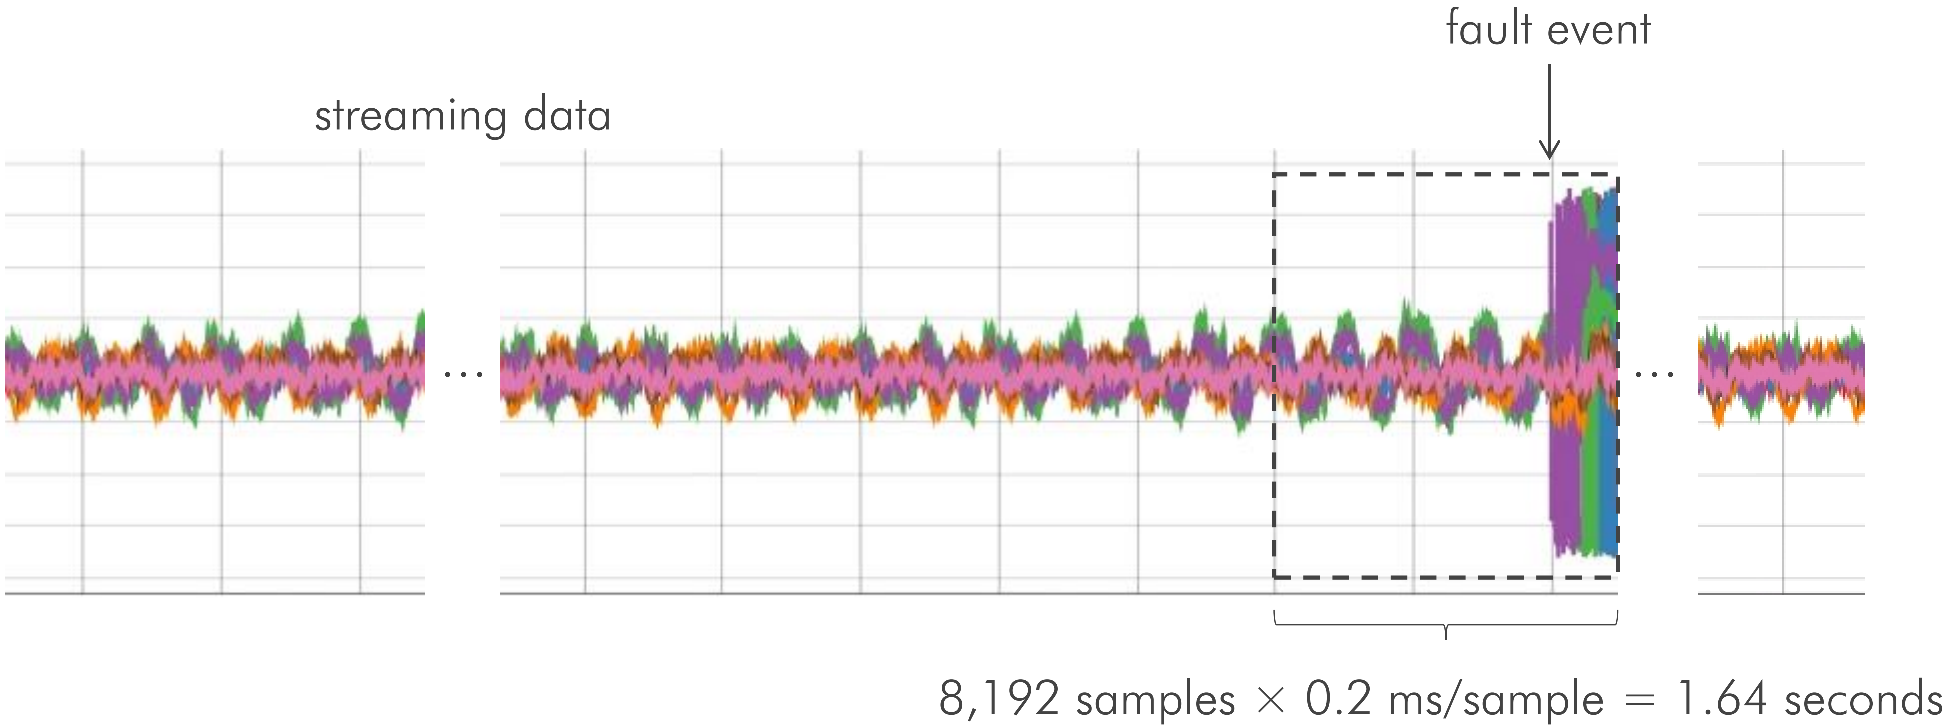

8,192 samples × 0.2 ms/sample = 1.64 seconds

Motivation

- labeling is hard
  - ✓ have a subject matter expert with 30+ years SRF experience to label fault events
  - ✓ closer to annotating medical images than distinguishing between cats and dogs

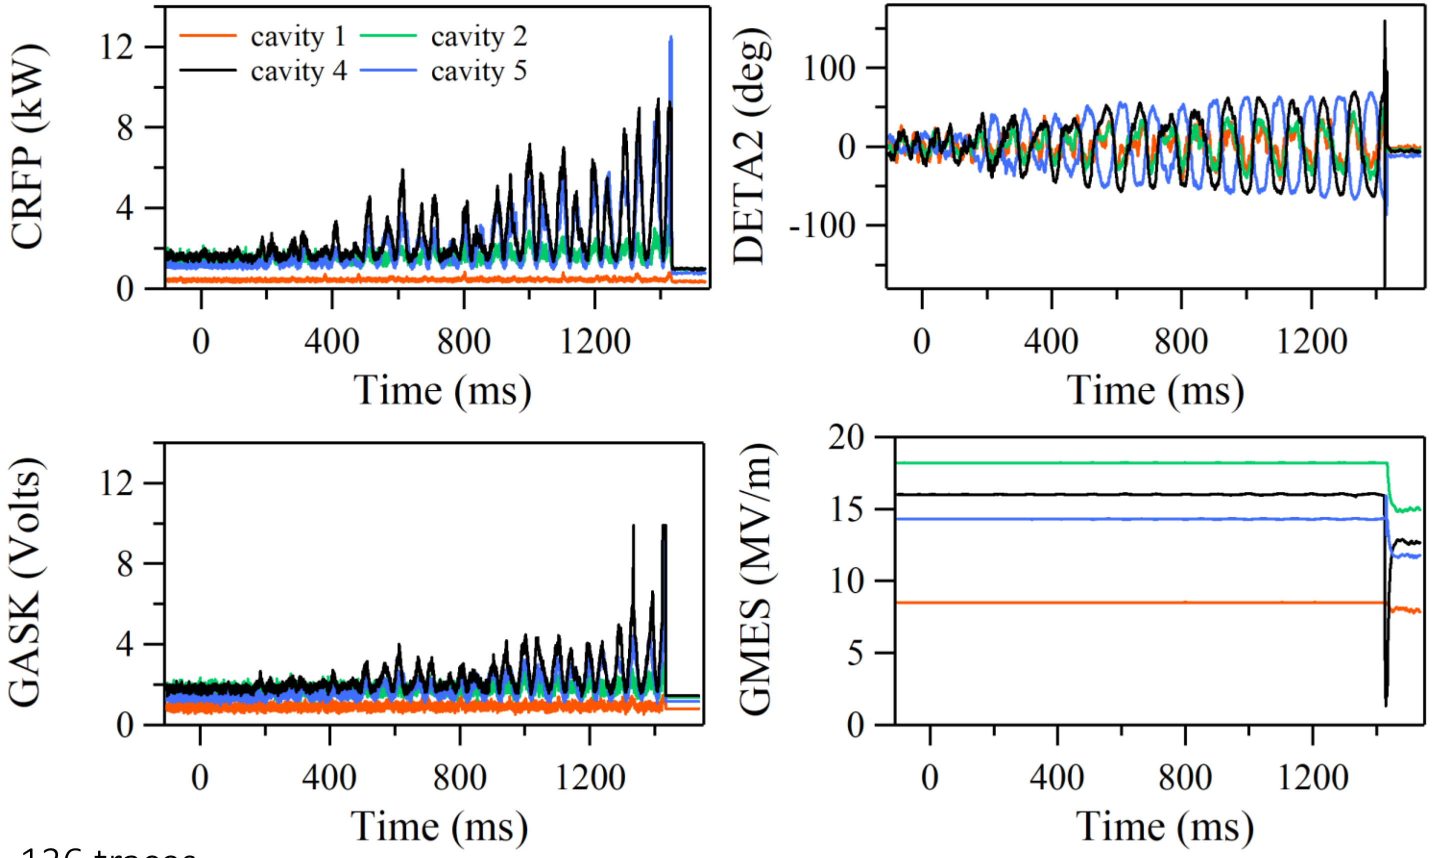

17 signals/cavity × 8 cavities = 136 traces

Benefit of Fault Isolation and Identification

Post-Run Analysis

- use aggregate statistics for data-driven guidance for maintenance and/or upgrade activities
  - ✓analysis of fall 2018 data indicated three cryomodules in the South Linac were prone to microphonic-based faults → provided justification to perform microphonics hardening (installing tuner dampers) → reduced microphonics-based trip rates → gradients could be increased in those cryomodules

Post-Fault Analysis

- provides critical feedback to control room operators
- fault types get mapped to actions for the operators
  - ✓"if Fault *A* happens *X* times within *Y* minutes, drop gradient in the cavity by *Z* MV/m"
  - ✓"if Fault *B* happens *X* times within *Y* minutes, contact a SME"

Workflow

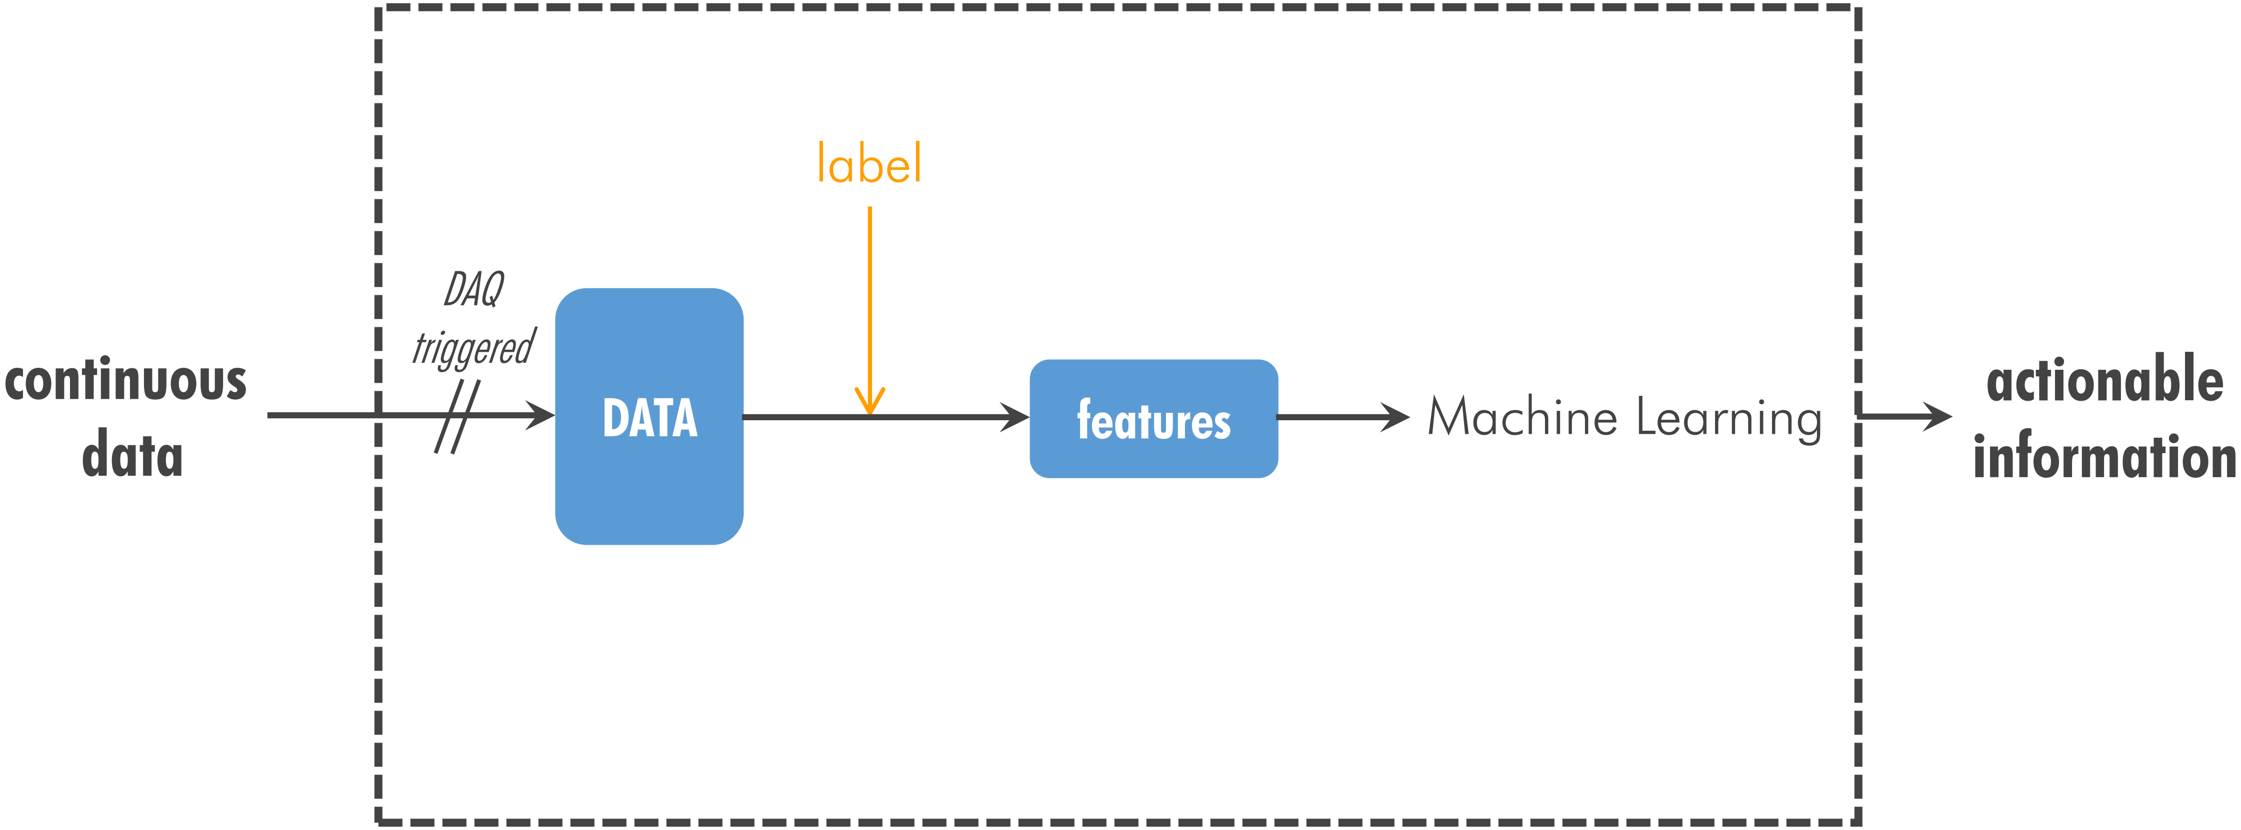
continuous data
DAQ triggered
DATA
label
features
Machine Learning
actionable information

Master Dataset

- perform feature extraction by fitting each time-series signal with 6 autoregressive coefficients:
  - ✓ 8 cavities × 4 signals/cavity × 6 features/signal = 192 features
- data from January 18, 2019 to March 9, 2020
  - ✓ event must include all 8 cavities
  - ✓ must be sampled at 5 kHz (0.20 ms sample time)
- 2,375 events × 192 features

data

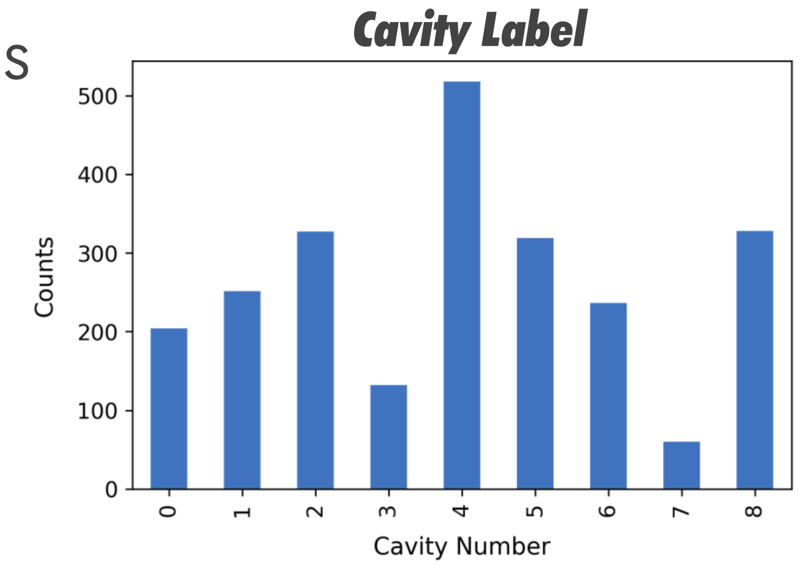

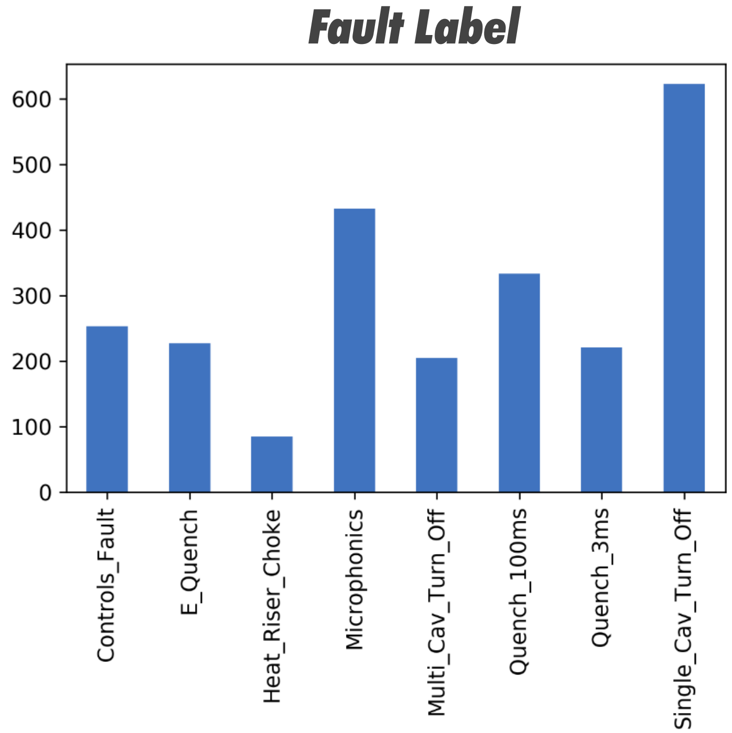

(L. Vidyaratne, JLAB-TN-20-014)

Workflow: Developing ML Models

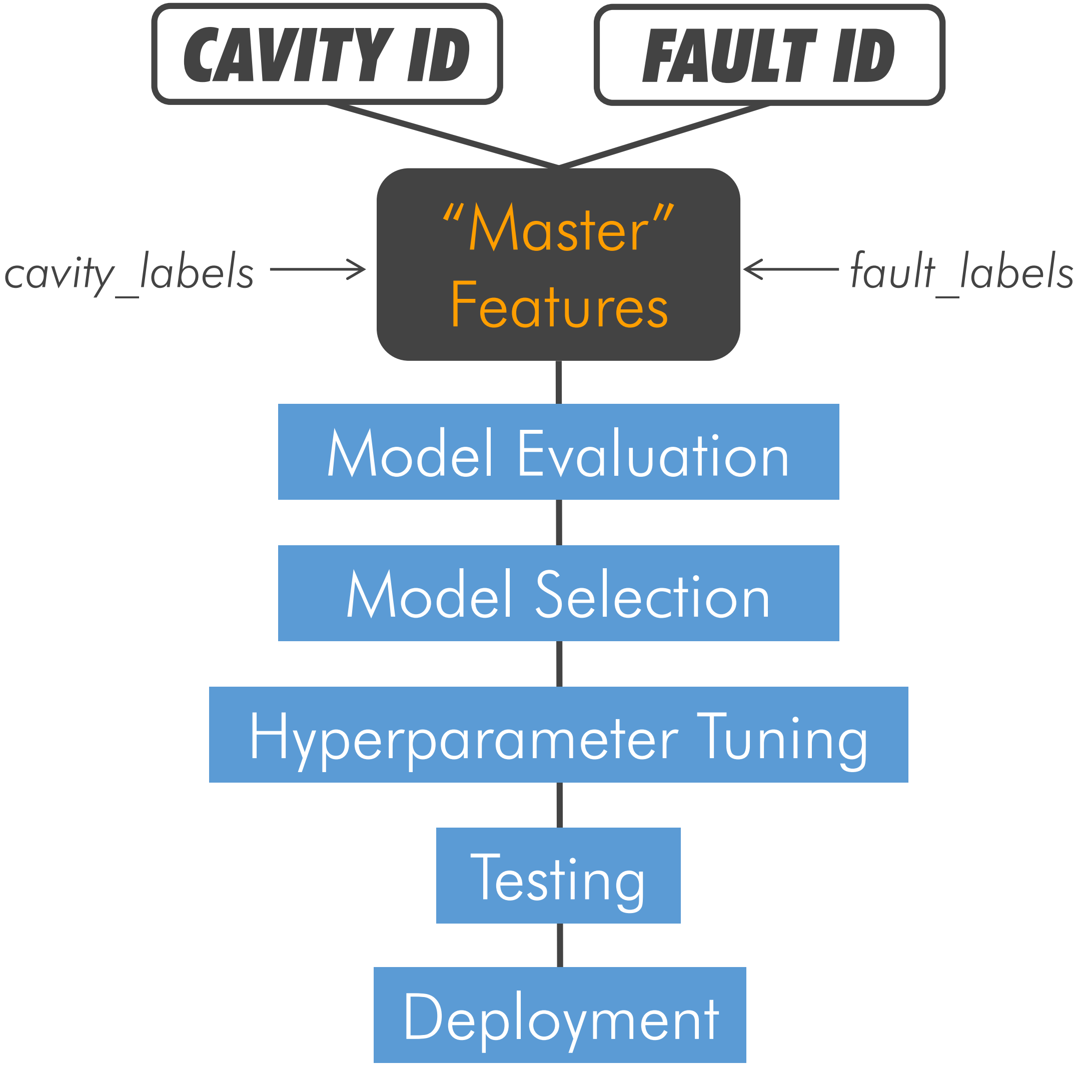

CAVITY ID

FAULT ID

cavity_labels →

"Master" Features

← *fault_labels*

Model Evaluation

Model Selection

Hyperparameter Tuning

Testing

Deployment

Model Evaluation and Selection

- split data into train/test (70%/30%)
- 10-fold cross-validation scores for several different algorithms
  - ✓ ensemble models excel
- perform hyperparameter optimization on Random Forest classifier

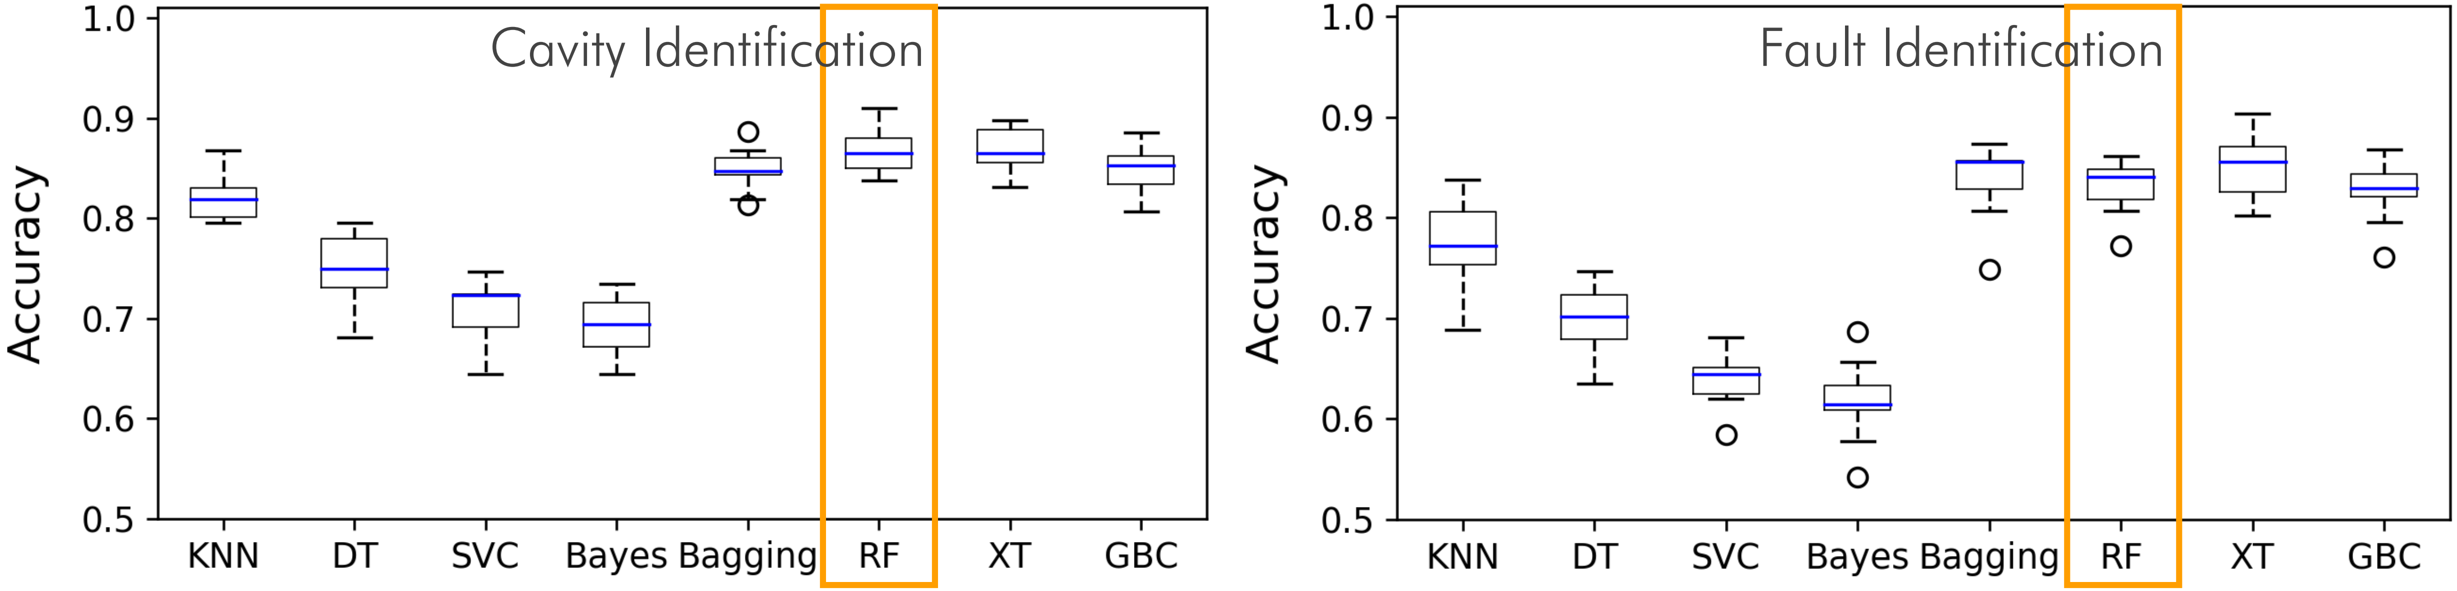

	Cavity Identification	**Fault Type**
10-fold cross-validation (%)	87.97 ± 1.81	85.52 ± 3.65
accuracy (test data) (%)	87.94	87.66

ML Model Performance

- models were applied to data collected from March 10-24, 2020
  - ✓ physics run was prematurely ended due to COVID-19
- 312 fault events were analyzed by the models
- summary of model performances compared to labeled data

	Agree	Disagree	Total
Cavity Model	265	47	312
Fault Model	244	68	312

- cavity model accuracy: 84.9%
  - ✓ testing accuracy: 87.9%
- fault model accuracy: 78.2%
  - ✓ testing accuracy: 87.7%

ML Model Performance

- confusion matrices showing ML model performance

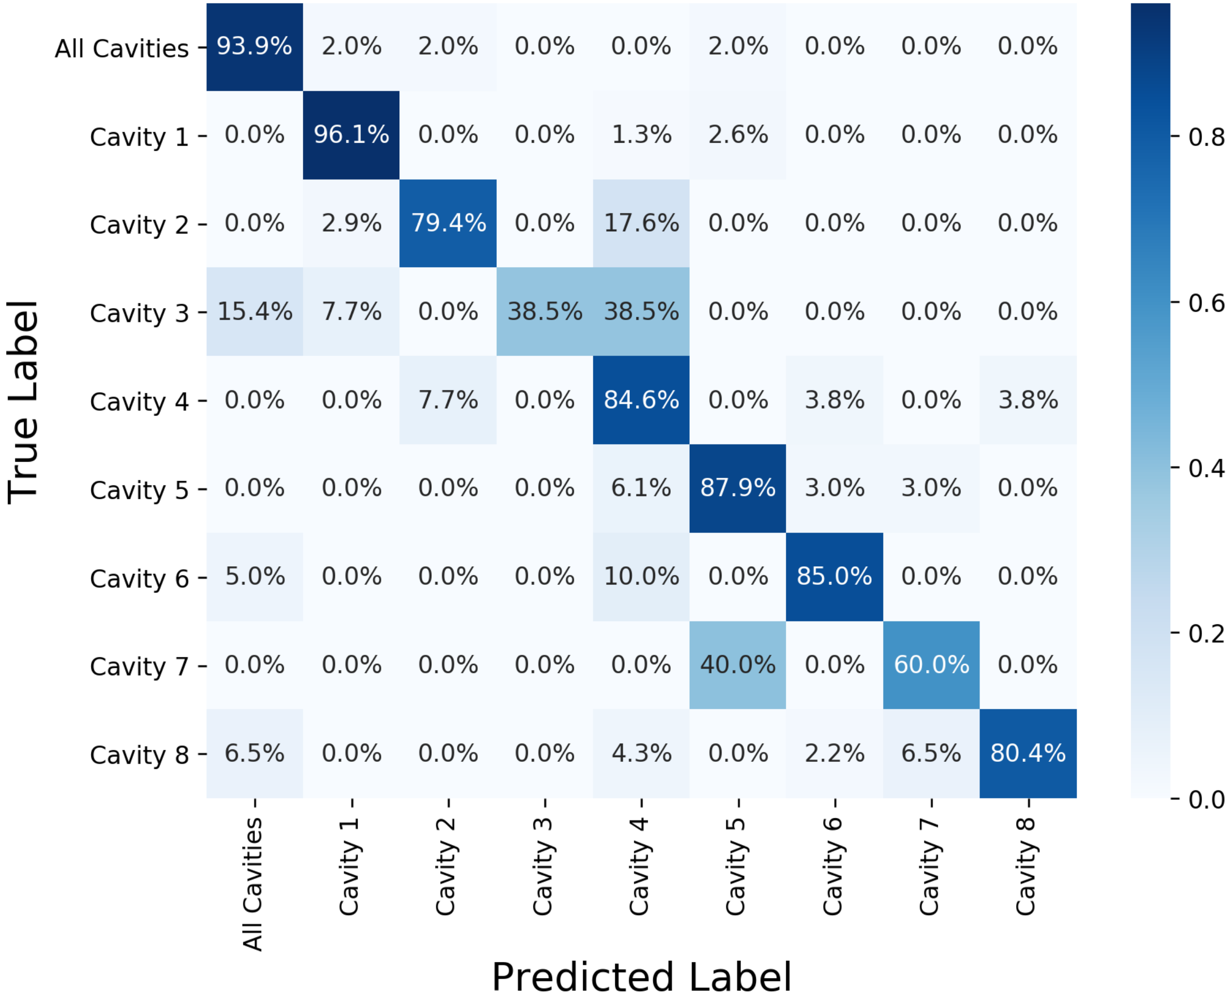

	All Cavities	Cavity 1	Cavity 2	Cavity 3	Cavity 4	Cavity 5	Cavity 6	Cavity 7	Cavity 8
All Cavities	93.9%	2.0%	2.0%	0.0%	0.0%	2.0%	0.0%	0.0%	0.0%
Cavity 1	0.0%	96.1%	0.0%	0.0%	1.3%	2.6%	0.0%	0.0%	0.0%
Cavity 2	0.0%	2.9%	79.4%	0.0%	17.6%	0.0%	0.0%	0.0%	0.0%
Cavity 3	15.4%	7.7%	0.0%	38.5%	38.5%	0.0%	0.0%	0.0%	0.0%
Cavity 4	0.0%	0.0%	7.7%	0.0%	84.6%	0.0%	3.8%	0.0%	3.8%
Cavity 5	0.0%	0.0%	0.0%	0.0%	6.1%	87.9%	3.0%	3.0%	0.0%
Cavity 6	5.0%	0.0%	0.0%	0.0%	10.0%	0.0%	85.0%	0.0%	0.0%
Cavity 7	0.0%	0.0%	0.0%	0.0%	0.0%	40.0%	0.0%	60.0%	0.0%
Cavity 8	6.5%	0.0%	0.0%	0.0%	4.3%	0.0%	2.2%	6.5%	80.4%

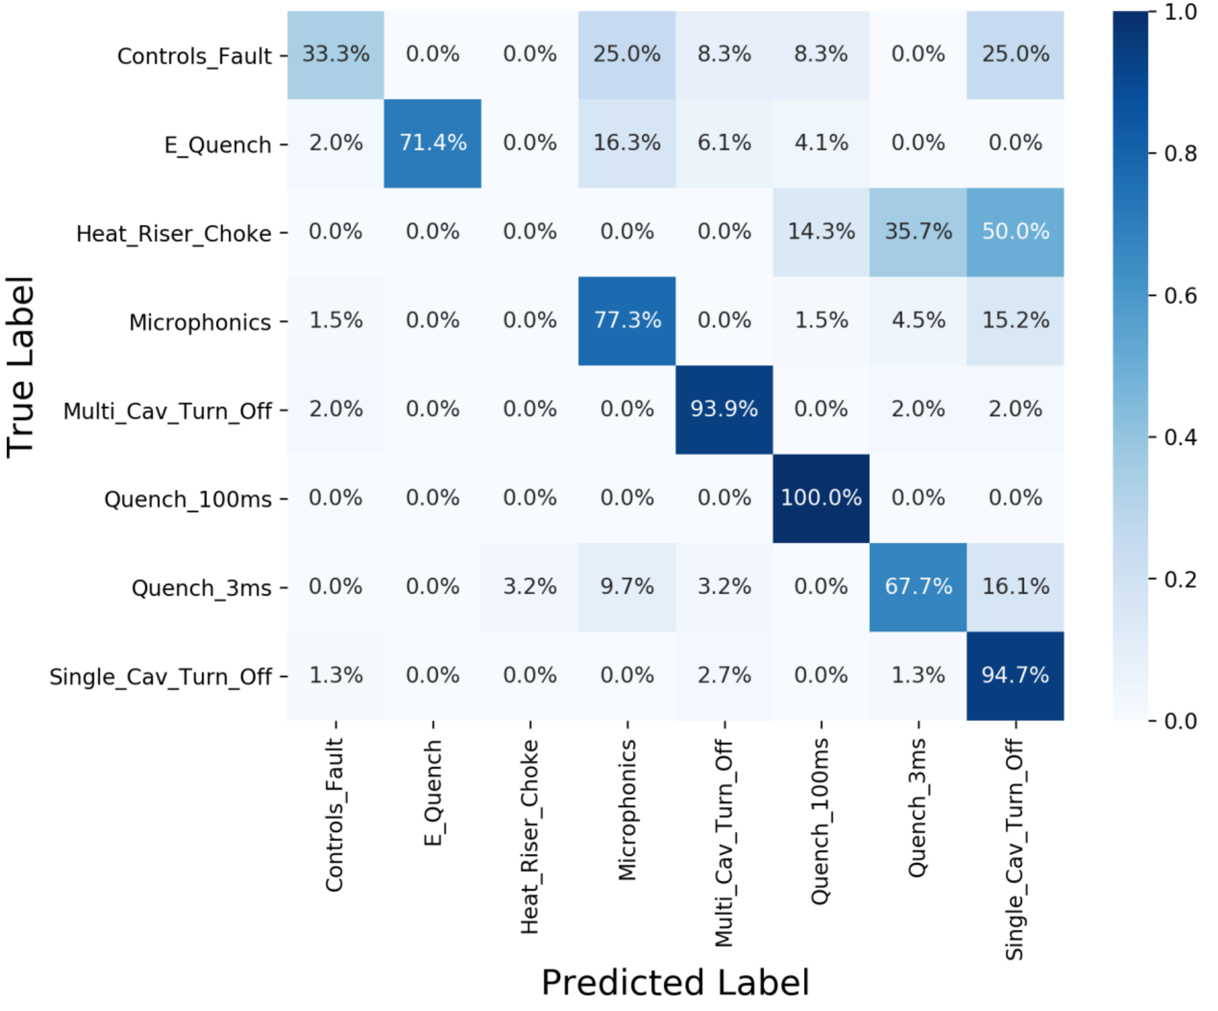

	Controls_Fault	E_Quench	Heat_Riser_Choke	Microphonics	Multi_Cav_Turn_Off	Quench_100ms	Quench_3ms	Single_Cav_Turn_Off
Controls_Fault	33.3%	0.0%	0.0%	25.0%	8.3%	8.3%	0.0%	25.0%
E_Quench	2.0%	71.4%	0.0%	16.3%	6.1%	4.1%	0.0%	0.0%
Heat_Riser_Choke	0.0%	0.0%	0.0%	0.0%	0.0%	14.3%	35.7%	50.0%
Microphonics	1.5%	0.0%	0.0%	77.3%	0.0%	1.5%	4.5%	15.2%
Multi_Cav_Turn_Off	2.0%	0.0%	0.0%	0.0%	93.9%	0.0%	2.0%	2.0%
Quench_100ms	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%	0.0%	0.0%
Quench_3ms	0.0%	0.0%	3.2%	9.7%	3.2%	0.0%	67.7%	16.1%
Single_Cav_Turn_Off	1.3%	0.0%	0.0%	0.0%	2.7%	0.0%	1.3%	94.7%

Visualization and Communication

- for ML models to be effective, information must be communicated clearly and concisely
- visualize spatial and temporal nature of model predictions

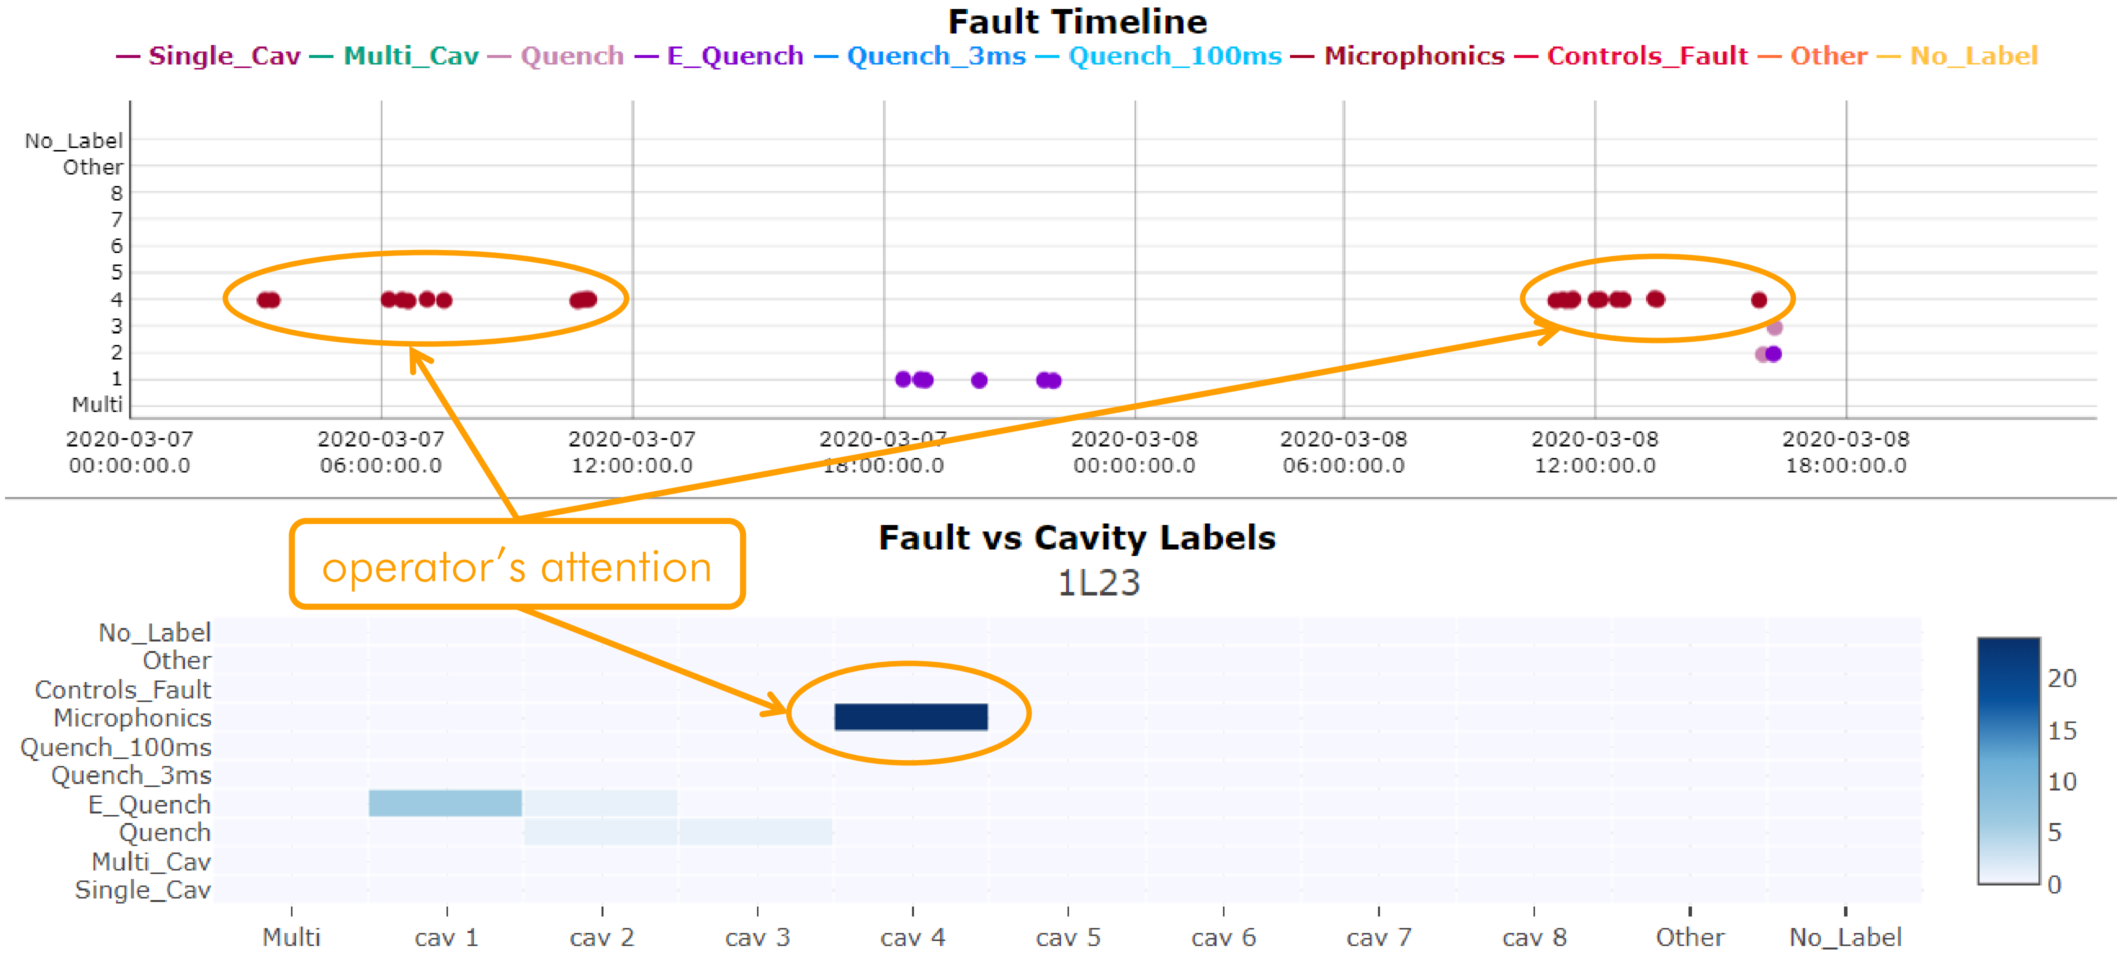

(C. Tennant, PRAB **23**, 114601 (2020))

Post-Fault: Actionable Information

- cavity 8 in cryomodule 2L26 plagued by electronic quenches

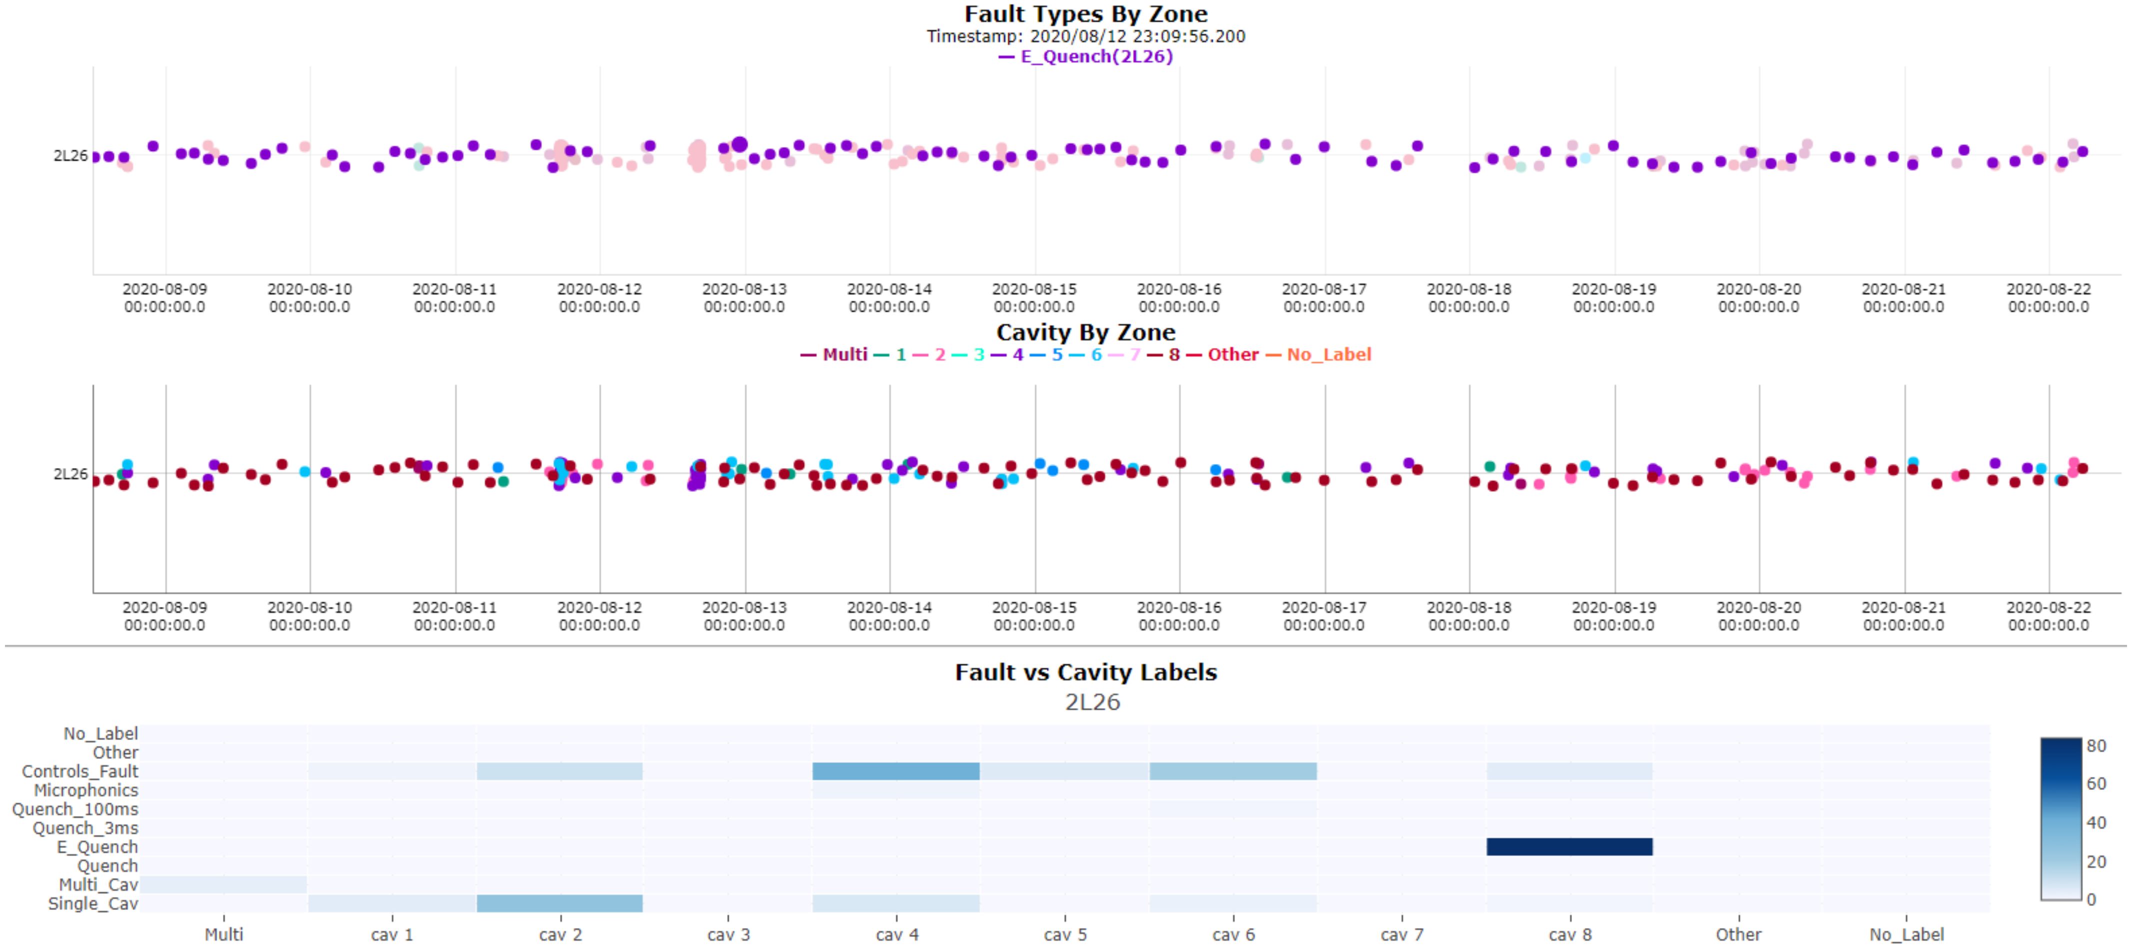

Post-Fault: Actionable Information

- turn down gradient 9/5/2020 and faults went away completely

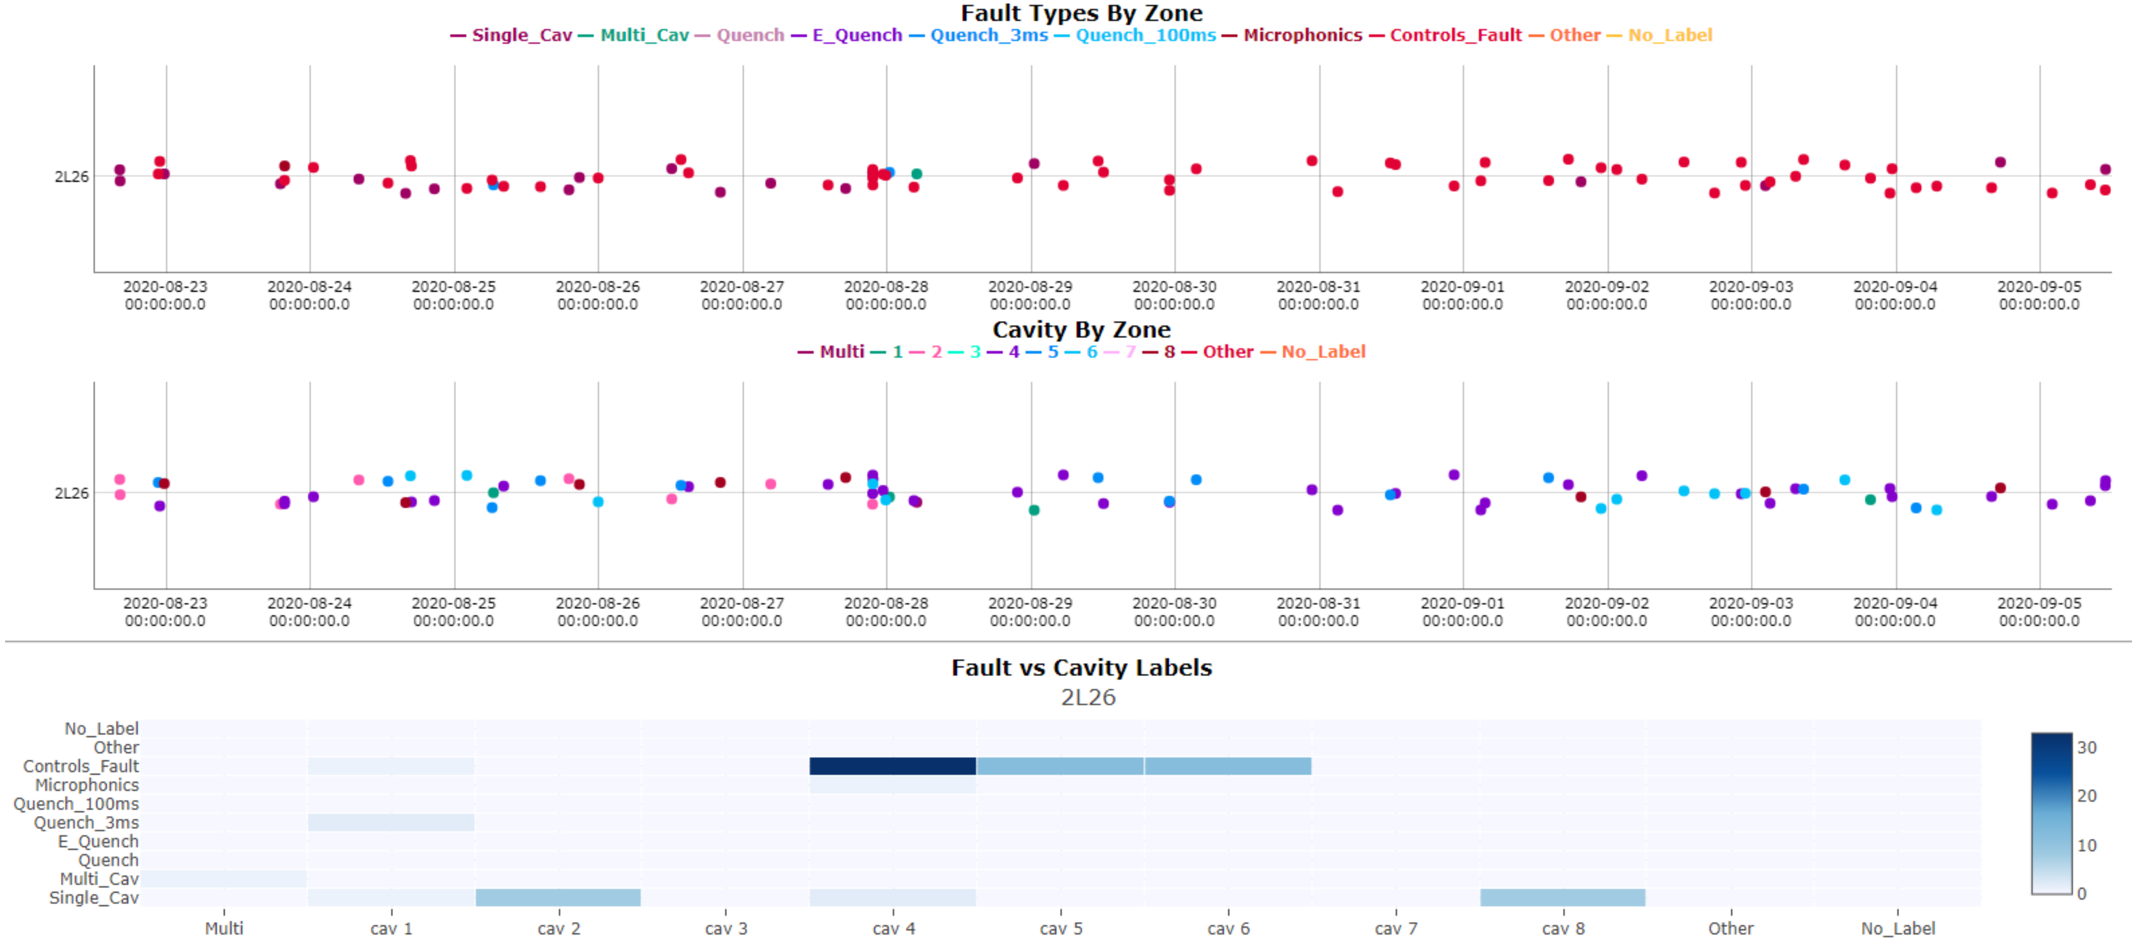

Post-Run: Actionable Information

North Linac

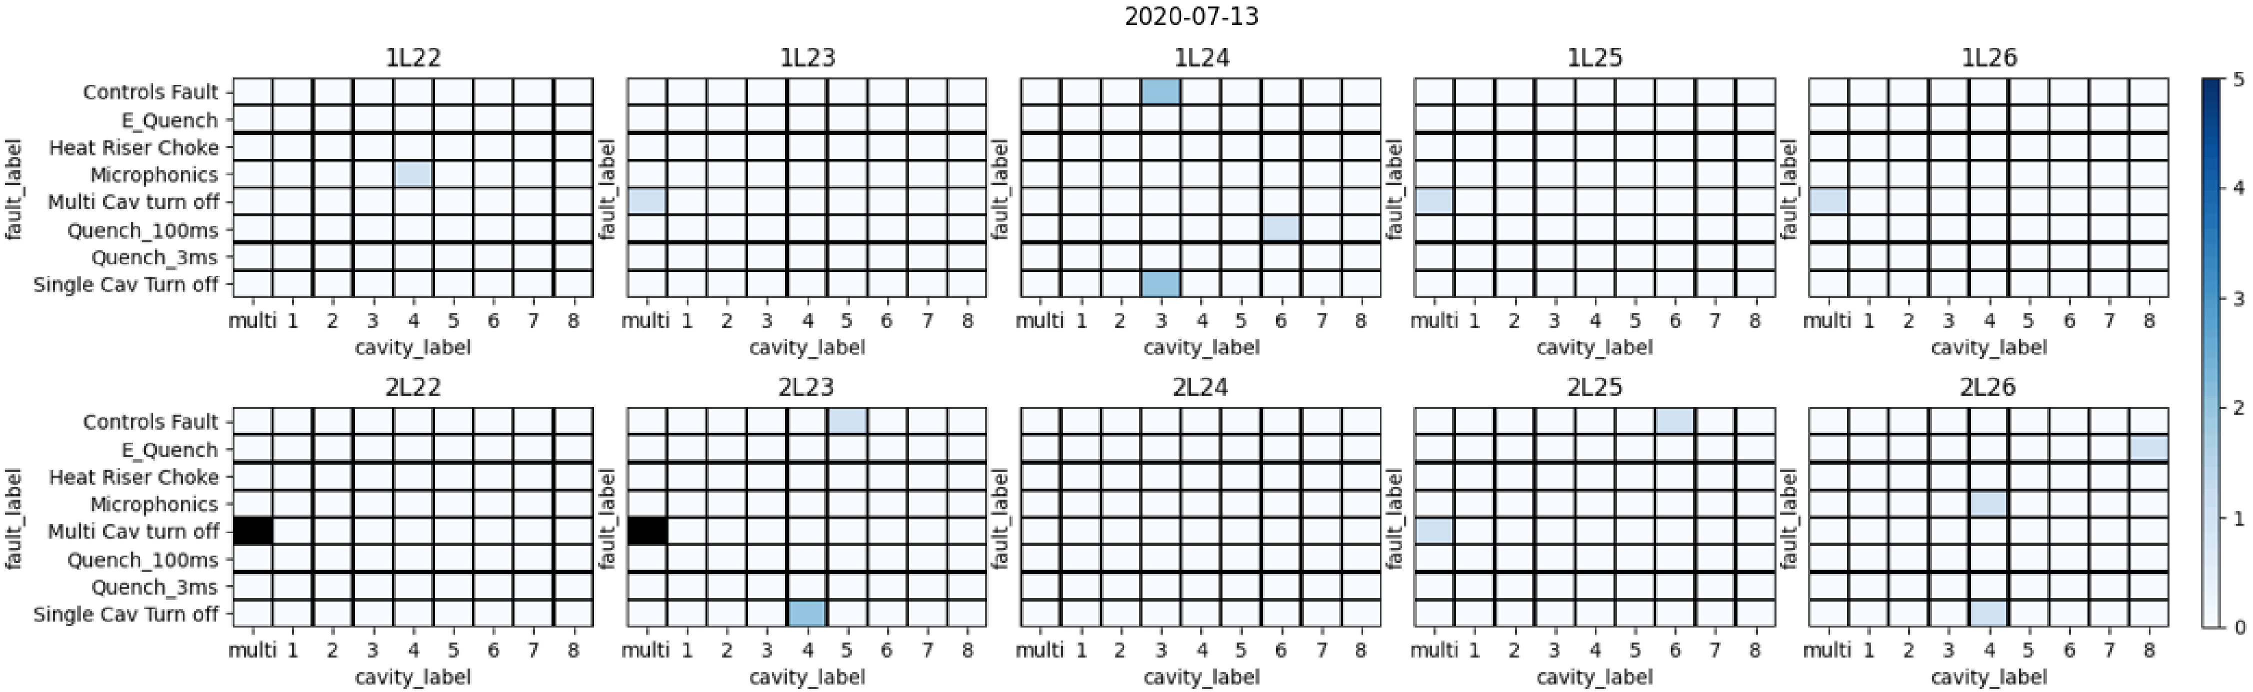

South Linac

Workflow

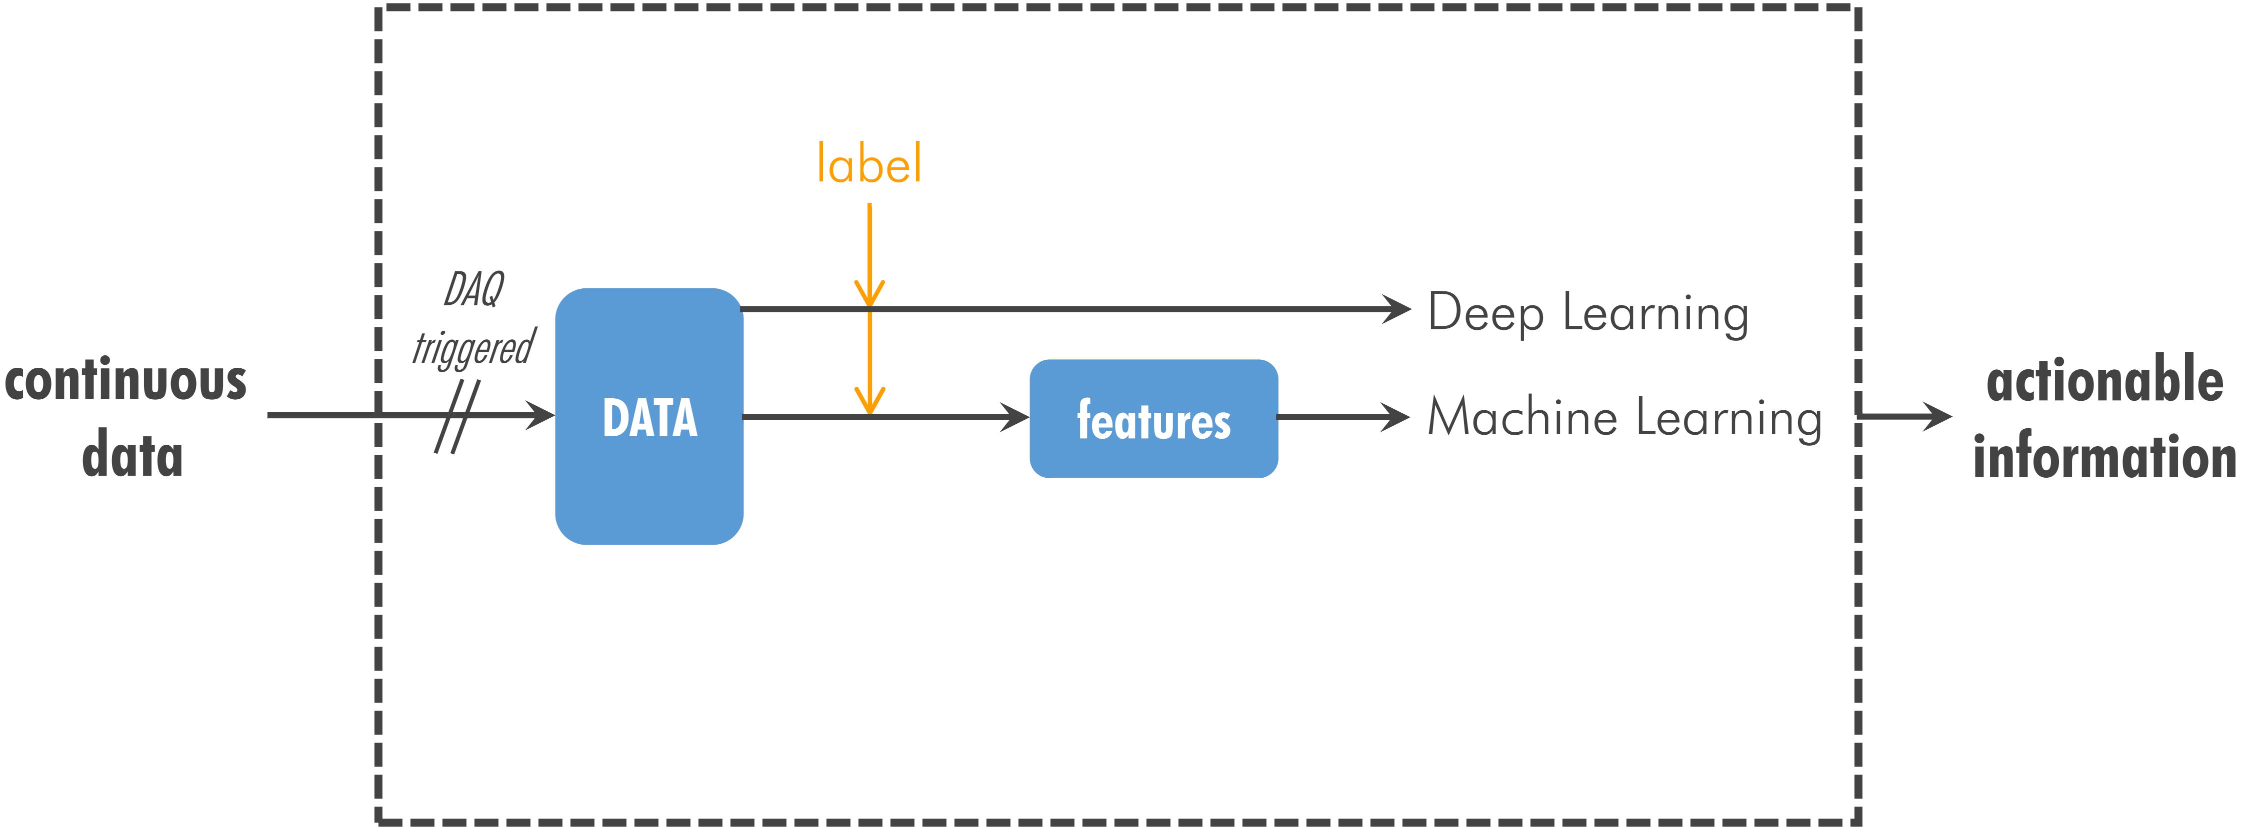

continuous data

DAQ triggered

DATA

label

features

Deep Learning

Machine Learning

actionable information

Deep Recurrent Architecture

- bidirectional LSTM layers for temporal feature learning
- training for simultaneous classification of cavity and fault: two-branch model
- training/validation/test (60%/20%/20%) stratified sampling

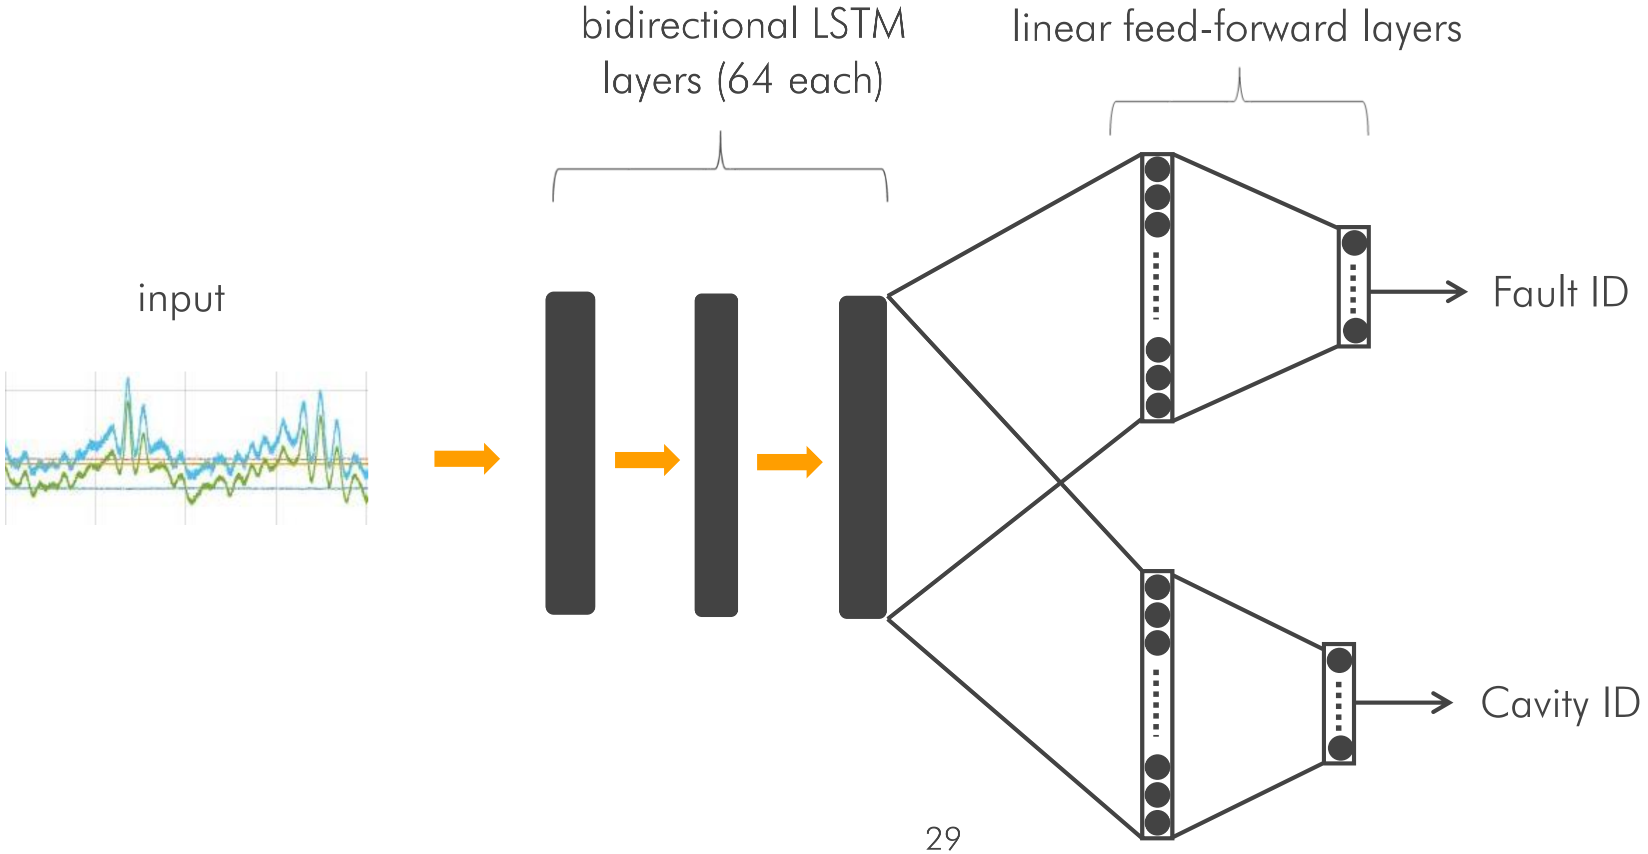

Deep Recurrent Architecture Results

	Cavity Classification		Fault Classification	
	Input Size	*Test Accuracy (%)*	*Input Size*	*Test Accuracy (%)*
17 waveforms/cavity	136×256	86.1	136×256	82.1
4 waveforms/cavity	32×256	87.7	32×256	81.3

- with more data, deep learning approaches the accuracies of the machine learning models
- additionally, several convolutional neural network (CNN) architectures were investigated yielding comparable results

Workflow

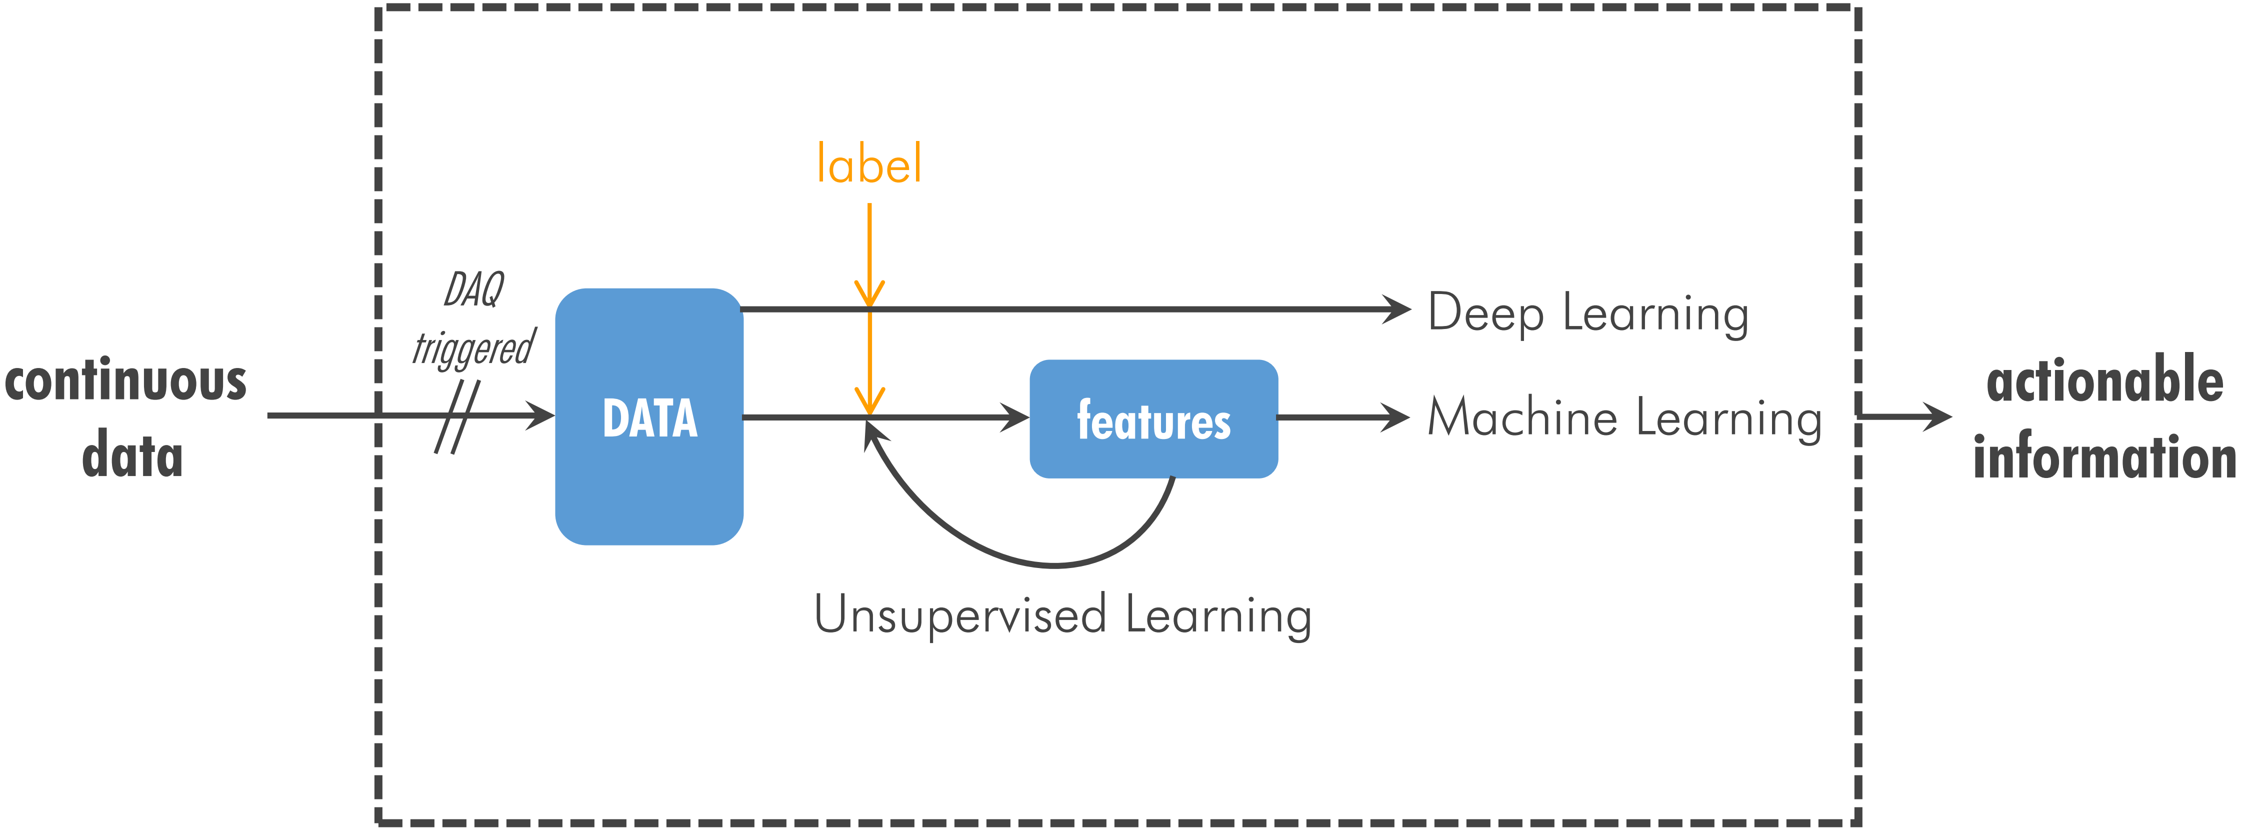

continuous data

DAQ triggered

DATA

label

Deep Learning

features

Machine Learning

Unsupervised Learning

actionable information

Unsupervised Learning

Dimensionality Reduction

- provides a means of compressing the features into a lower dimensional space
- allow for visualization of higher-dimension datasets
- speeds up training and inference time of machine learning models

Clustering

- by grouping data that are similar into clusters, underlying structure and patterns emerge that offer useful insights into the dataset

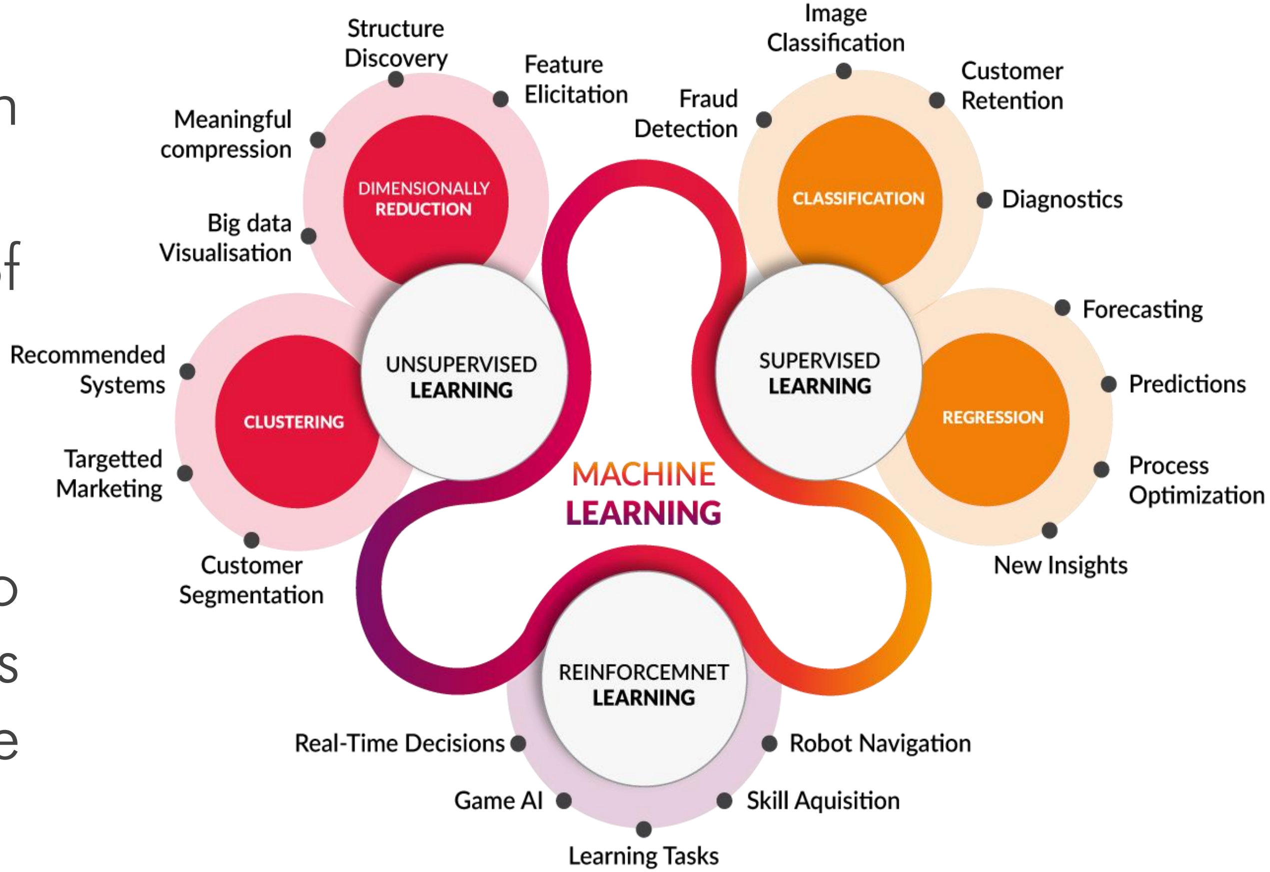

MNIST Example

- consider familiar MNIST, handwritten digit dataset
- this is a very high dimensional space
- with dimensionality reduction can visualize in 2D
- if unlabeled, could apply clustering techniques

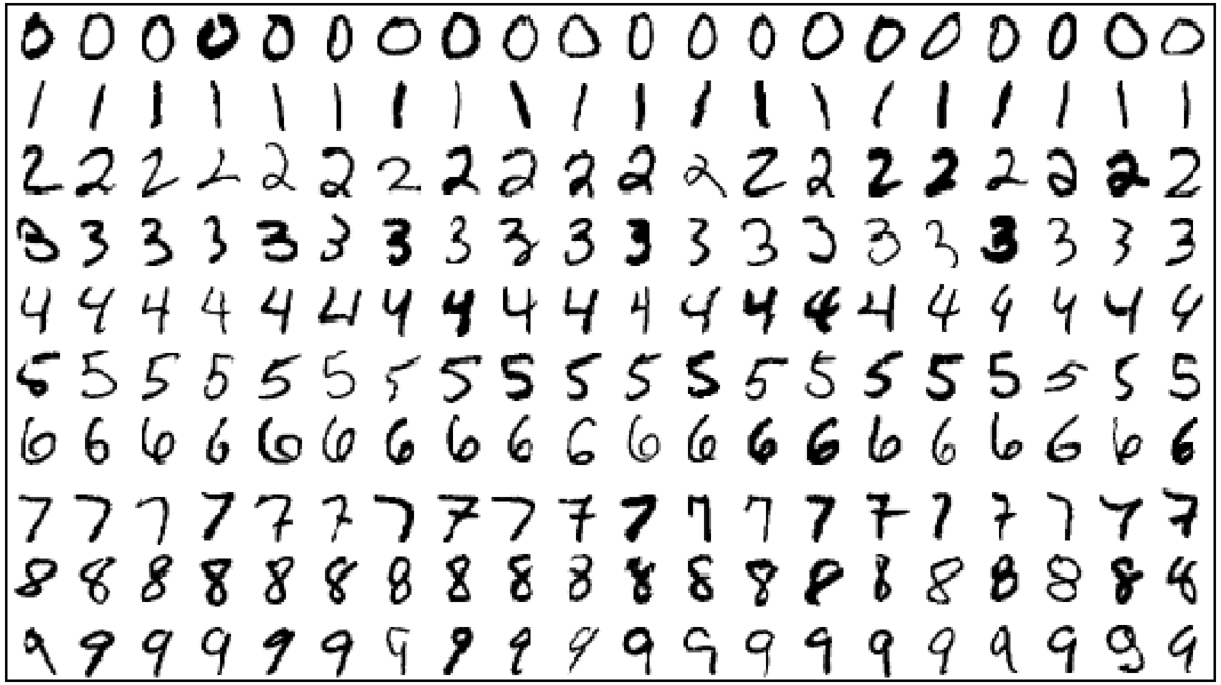

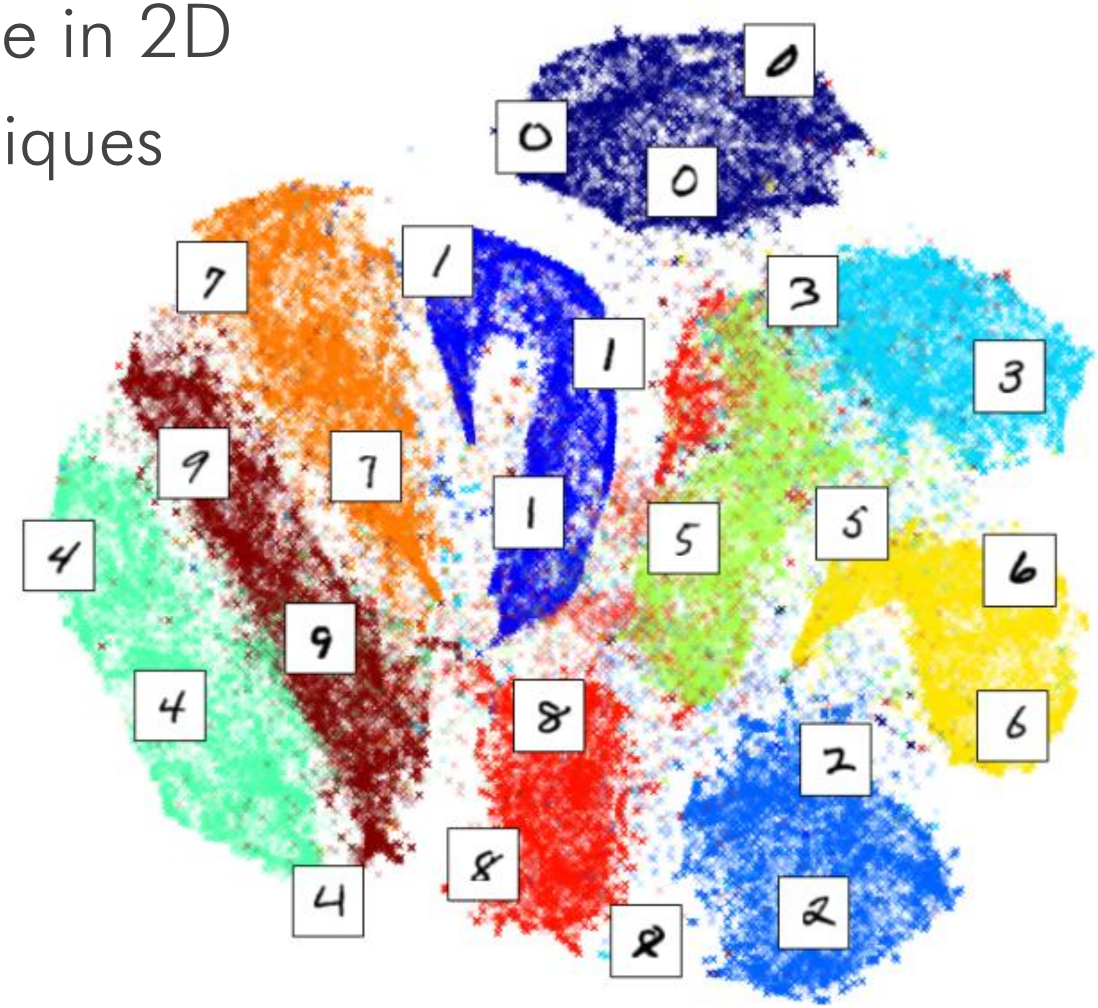

Dimensionality Reduction: 2D

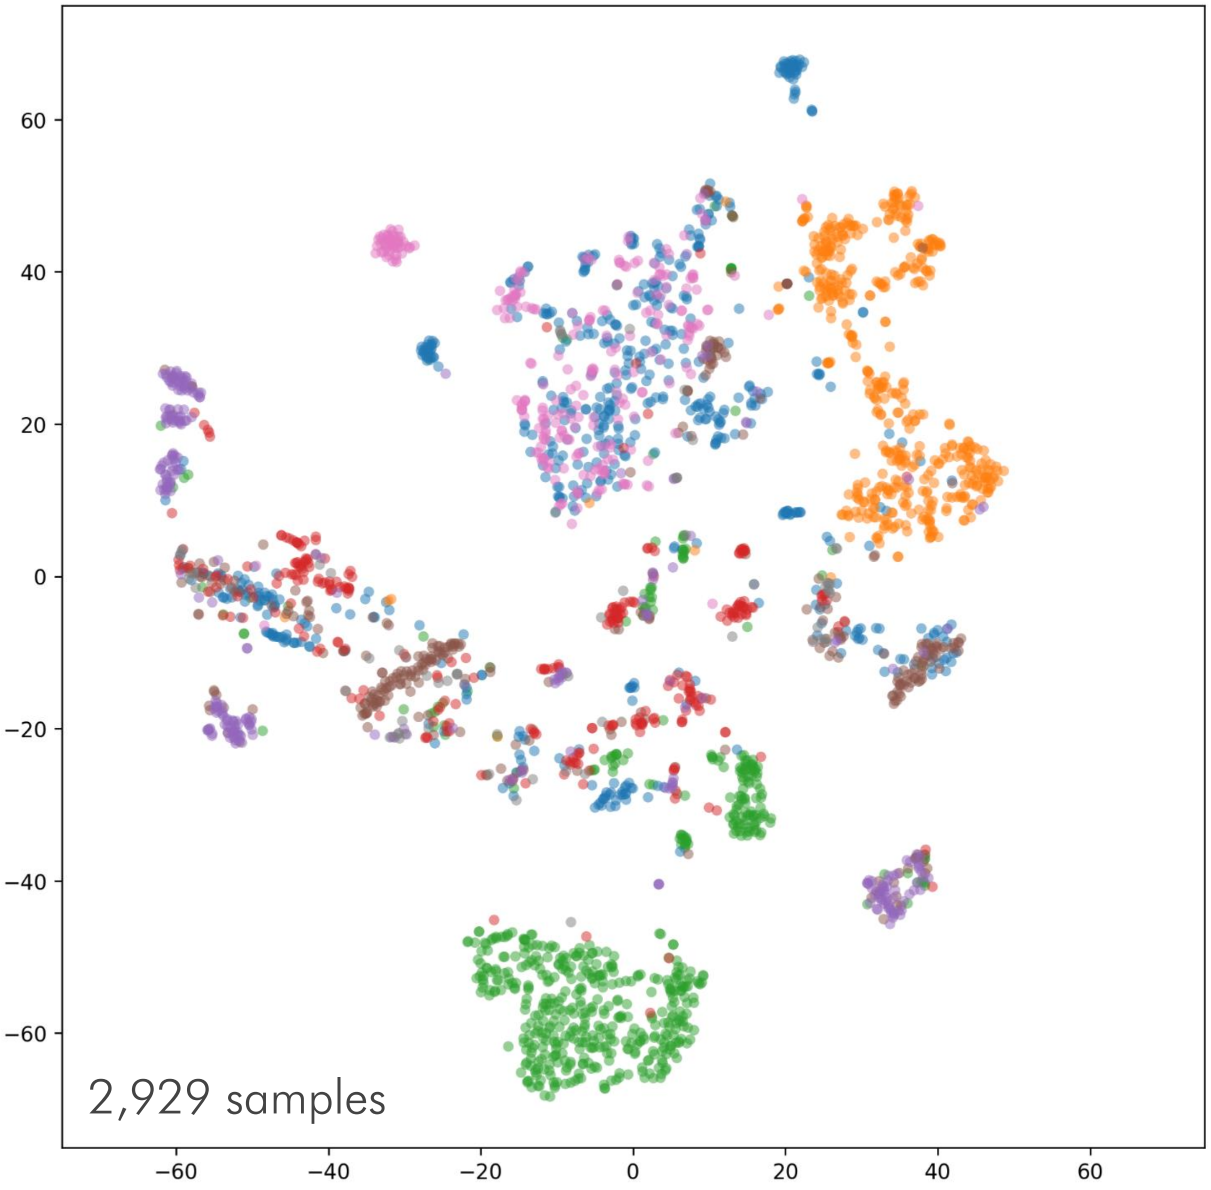

Heat Riser Choke (636)
Controls Fault (598)
Electronic Quench (469)
3 ms Quench (330)
Microphonics (284)
100 ms Quench (278)
Single Cavity Turn-Off (270)
Unknown (64)

Dimensionality Reduction: 3D

- clusters are evident
- however, clusters of same fault-type are often separated

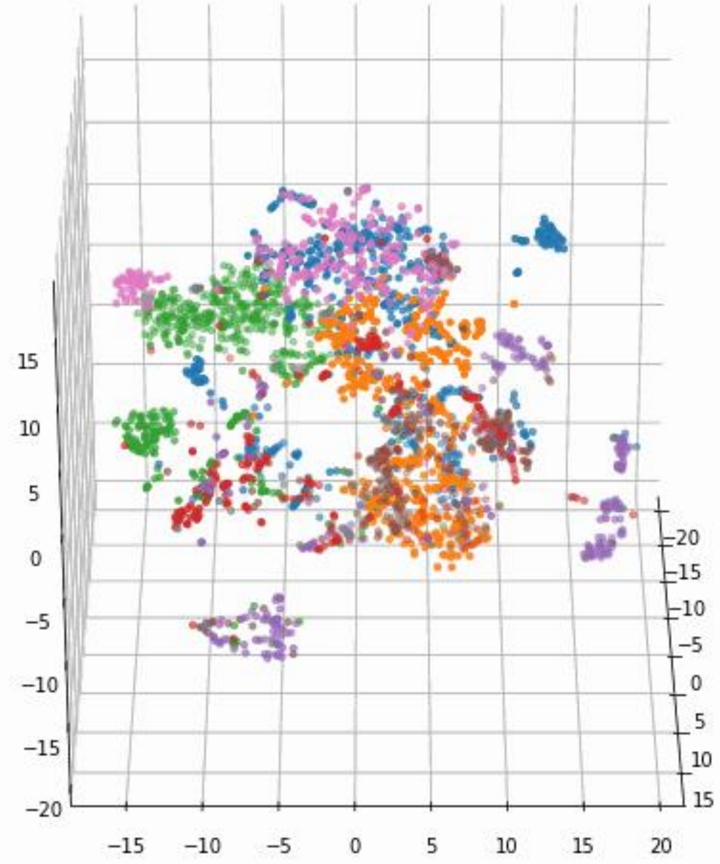

Heat Riser Choke (636)
Controls Fault (598)
Electronic Quench (469)
3 ms Quench (330)
Microphonics (284)
100 ms Quench (278)
Single Cavity Turn-Off (270)
Unknown (64)

2,929 samples

Dimensionality Reduction

Single Cavity Turn-off

Controls Faults in Cavities 3 and 4

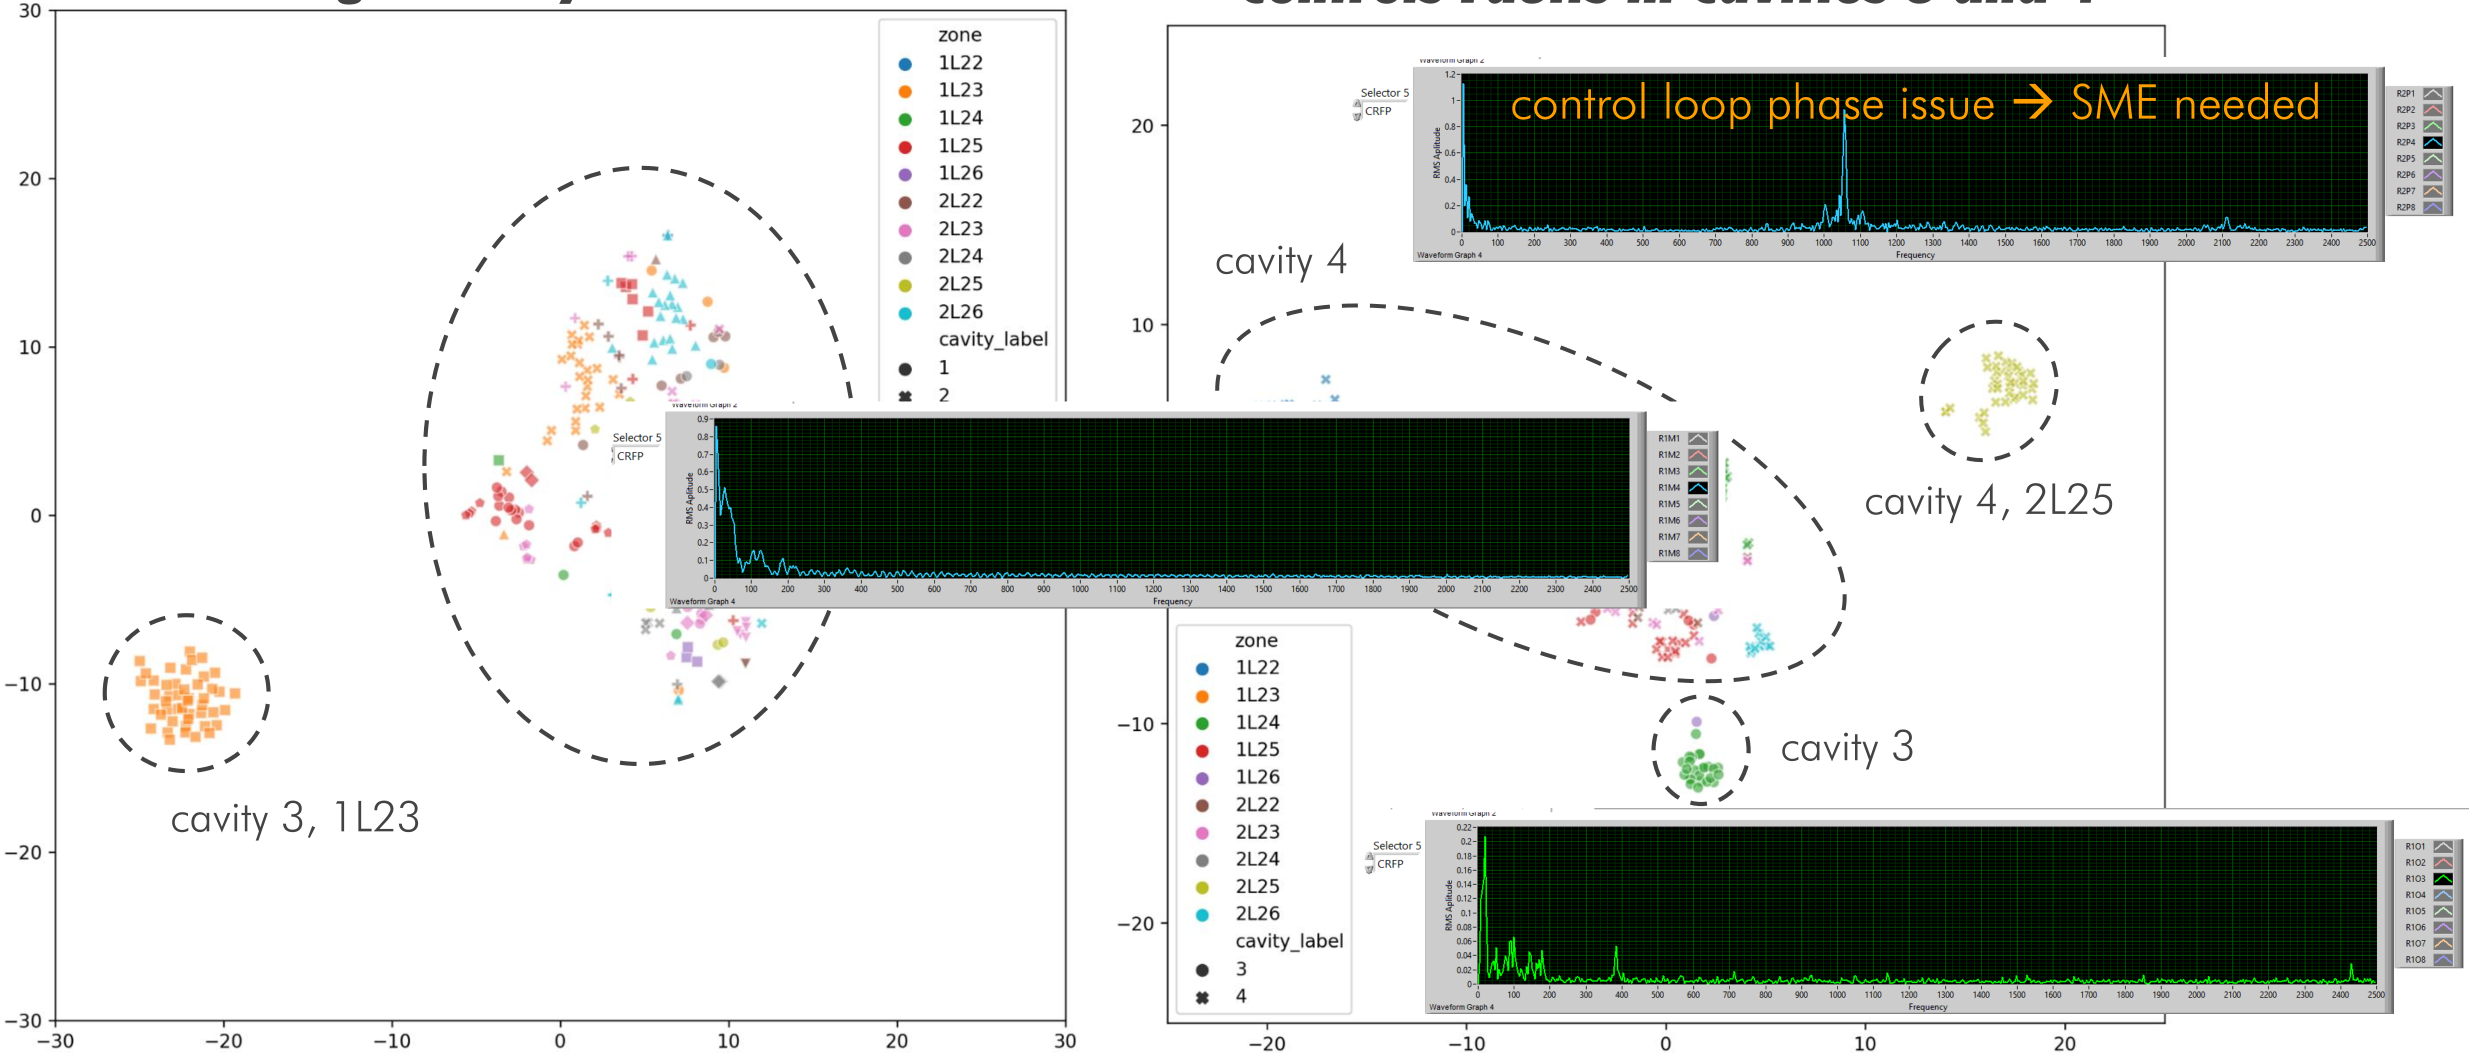

cavity 3, 1L23

cavity 4

cavity 4, 2L25

cavity 3

control loop phase issue → SME needed

Workflow

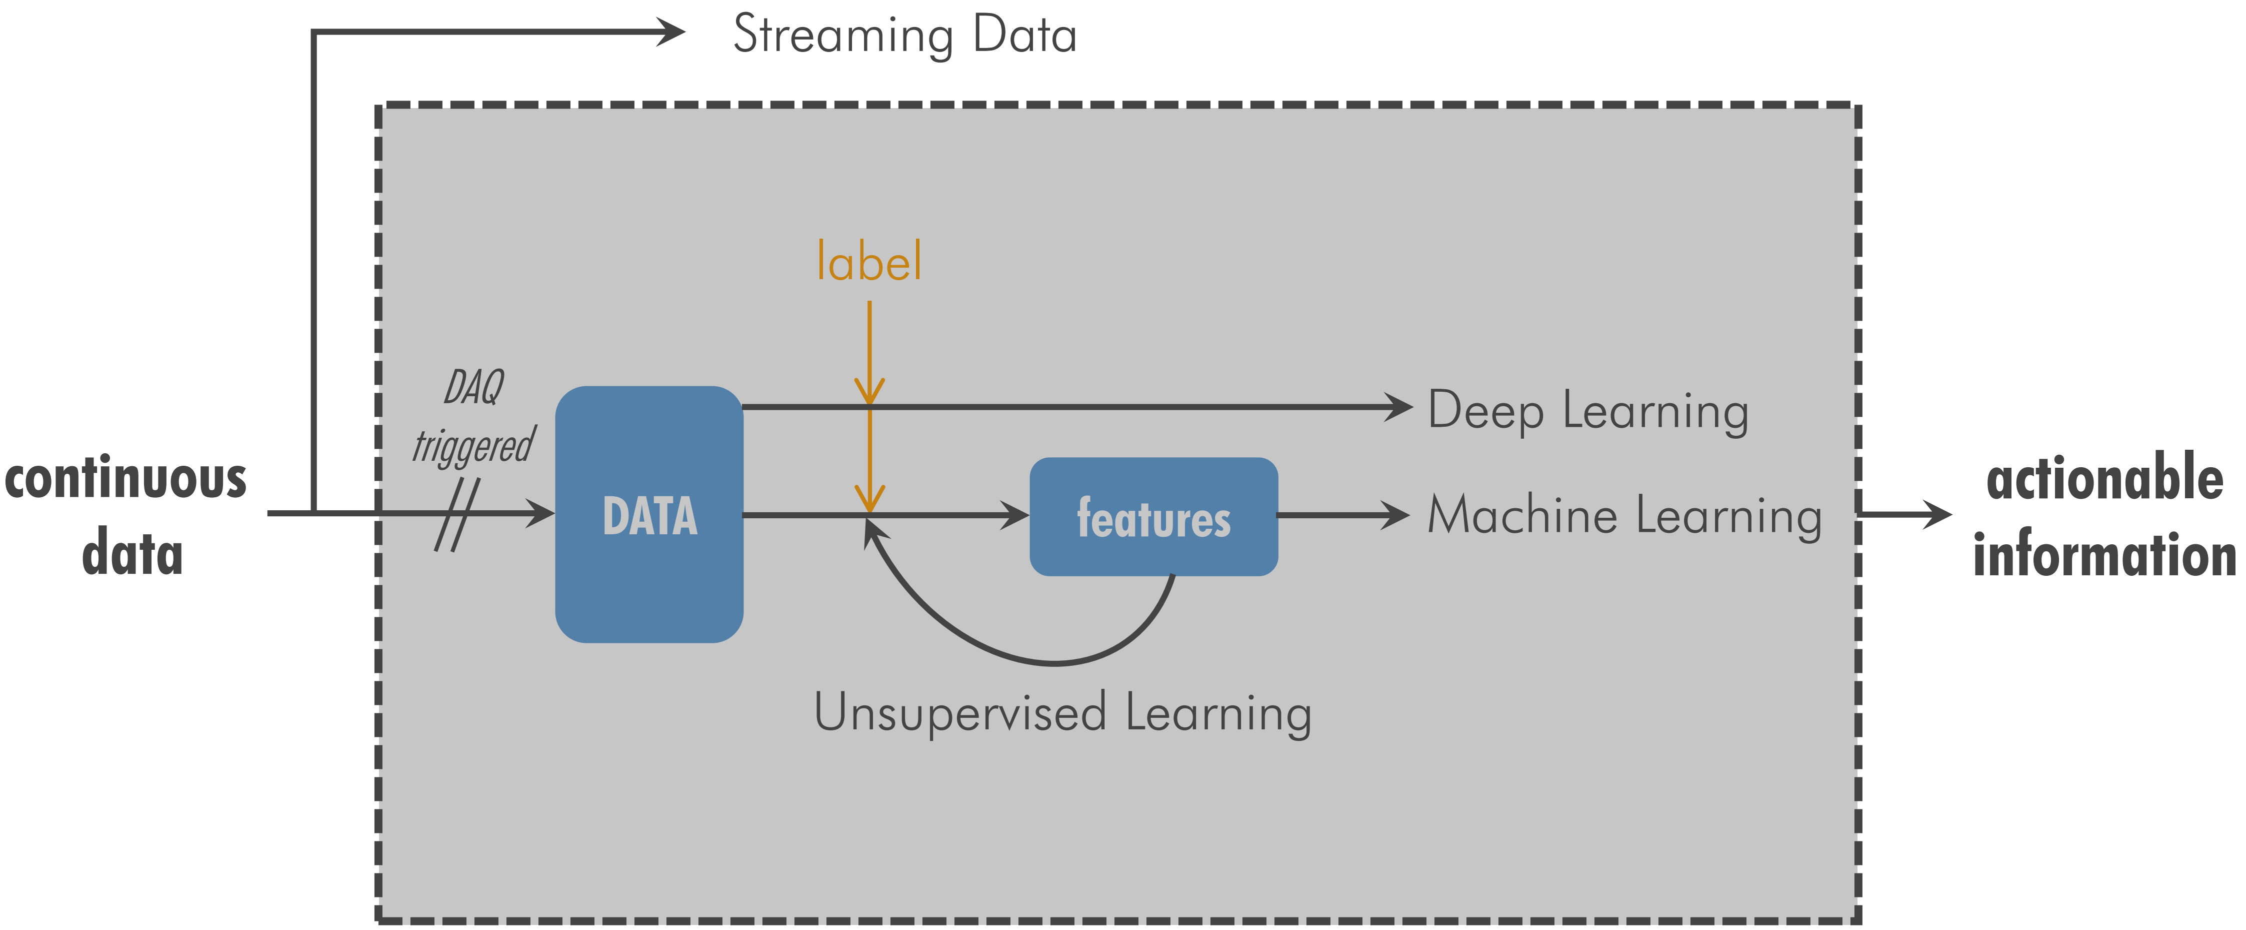
Streaming Data
label
DAQ
triggered
continuous
data
DATA
Deep Learning
features
Machine Learning
actionable
information
Unsupervised Learning

C100 Fault Prediction: Future

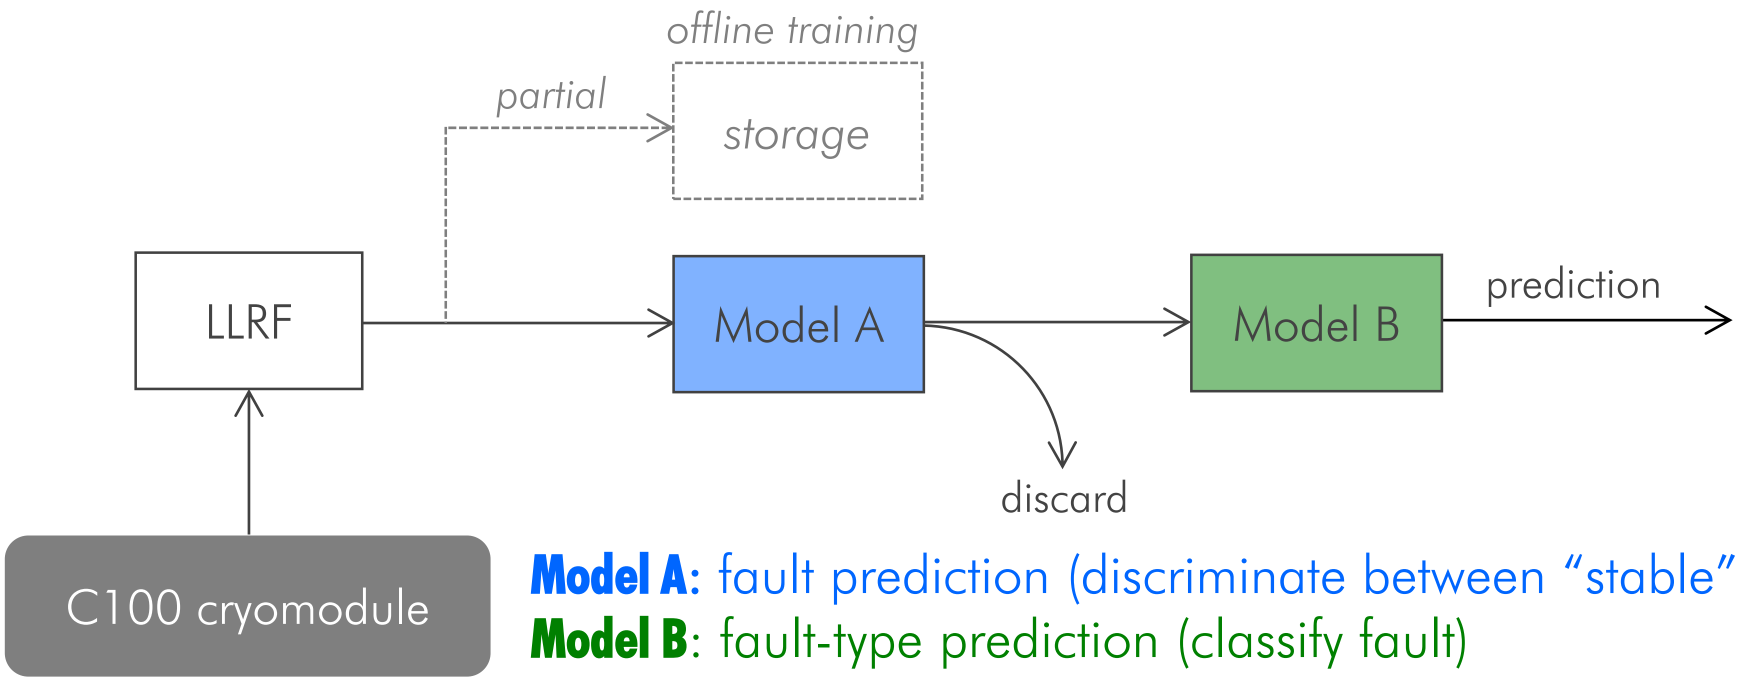

Model A: fault prediction (discriminate between “stable” and “impending”)
Model B: fault-type prediction (classify fault)

- learning from data streams requires:
  - ✓ ability to process an example, inspect it only once, after which the data is discarded
  - ✓ using a limited amount of memory
  - ✓ the ability of models to predict at any point

From Isolation and Identification to Prediction

- initial step: discriminate between “stable” and “impending” fault conditions
  - ✓ use saved waveforms

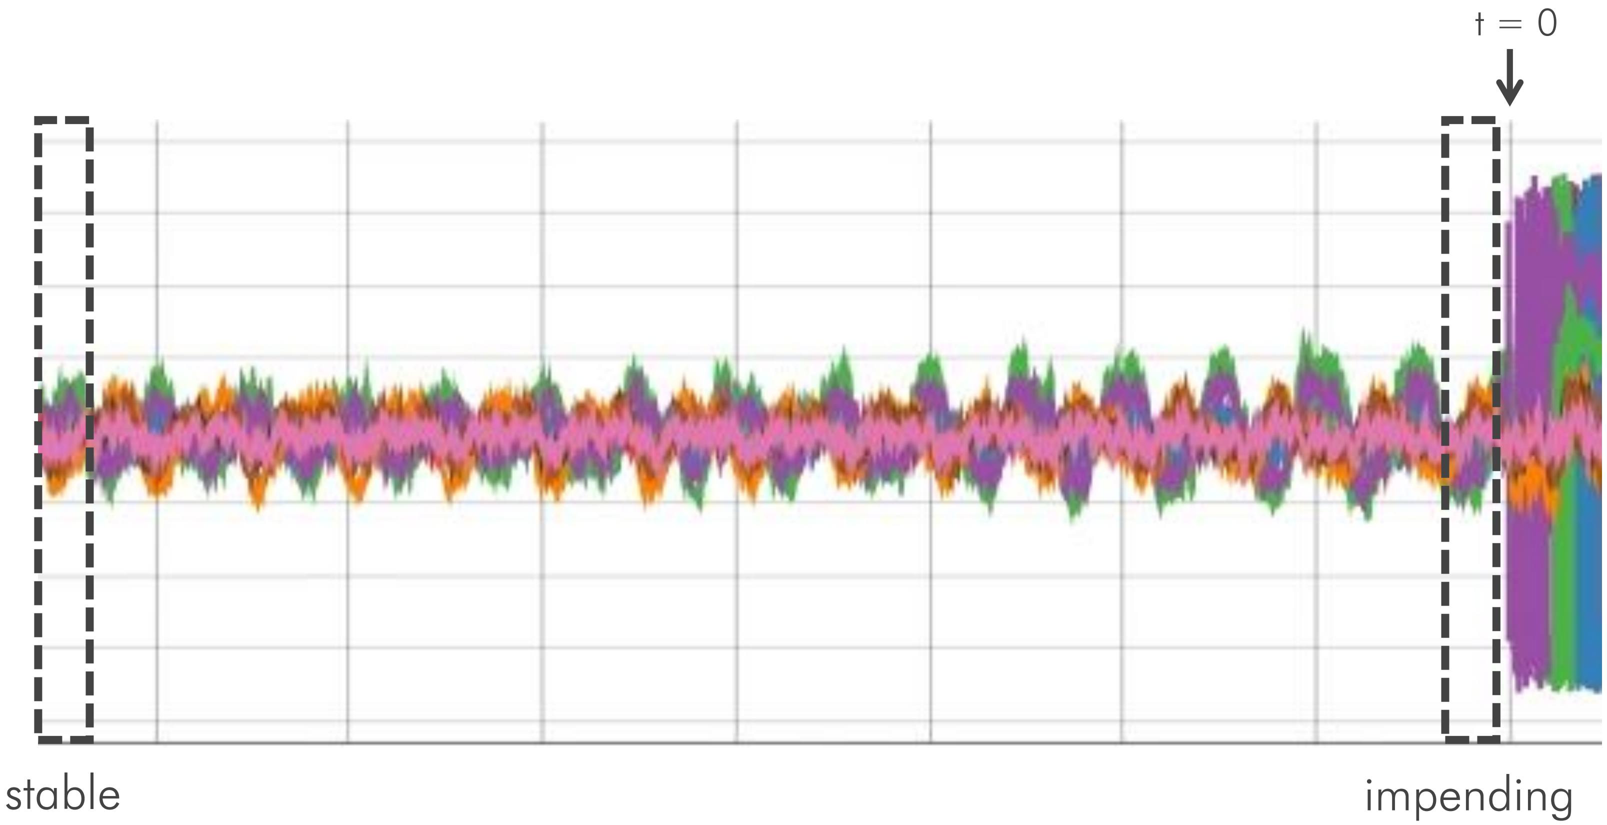

Initial Step: Binary Classifier

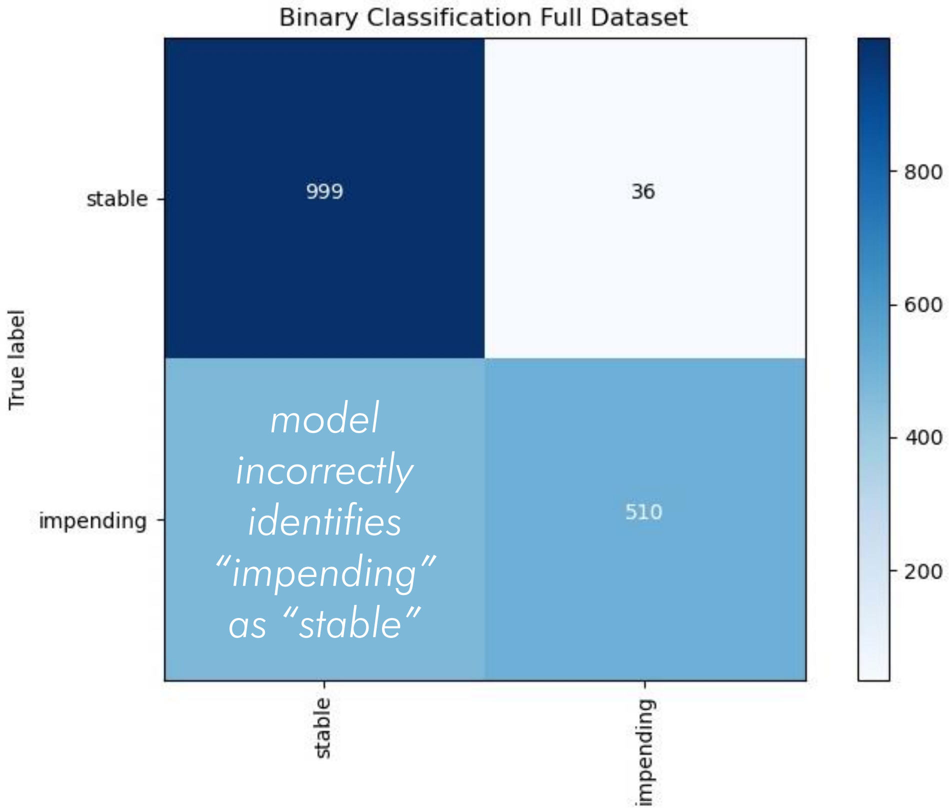

model incorrectly identifies "impending" as "stable"

accuracy = 74.74%

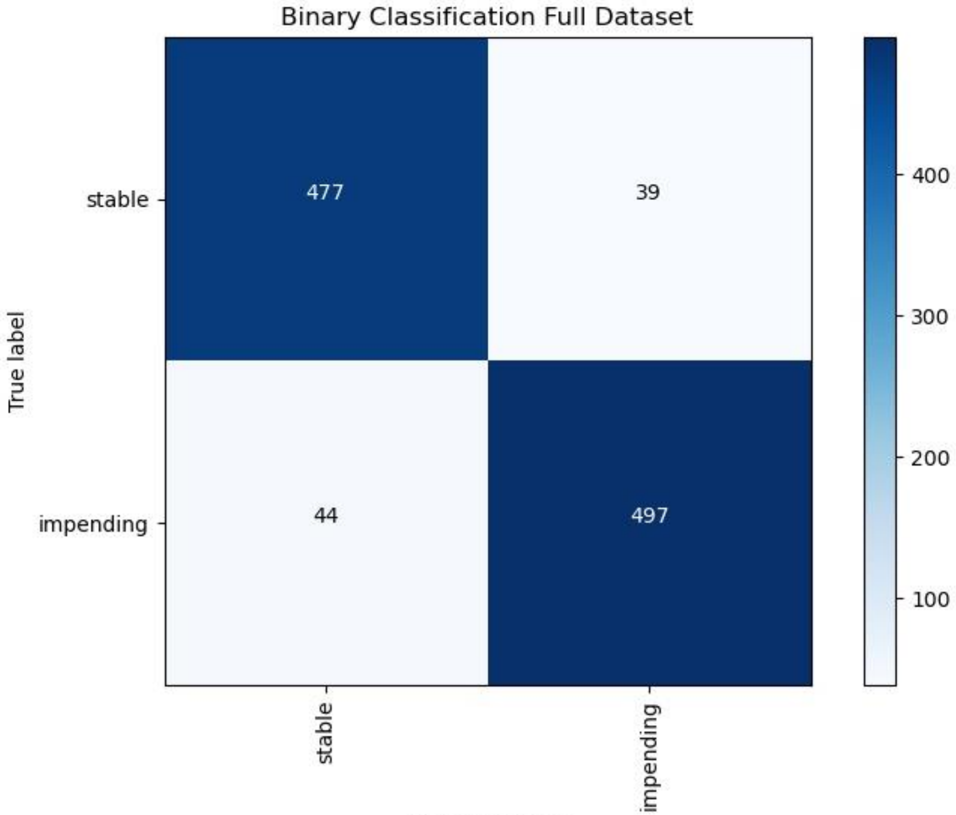

accuracy = 92.13%

- remove fault types which do not show any precursors

➡

	Precision	Recall	f1-score	Support
Stable	0.9155	0.9244	0.9199	516
Impending	0.9272	0.9186	0.9229	541
Accuracy	0.9213			

Intermediate Step: Sliding Window

- can data prior to event accurately predict the fault type?
  - ✓ use saved waveforms

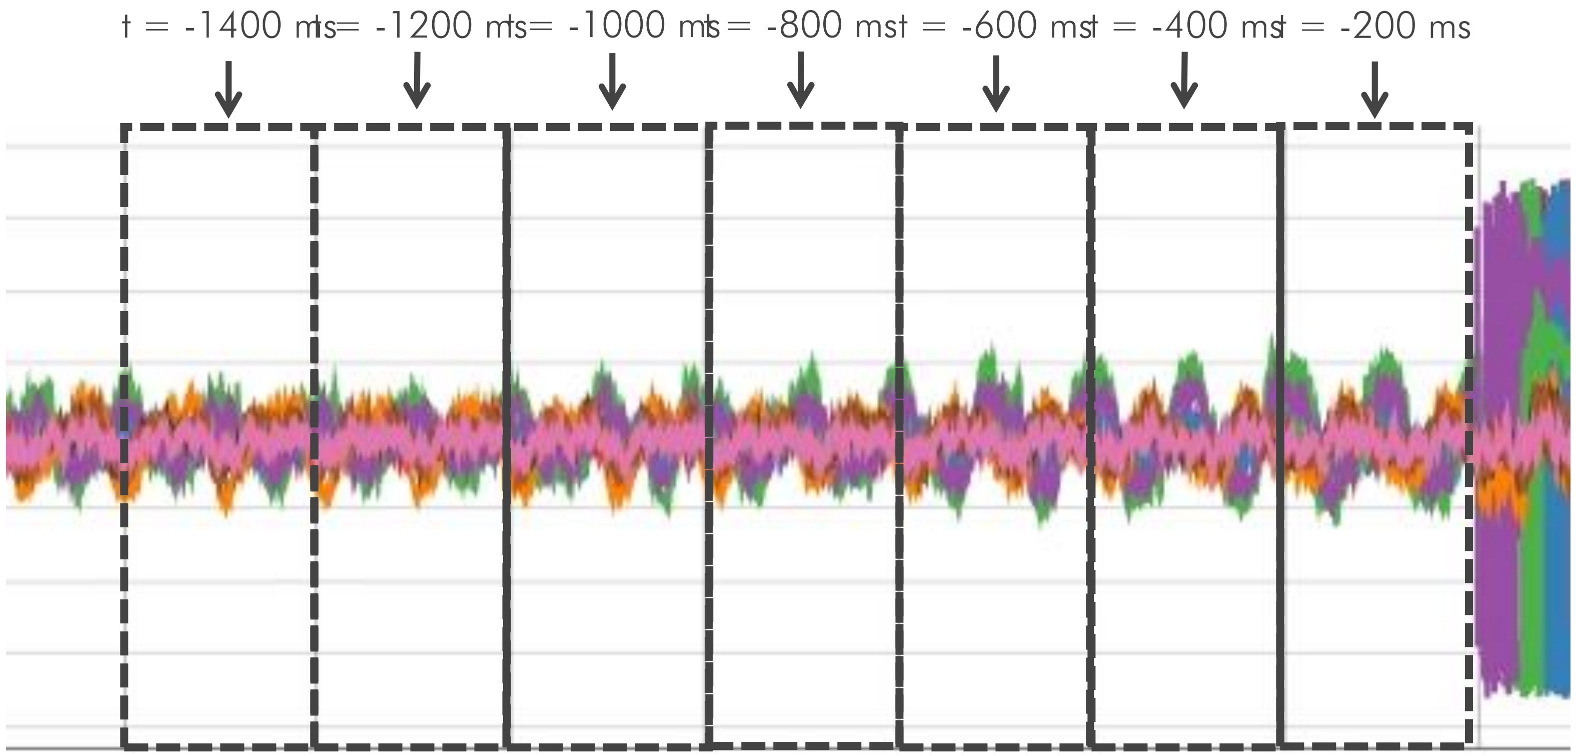

Intermediate Step: Sliding Window

- initial results suggests that for some fault types, prediction is possible

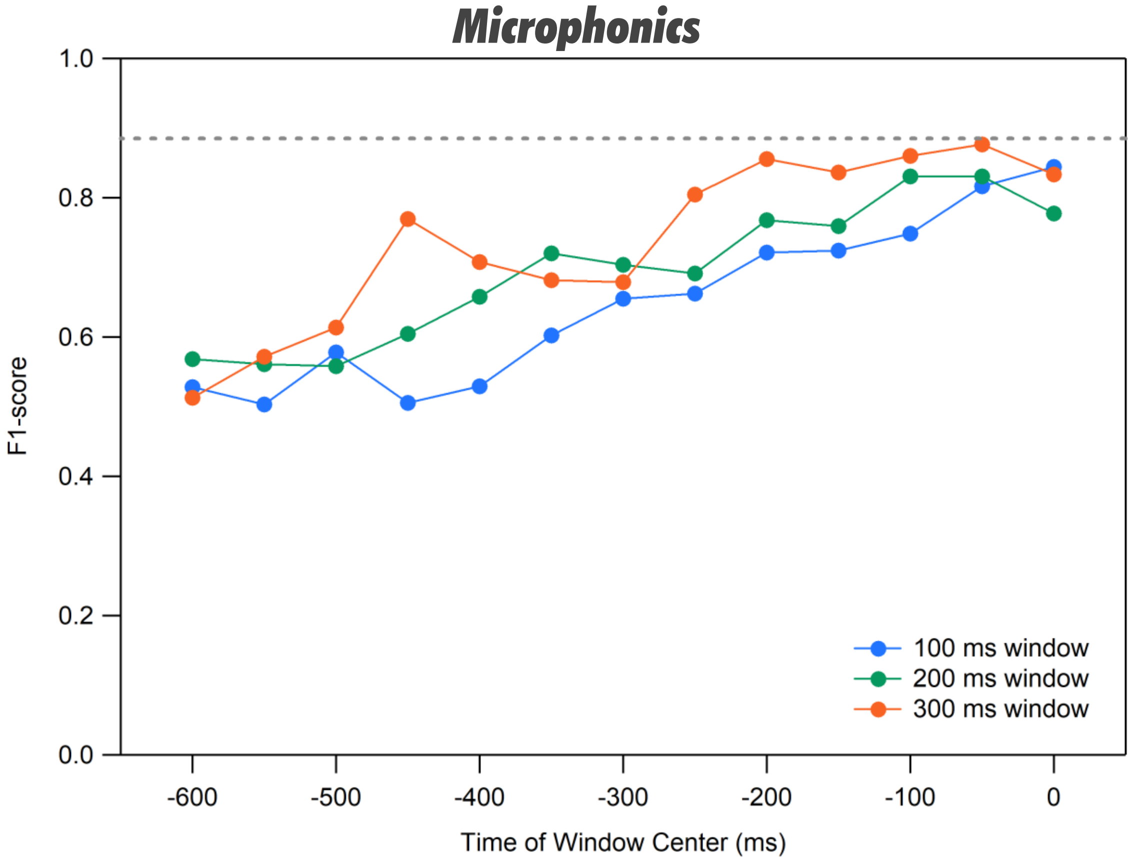

Microphonics

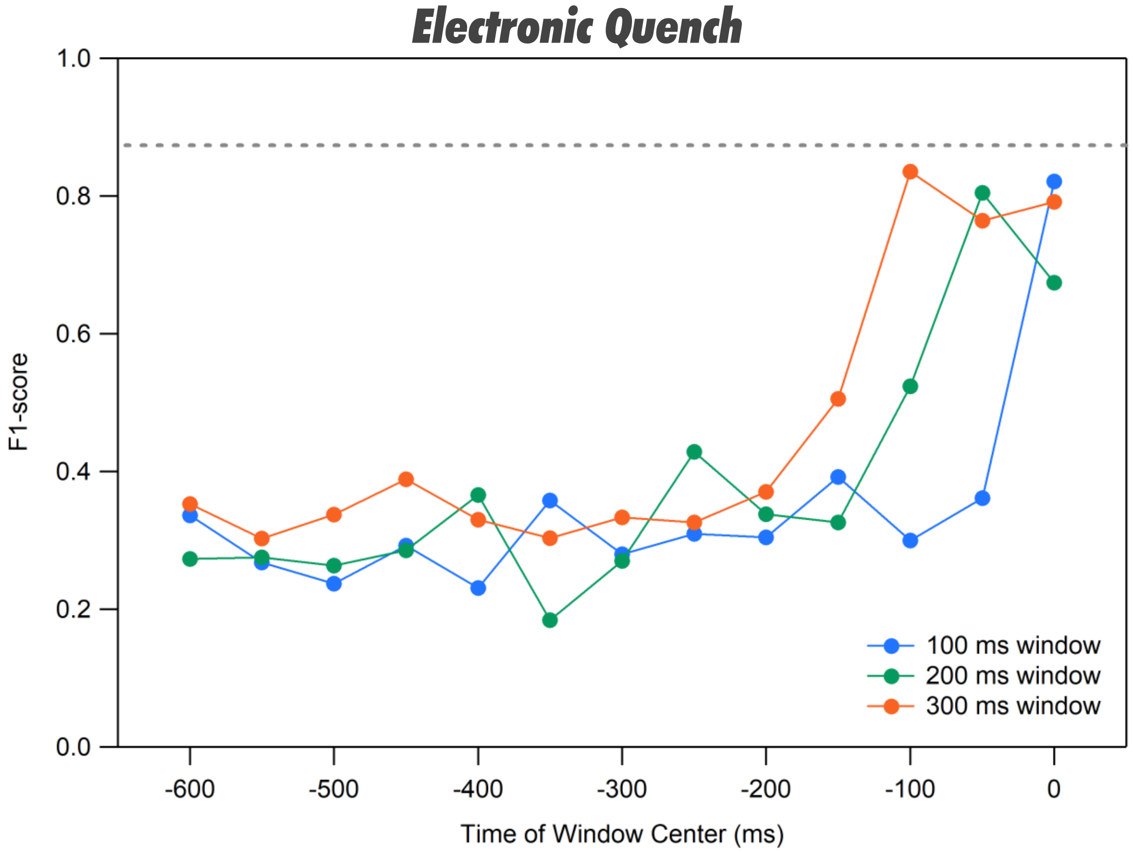

Electronic Quench

- motivates continued study
  - ✓what kind of targeted mitigations could be implemented in those time-scales?

Data: Fueling AI

Detection → ***Isolation*** → ***Identification*** → ***Prediction***

beam off

high sample frequency "snapshots"

high sample frequency streaming data

- commensurate increase in data fidelity required

SRF cavity instability detection in legacy cryomodules

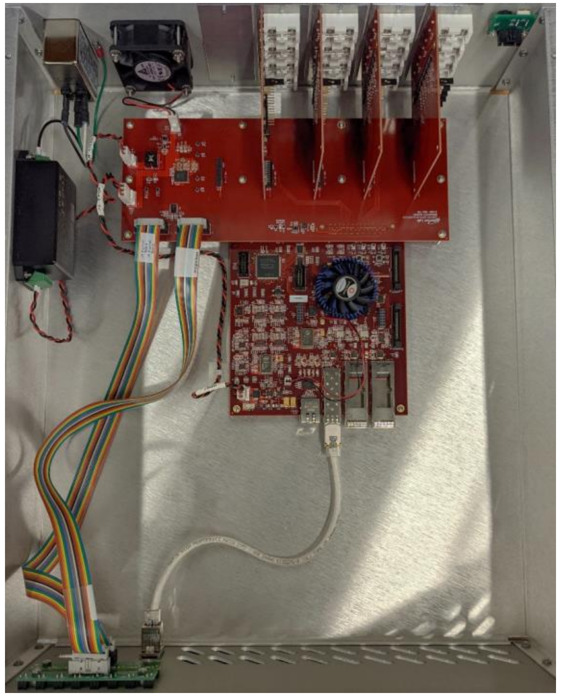

prototype DAQ for legacy CEBAF cryomodules

field emission management

JLab designed radiation detector

Cavity Instability Detection

Problem

- ✓SRF cavities can become unstable without presenting faults
- ✓cavity instability causes beam energy instability, which can lead to beam loss and limited availability of beam for experiments
- ✓identifying an unstable SRF cavity with the present diagnostics at CEBAF is difficult and time-consuming
  - *present diagnostics for the legacy cavities are not fast enough to record fast transient instabilities*

Solution

- ✓develop and install a new fast DAQ system for the legacy SRF cavities
- ✓apply AI to the data acquired by the new DAQ to identify unstable cavities
- ✓the goal is to quickly identify misbehaving cavities and therefore improve beam quality and availability

Cavity Instability Detection

RF Analyzer Tool

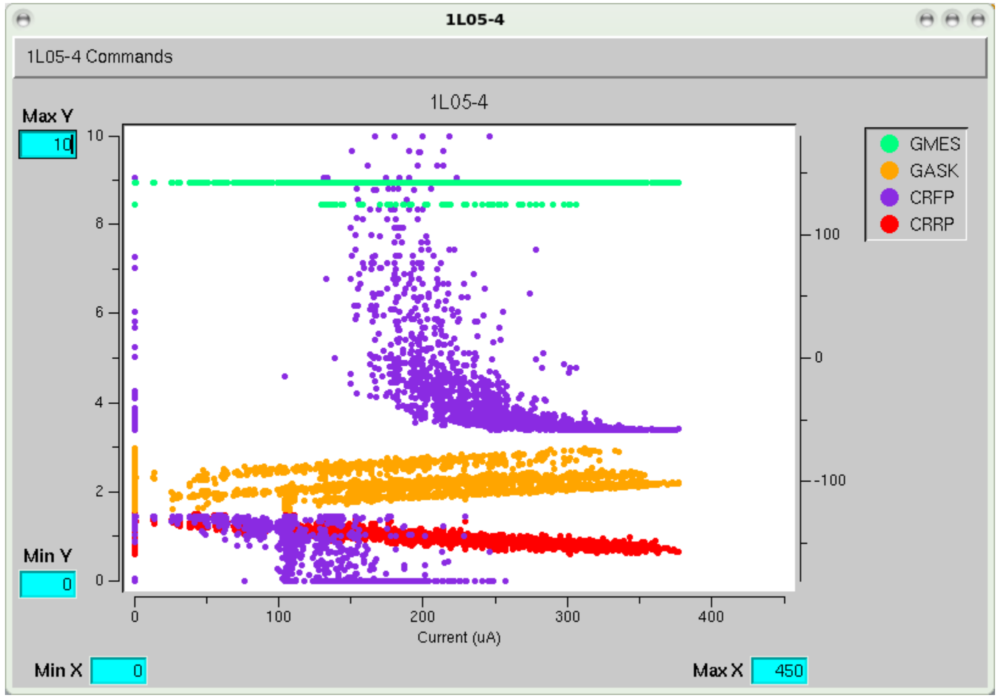

Max Y 10

Min Y 0

Min X 0

Max X 450

- note, this represents an obvious example
- not all instances are so easily detectable by an operator

Field Emission Management

- Goal:

  Maintain low levels of field emitted (FE) radiation without invasive interruptions to physics and prevent damage to beamline components

- Description:

  Use machine learning models – trained on data acquired with newly installed radiation monitors – to model radiation levels, identify cavities that are the source of excessive FE and/or cavities where field emission onsets have changed

radiation area

damaged beamline valve

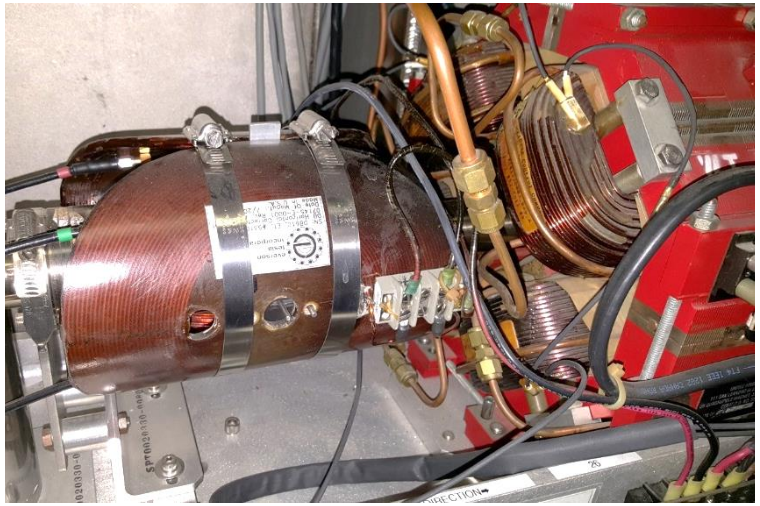

damaged magnet and cables

Field Emission Management

Problem

- ✓field emission is a notorious problem resulting component damage, trips, activation, etc.
- ✓a single cavity produces field emitted electrons with a non-linear response to gradient above a threshold (FE onset)
  - *these may change over time due to various factors*
- ✓FE electrons can have complicated interactions with neighboring cavities and/or cryomodules and can be transported substantial distances up or downstream

Solution

- ✓use machine learning models to help manage this radiation problem non-invasively
  - *can we model radiation levels given an RF configuration (GSETs, etc.)?*
  - *can we identify cavities that are the source of lots FE-related radiation?*
  - *can we identify cavities with changed radiation onset thresholds?*
  - *can we identify new field emitters and localize them in a linac?*

Field Emission Management: Data Requirements

- Jefferson Lab designed, installed, and commissioned a new neutron and gamma radiation detection system focused on FE radiation
  - ✓operational August 2021
  - ✓measure neutron dose rates correctly in the presence of photon radiation
  - ✓detectors are "blind" to low energy photons and electrons
  - ✓integrated into EPICS with signals for gamma and neutron dose rates
  - ✓wide dynamic range
  - ✓currently have 21 detectors installed

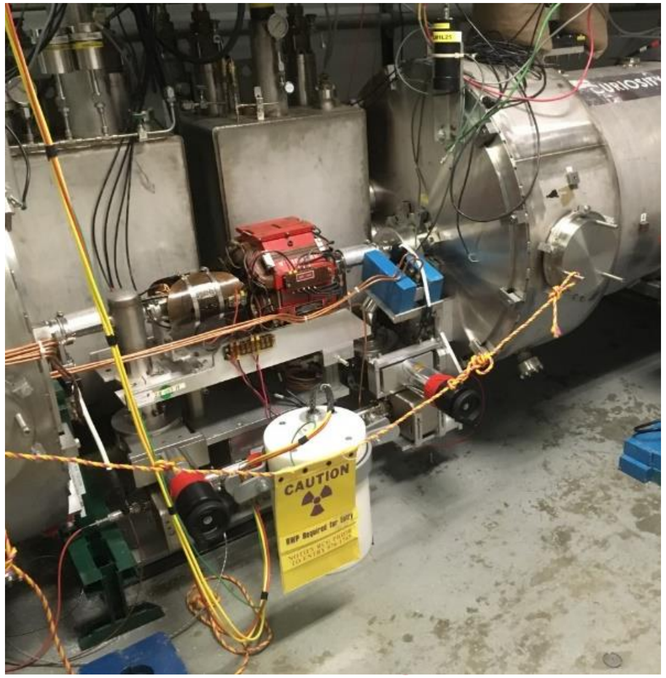

Deep Learning Model

- develop deep learning models that *do not rely on feature engineering*
  - ✓ getting similar performance as ML model

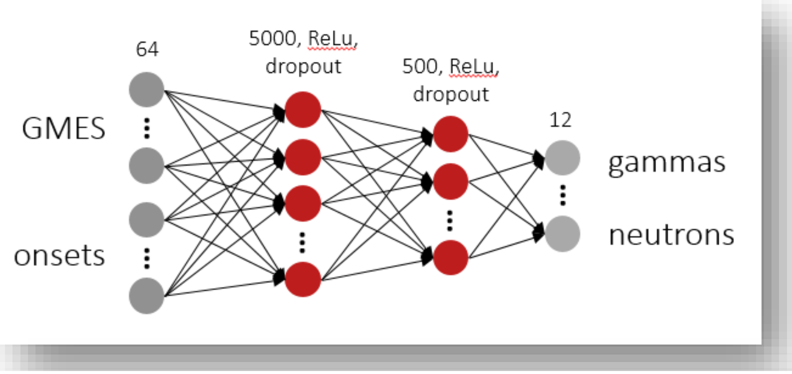

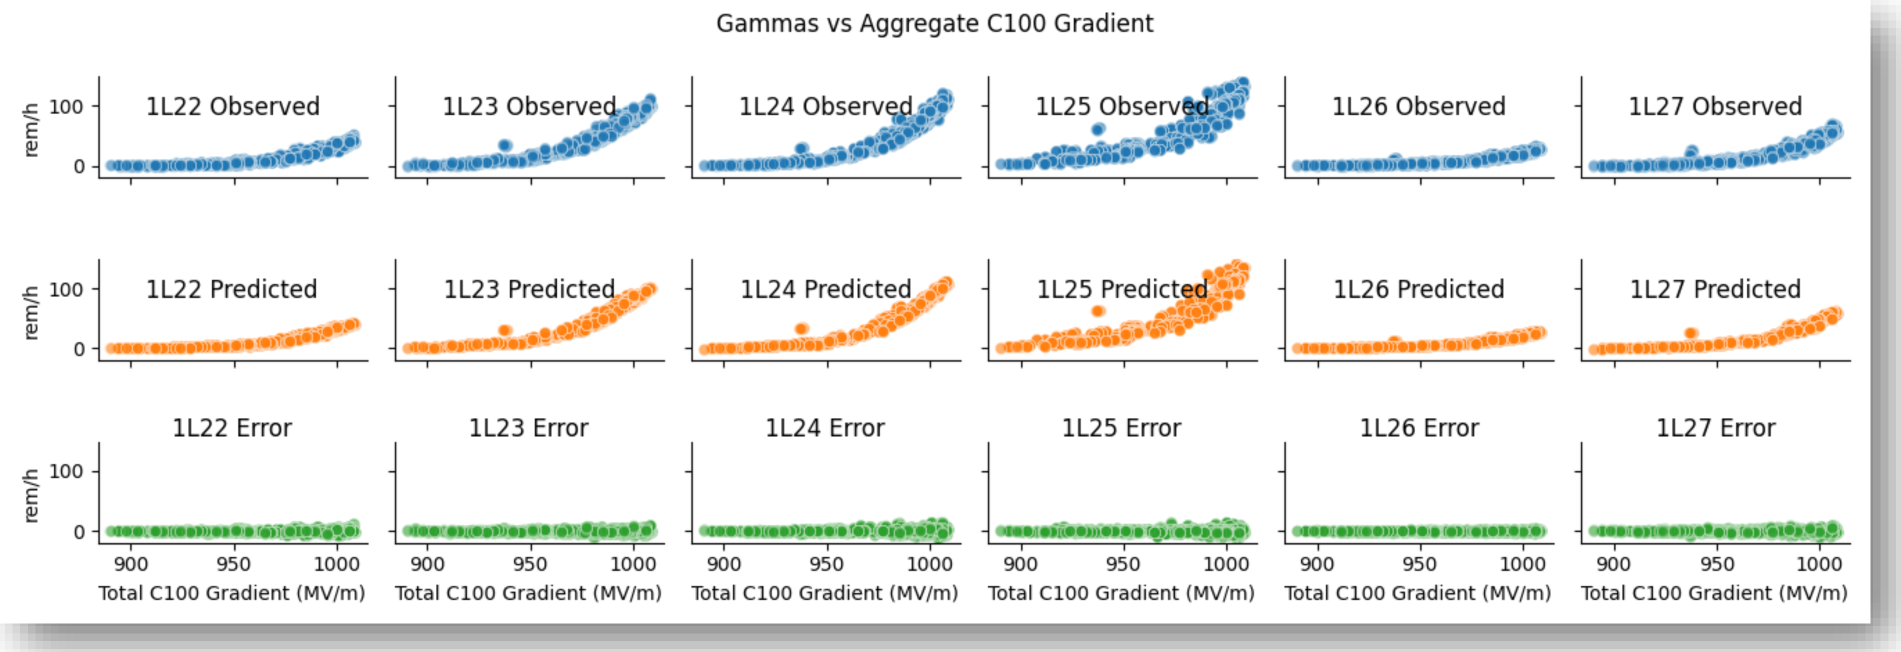

Summary

- detecting, localizing (isolation) and classifying (identification) represent areas ripe for AI/ML application
- the transition to fault prediction represents an ultimate goal
- cannot overemphasize the importance of access to information-rich data
  - ✓higher fidelity data is needed as you move along the spectrum from detection to isolation to identification to prediction
  - ✓to achieve good performance, in addition to higher fidelity data, may also need additional and/or different data

Thank You.

tennant@jlab.org