Element Aero

HARNESSING DATA SCIENCE FOR THE *INFORMED* CONTROL OF SYSTEMS

SANDRA G. BIEDRON, PH.D.

ELEMENT AERO

ICALEPS 2019, BROOKLYN, NEW YORK

6 OCTOBER 2019

ELEMENT AERO

THANKS FOR THE INVITATION!

- Thanks to Marco Lonza and Manuel Gonzalez Berges for organizing this gathering and the inivtaiton.
- Thanks to Richard Farnsworth and the entire ICALEPS team on the organization on the conference **including support for students** (including mine).

TODAY

Element Aero

- How does an accelerator scientist and engineer arrive at the realization of data science.
- General overview of intelligent methods of examining and using data and reminding ourselves we cannot be decoupled from computing advances.
- Examples of several applications of using intelligent techniques in the physical sciences including in the engineering of systems (e.g. control of scientific facilities including accelerator science and engineering).
- Point the attendee to the many resources in our global community, to help better guide them to their own solutions, including reference materials for recent workshops, resources of past efforts in related fields, and to recent initiatives for ethical design of intelligent systems (e.g. IEEE – see back-up charts).
- Thoughts for the future of how to better weave data science processes into our future scientific systems – specifically particle accelerators and laser facilities and well as these systems' research outputs.

MY STORY

ELEMENT AERO

- Those of you that know me understand that I have a wide variety of interests.
- I have used data science techniques for building simulation codes as well as controlling aspects of accelerators.
- I actually came to use intelligent data science methods in the control of general aviation aircraft and other aerial platforms for security purposes in the post 9/11 era.
- For accelerators and peripheral sub-systems, I recognized that deployed systems and the ever more complex light source systems would require additional attention.
- Having worked on defense and industrial projects since the very beginning of my career, I apply approaches such as model-based systems engineering. The "model" is increasingly important and is not steady state.

Ship-board FELs for defense

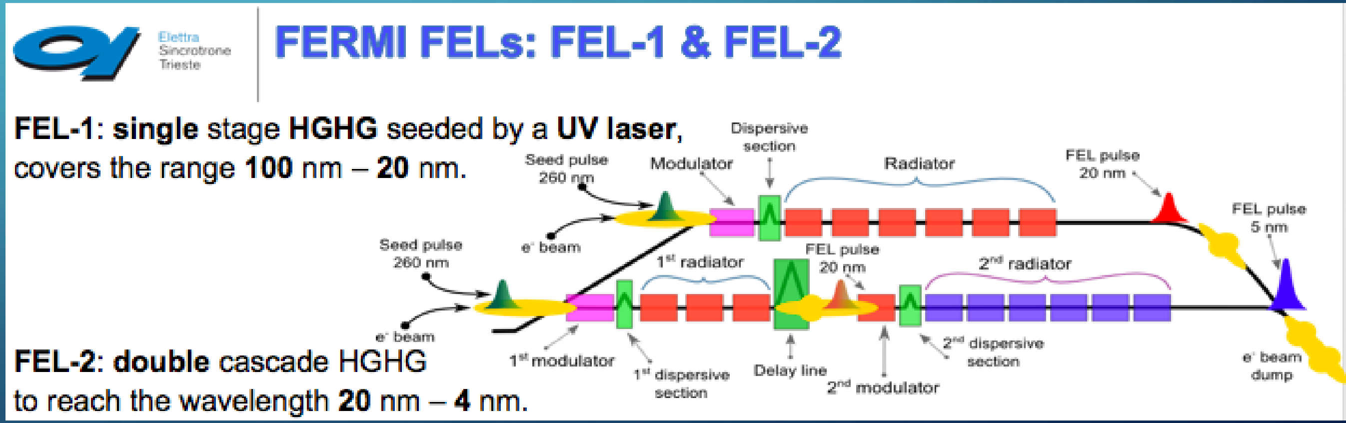

User facilities requiring high quality, repeatable performance for several thousands of hours per year.

DIGRESSION FOR A MOMENT TO THE CONCEPT OF A MODEL

MODELS

ELEMENT AERO

- Models of systems can help provide insight into the dependencies and interfaces between the various subsystems.
- In systems engineering practices, ***models reign***.
  - One definition - "**Model-Based Engineering (MBE): An approach to engineering that uses models as an integral part of the technical baseline that includes the requirements, analysis, design, implementation, and verification of a capability, system, and/or product throughout the acquisition life cycle.**" (Source Final Report, Model-Based Engineering Subcommittee, National Defense Industrial Association, Feb. 2011).
  - This definition does not capture the model being used after a system is delivered but the model can be used later – like being updated constantly for refining the understanding of the system and being used to monitor its health as well as controlling it.
- If I just looked at the FEL for defense, I saw clearly that a more comprehensive model was needed that could help better design and understand the system *as well as for health monitoring of the system and control the system* out to the use case.
- Data science helps build off of our first principles models, help systems that might not have first principles models, and help understand systems that might not have much experimental data, etc.

"All models are wrong, but some are useful" George E. P. Box

but we can to our best try to build the best model possible.

DATA SCIENCE

- *NOT ANY ONE SIMPLE DEFINITION FITS ALL*
- *IS ONE TOOL WE CAN USE*

ELEMENT AERO

Numerical Analysis

- Optimization (including genetic algorithms)
- Interpolation, extrapolation, and regression
- Systems of equations
- Eigenvalue and singular value problems
- Numerical differentiation
- Numerical integration
- Etc.

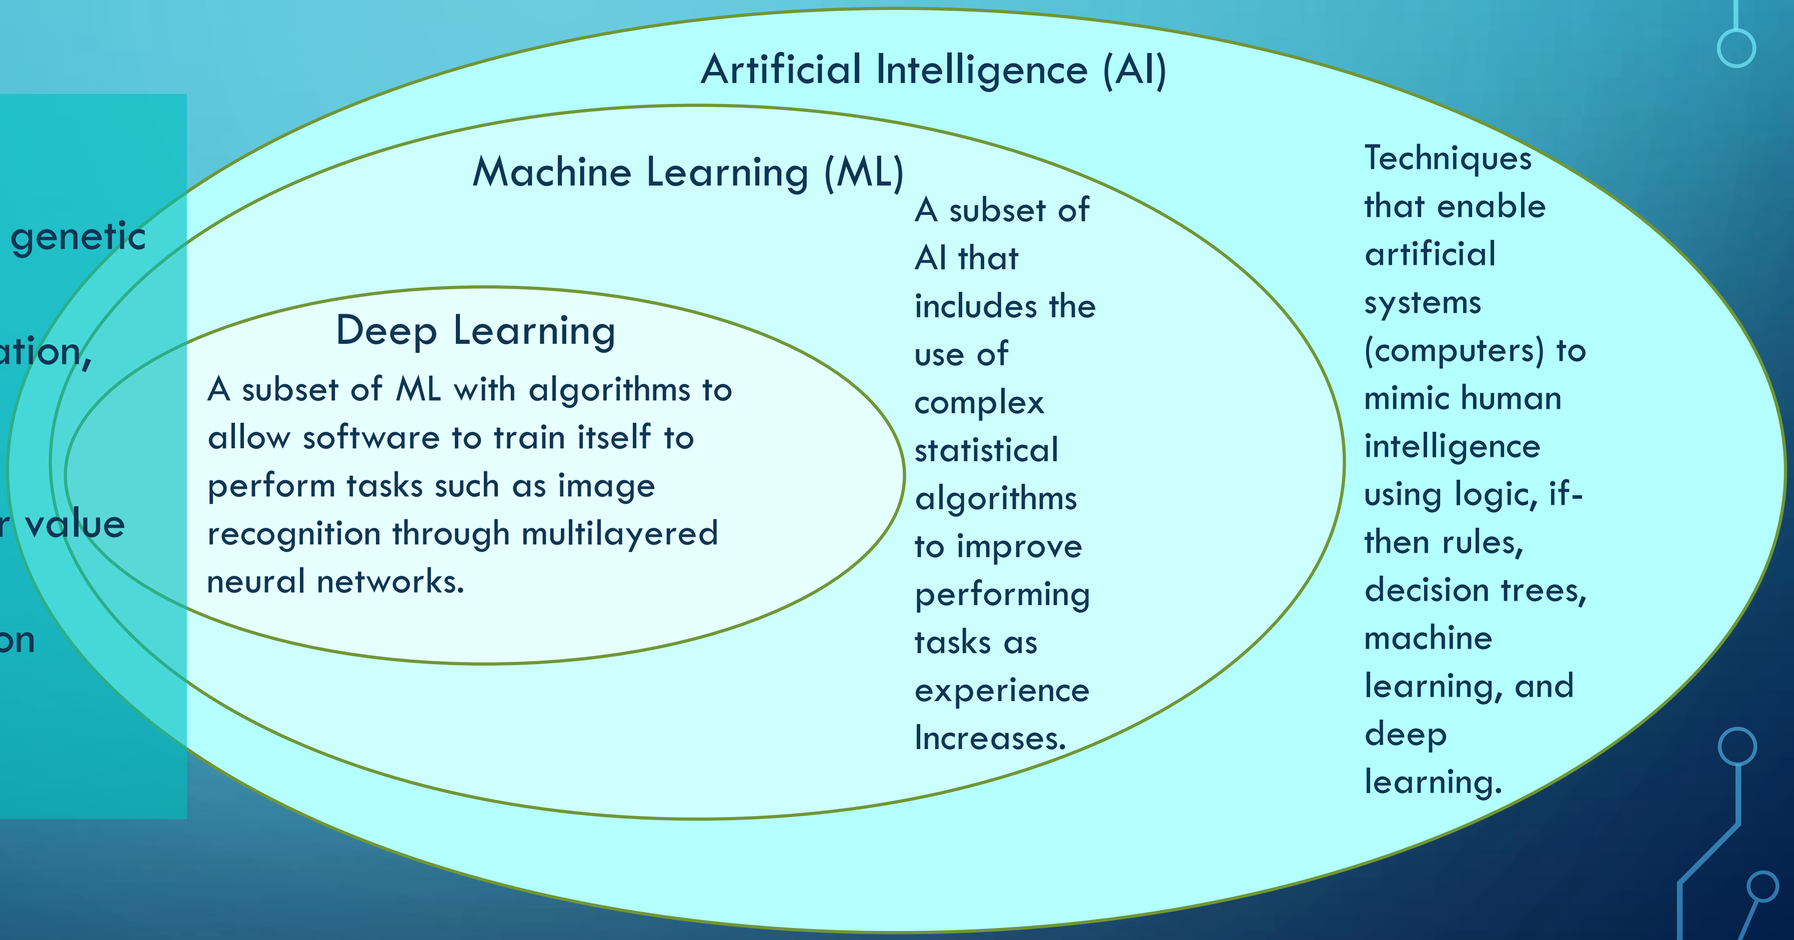

"[Machine Learning is the] field of study that gives computers the ability to learn without being explicitly programmed." - Arthur Samuel in 1959

MACHINE LEARNING

ELEMENT AERO

- Machine learning is a data analysis method that can identify patterns, learns from them and utilizes for making accurate predictions and better decisions without or with minimal human guidance.

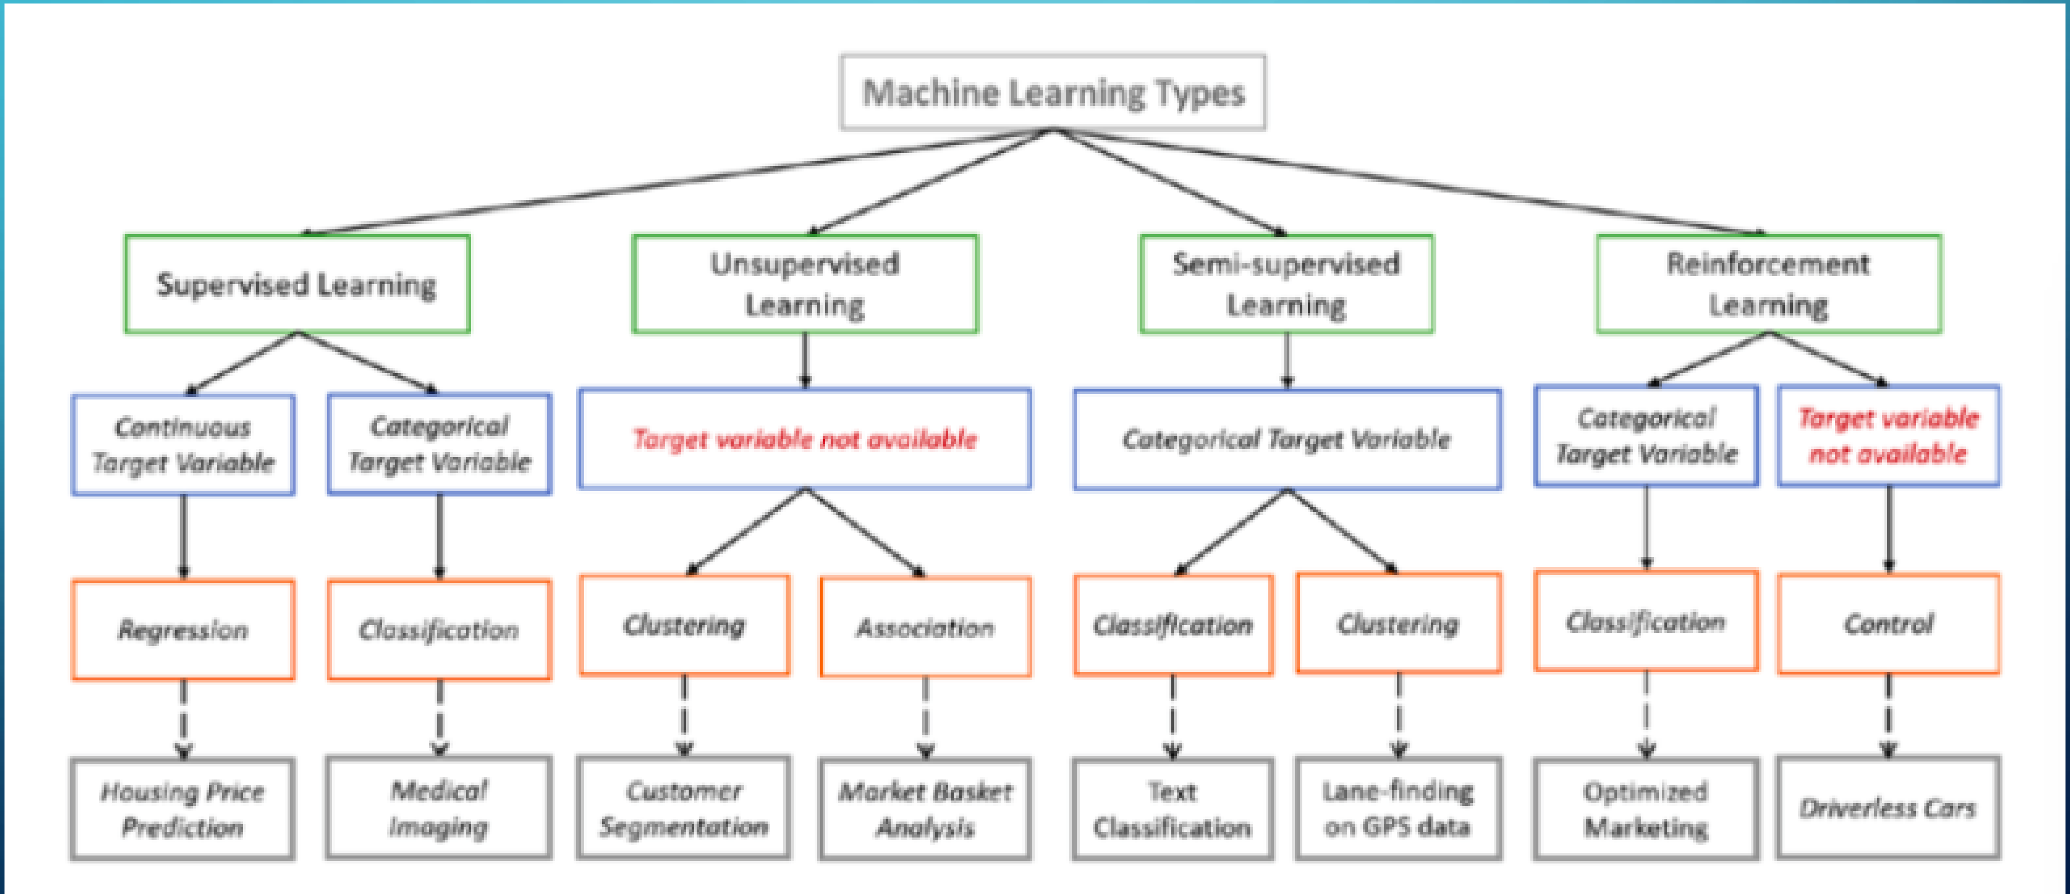

WHAT ELSE MIGHT WE NEED FOR DATA SCIENCE?

- Knowledge of where the data originates (in data science terms "its domain")
- Exploratory data analysis
- Mathematics
- Statistics
- Visualization
- High performance computing
- Specialized computer hardware – FPGAs, application-specific integrated circuit (ASIC), including tensor processing units (TPUs), neural engines
- Etc…

KEEP IN MIND!

- One of the things to keep in mind when thinking about which algorithms to use is that the people that developed a certain algorithm were concerned about making a better algorithm and not about your exact data set. The algorithm might have been developed on random data.
- This is one reason I mention later about technically diverse teams.

BACK TO MODELS FOR A MOMENT

Element Aero

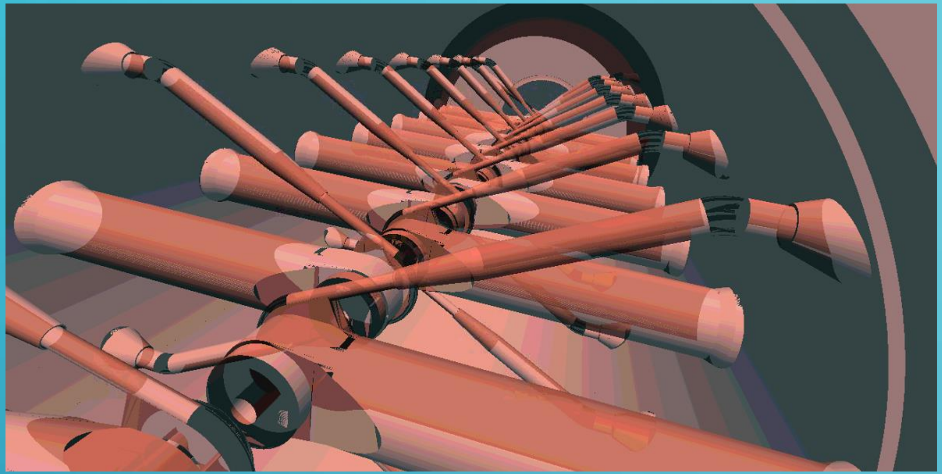

MODEL

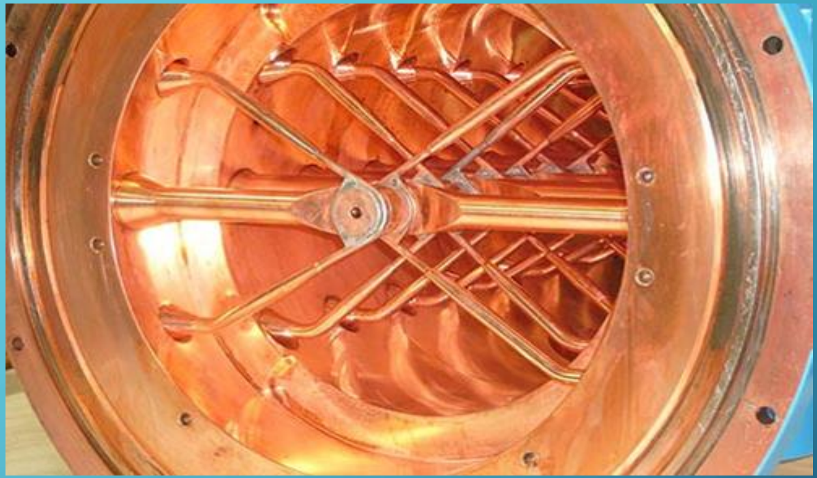

DEVICE (SYSTEM)

TECH-X

Tech-X's VSim engineering and physics models courtesy of Salvador Isaac Guitron Sosa

BACK TO MODELS FOR A MOMENT

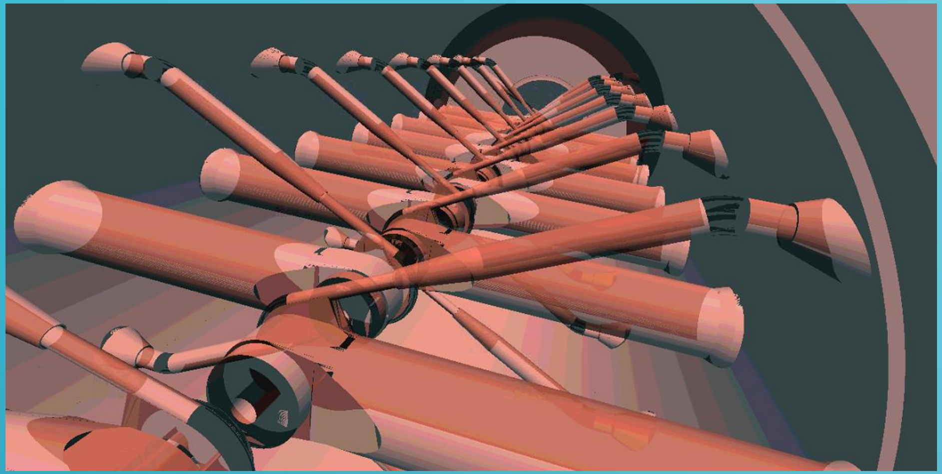

MODEL

First principles, design to requirements, etc.

GOAL

is to make

identical

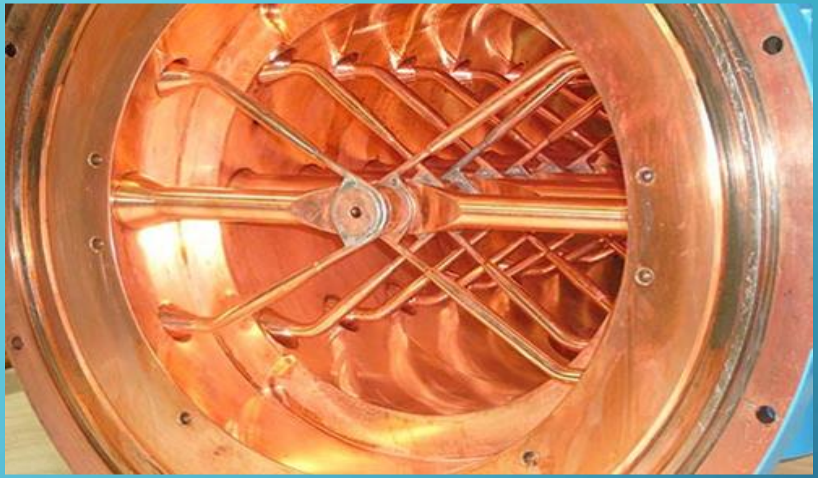

DEVICE (SYSTEM)

Operation, controls, maintenance records, etc.

HPC

DATA SCIENCE

There can be a data exchange process that begins from the first principles simulations and eventual model of a device with the behaviour and control of the device if well monitored can improve the model continuously through data science.

TECH-X

Tech-X's VSim engineering and physics models courtesy of Salvador Sosa

WHERE CAN DATA SCIENCE BE APPLIED

AN EXAMPLE FROM THE PHYSICAL SCIENCES

- *High-Energy Physics*

HIGH ENERGY PHYSICS

- Employing AI/ML algorithms has demonstrably sped up calculations (e.g. one million times for analysis of strong lenses) and produced clear cost-savings (30% increase in effective detector volume for neutrino flavor tagging with neural networks).
- These faster algorithms enable accelerated science, and in some cases they lead to paradigm shifts: in situations where it once required one day for a human to analyze a single object.
- Also see a recent talk/work by Gabriel Perdue (FNAL) -> next chart

- Y. D. Hezaveh, L. P. Levasseur, and P. J. Marshall, “Fast automated analysis of strong gravitational lenses with convolutional neural networks,” Nature, vol. 548, pp. 555–557, Aug. 2017; P. Adamson et al., “Constraints on oscillation parameters from ne appearance and nμ disappearance in nova,” Phys. Rev. Lett., vol. 118, p. 231801, Jun 2017. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett. 118.231801; M. A. Acero et al., “New constraints on oscillation parameters from ne appearance and nμ disappearance in the nova experiment,” Phys. Rev. D, vol. 98, p. 032012, Aug 2018. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevD.98.032012

Courtesy Jim Amundson

Why deep learning is interesting for neutrinos in one slide...

Because, just *look* at them!

Event reconstruction

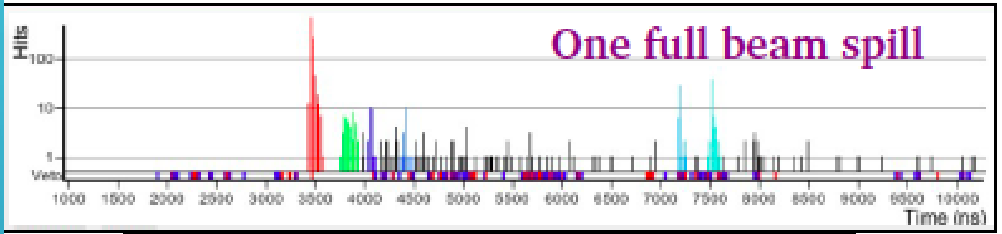

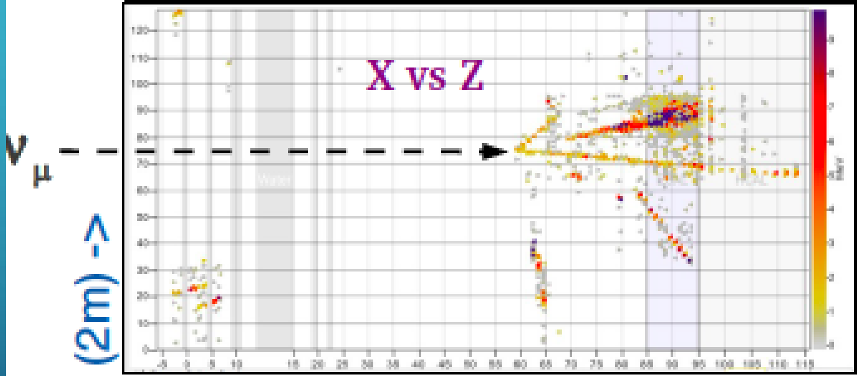

Position (5m) ->

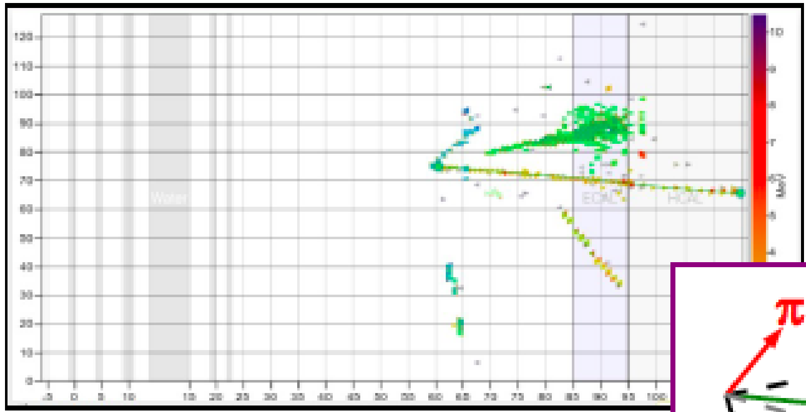

My guess, just for fun

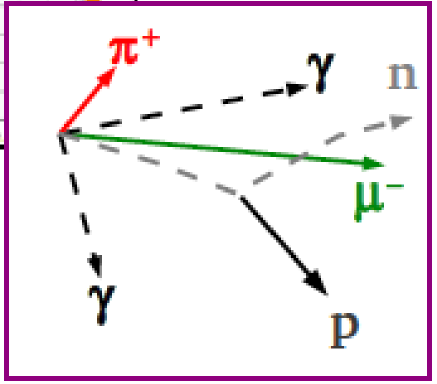

Modern neutrino detectors are imaging detectors. Domain experts can discern an enormous amount of information visually. But feature engineering is hard.

Fermilab

Courtesy Jim Amundson

AN EXAMPLE OF A SYSTEM

- *Example of a life saving system with humans in the loop*
- *Fire segmentation in thermal imaging for situational awareness*

Motivation

- Fire fighters carry sensors such as thermal imaging cameras
- There is a need for creating a situational awareness system for first responder crews
- Automatic detection of fire and smoke is part of this situational awareness - *Fire segmentation* rather than just *fire detection* allows one to **determine the extent of the fire and the smoke**.
- Neural network procedures can fully characterize the fire ground by producing natural language descriptions of the scene.
- The method must be real-time and unsupervised

Courtesy of Manel Martinez-Ramon

Proposed algorithm

Sequence acquisition
(thermal images at 30 fps)

→

Feature Extraction
- Intensity
- Velocity
- Divergence SIFT

→

Segmentation
Comparison: K-means, GMM, GMRF, MRF

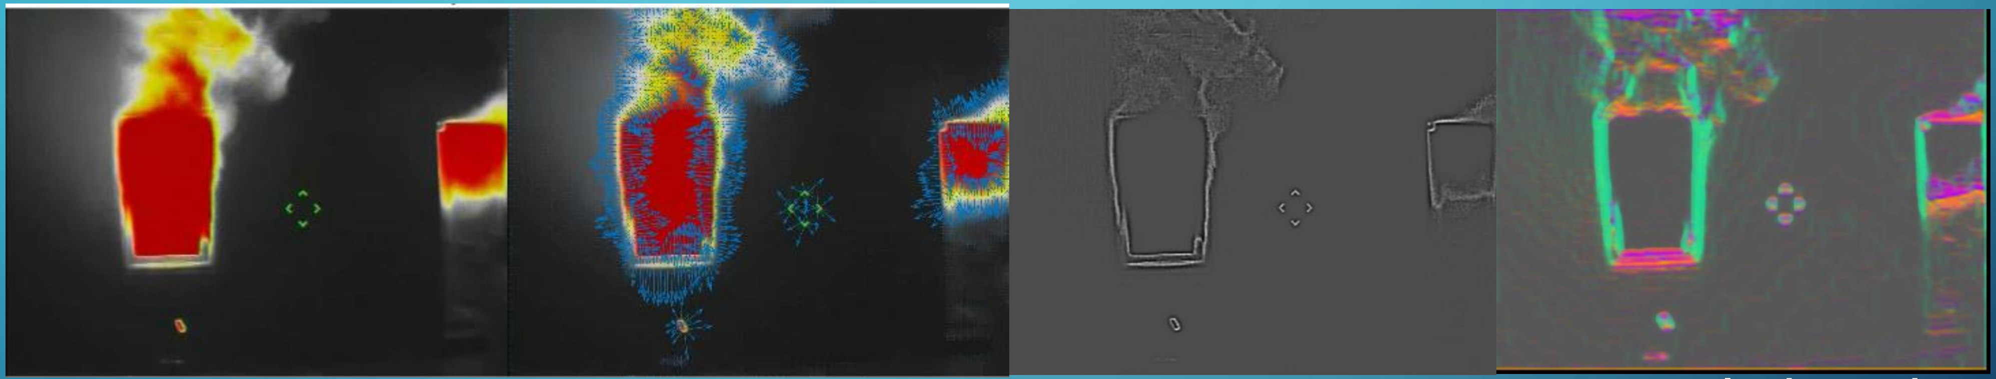

Image | Velocity | Divergence | SIFT - scale invariant feature transform

Features processed by ML: Construction of a probabilistic model of the features
Three conditional probability distributions: background, fire, smoke

Different models are compared. **Markov Random Fields** is the best algorithm

Courtesy of Manel Martinez-Ramon, to be published IEEE Access

Courtesy of Manel Martinez-Ramon, to be published IEEE Access

Results

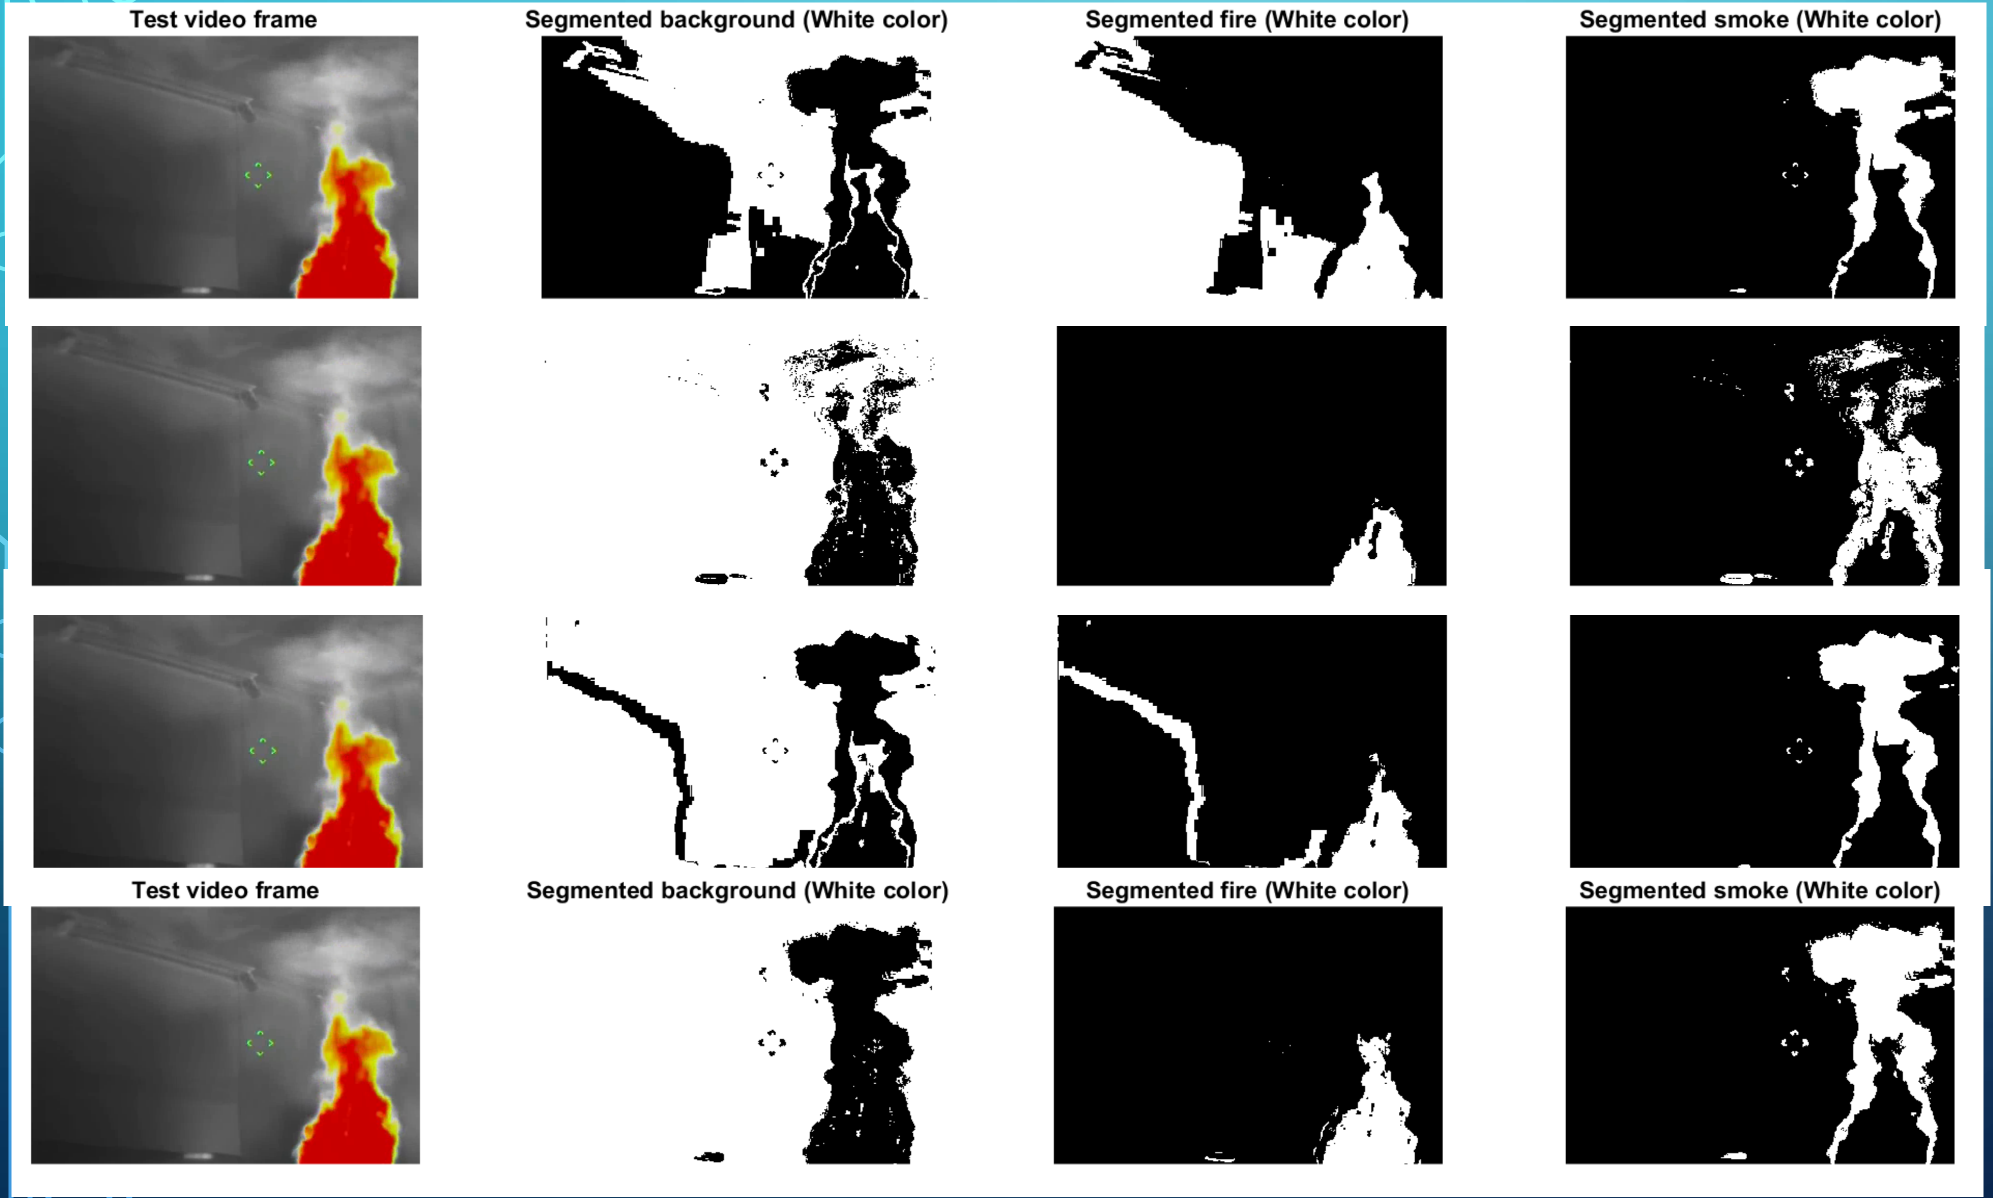

KNN

Gaussian mixture models (GMM)

Gaussian markov Random Fields (GMRF)

General Markov Random Fields (MRF)

Results

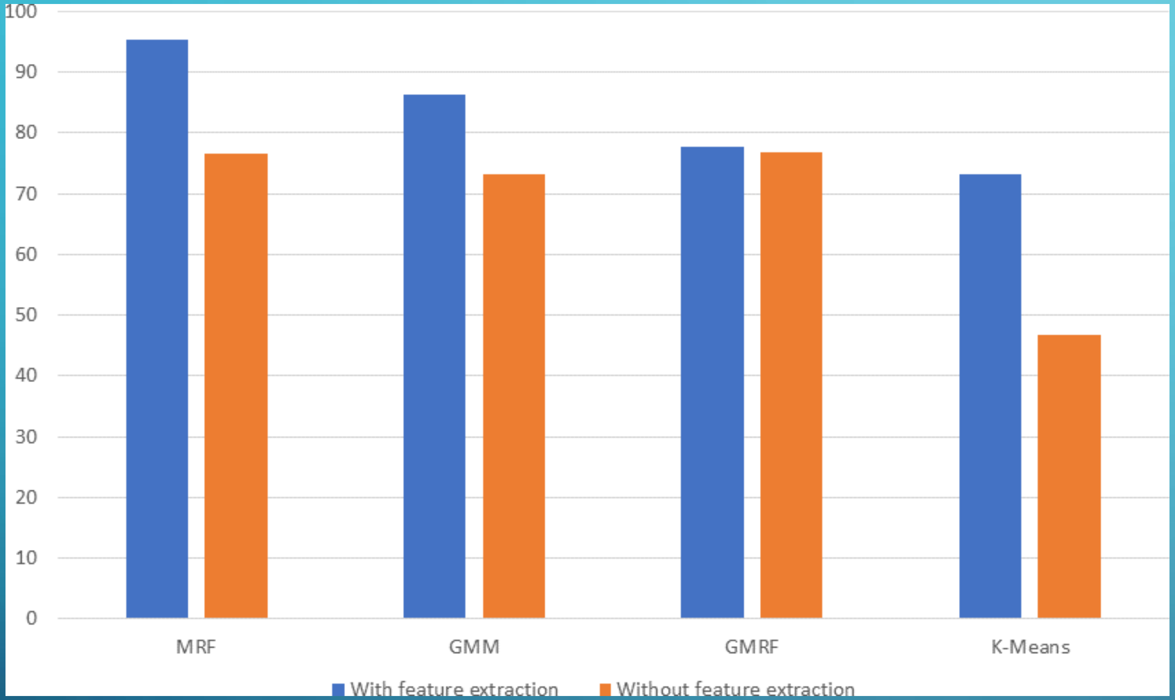

Performance comparison among algorithms

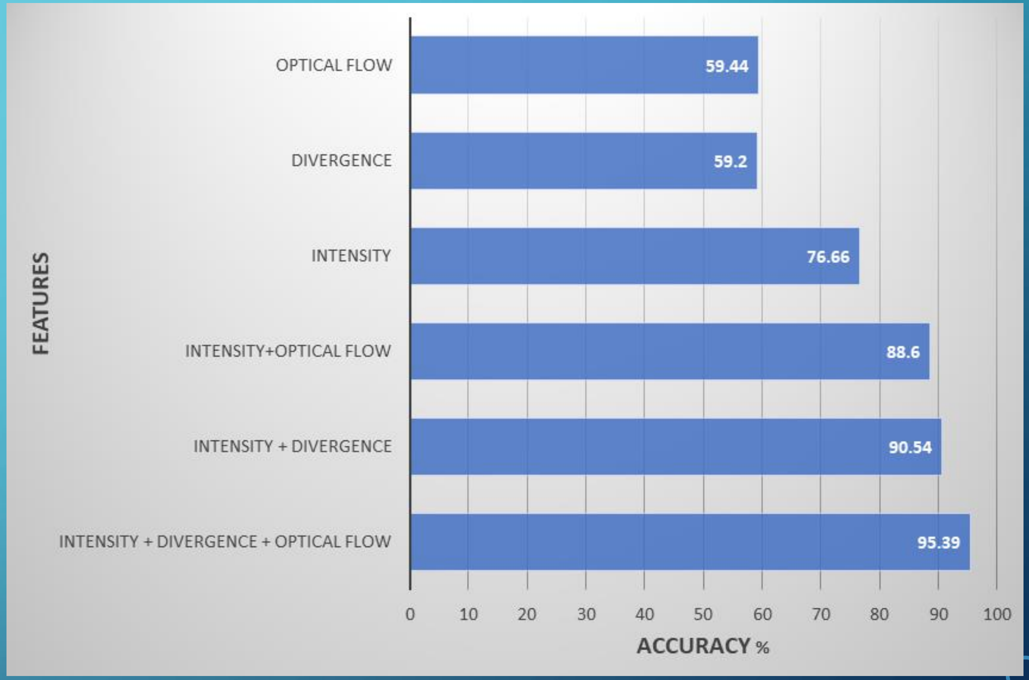

MRF performance with different feature combination

Courtesy of Manel Martinez-Ramon, to be published IEEE Access

SO WHAT ABOUT SCIENTIFIC SYSTEMS SUCH AS PARTICLE ACCELERATORS AS INTELLIGENT SYSTEMS?

- How can we best use data science to better design, control, and understand our machines and make better models.
  - *AI/Intelligence/ML is* ***one way*** *and it is more than just algorithms, it is data, sensors, computing platforms and analysis techniques! (Stated succinctly by J. Amundson, Fermilab)*
  - *Coming back to the point of algorithms usually developed on random data – the sciences, including the machines enabling science, have "more interesting" data sets, including errors in data and embedded physics.*

ACCELERATOR TOOLBOX

PHYSICS, CHEMISTRY, MECHANICAL ENGINEERING, ELECTRICAL ENGINEERING, ETC.

ELEMENT AERO

- Materials science
- Vacuum science
- Collective effects
- Electromagnetism
- Vacuum electronics/RF sources
- Instrumentation
- Controls engineering
- Computational techniques

- Lasers
- Radiation and dosimetry
- Pulsed power
- Survey and alignment
- High performance computing
- Data science, including mathematical optimisation and machine learning
- Etc…

Since we as a community already have such a filled skill set (toolbox), this has led to a slow adaptation of data science as we thought we could solve problems in other ways (already in our toolbox).

ACCELERATOR TOOLBOX

ELEMENT AERO

PHYSICS, CHEMISTRY, MECHANICAL ENGINEERIN... TC.

- Materials science
- V...
- C...
- Ele...
- Vac... sour...
- Instru...
- Contr...
- Compu... techniq...

Since we a... a filled skill set (toolbox), this may have led to a slow adaptation ... us we thought we could solve problems in other ways (already in our toolbox).

One of the hard parts of examining data is that one needs to understand where the data comes from and what it means. **The data is not disconnected from the system from which it originates.**

This makes choosing the right data science tools for each accelerator "challenge" difficult. Why? Our toolbox of systems is so large and so integrated/overlapping so the data science techniques have to be chosen very carefully.

This means that a data scientist cannot just walk in off the street from a different field. **The solutions we need are accelerator-domain specific, complicated, chosen by experience, and are in no way a single, magic band-aid.**

As was stated over and over by the community last week at the Department of Energy's INNOVATION XLAB Artifical Intelligence Summit in Chicago – NOT ONE PERSON CAN MAKE ML/AI USEFUL - WE ADD DATA SCIENTISTS TO EXPERT TEAMS AND EXPERTS TO DATA SCIENCE TEAMS TO SEE THE BENEFIT.

INTELLIGENT CONTROLS FOR MEDICAL ACCELERATORS

- ILS has developed a medical linac for **cancer therapy** based on the Neutron Capture nuclear interaction:
  - Cancerous cells are tagged and then radiated with a beam of (slow) epithermal neutrons.
- The NC process releases a large radiation dose within a micrometer range, i.e. **no radiation damage to healthy cells.**
- NC has a small cross section. A **high-flux of neutrons** is then required **to maximize** the number of **NC interactions**
  - The linac needs to be controlled to produce and hold steady a high-current proton beam for producing neutrons.

NM ENGINEERING

ILS ION LINAC SYSTEMS, INC.

TECH-X

ELEMENT AERO

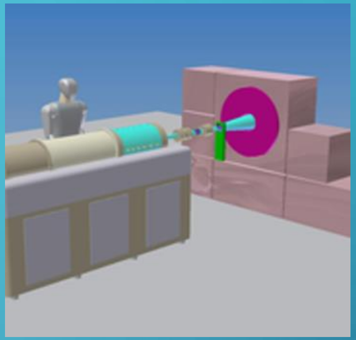

Scheme for BNCT cancer therapy with ILS accelerator.

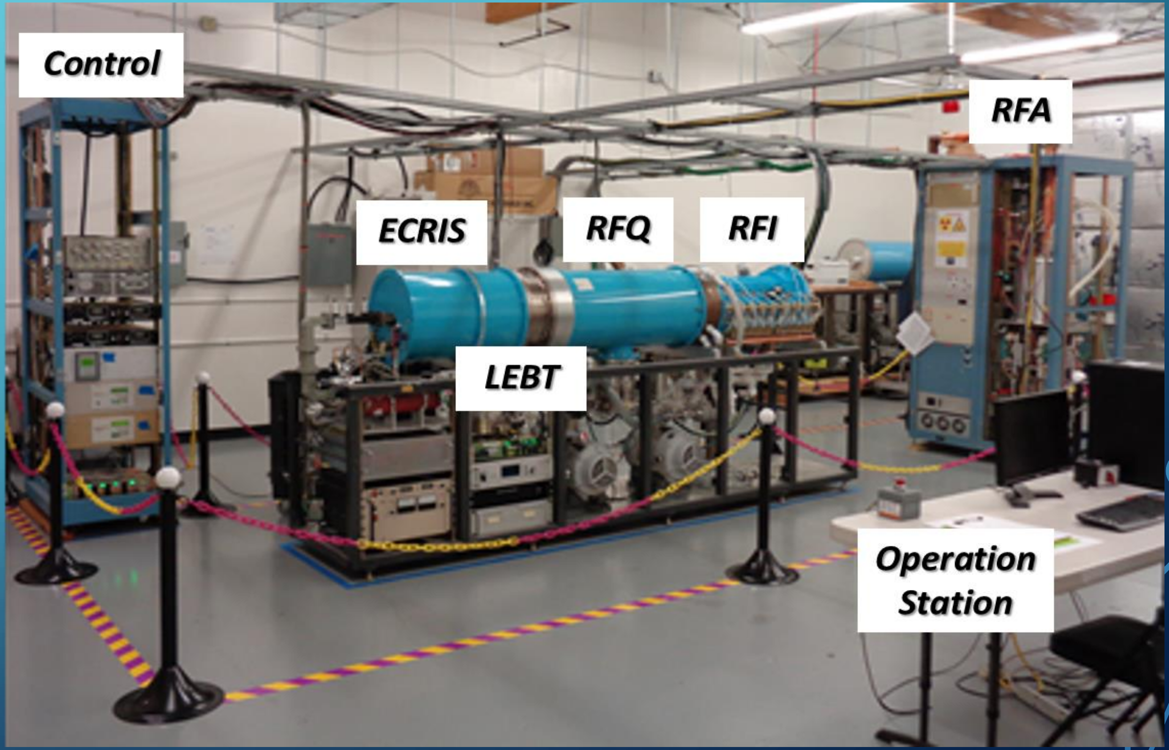

ELEMENT AERO

INTELLIGENT CONTROLS FOR MEDICAL ACCELERATORS

- AI controllers can help the linac operate in an optimal configuration that, as one example, maximizes beam current.

- A Neural Network (NN) can help find the a very good suite of set-points for a set of requirements based on the available diagnostic instrumentation readouts.

- The NN can **adjust the machine online** provided (e.g. pre-trained offline).

- The NN can be **trained** with large amounts of data from two sources:
  - Recorded data from **operations and experiments.**
  - Generated **data from first principles simulations.**

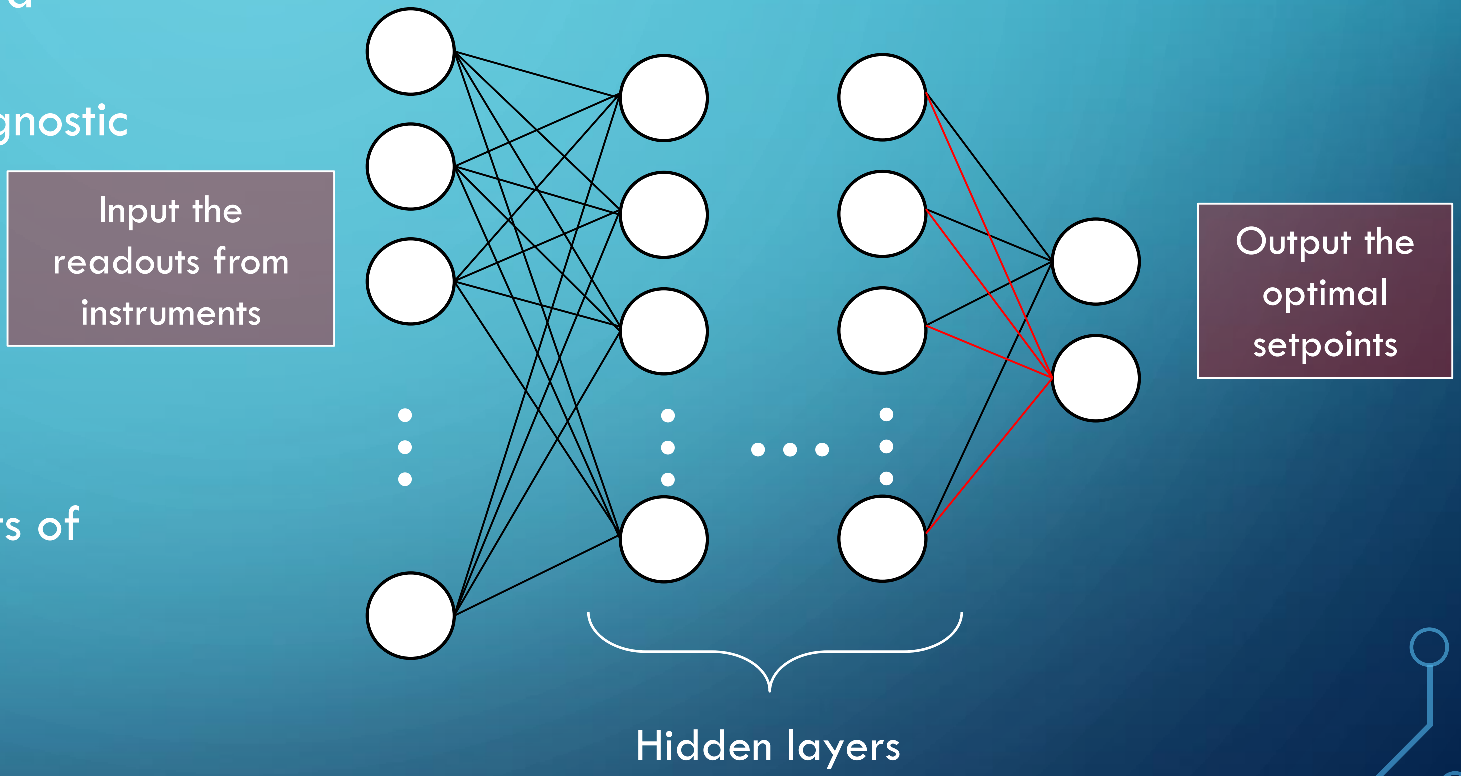

This technology has the potential of making **medical accelerators more affordable and efficient.**

INTELLIGENT CONTROLS FOR MEDICAL ACCELERATORS

- Production of **training data** using electromagnetic and PIC simulation software ***VSim10***.
- Scanning the **parameter space** assuming numerous injected beams, which depend on ion source and LEBT configurations:
    - Hydrogen flow
    - Power supplies
    - Water temperature
    - RGA data
    - Vacuum levels, etc.
    - Target material, condition etc.
- Propagate numerous beams and determine **optimal machine states** being used as training data for the NN controller.
- Use of HPC (**THETA at the Argonne Leadership Computing Facility**) in order to generate thousands of training points.

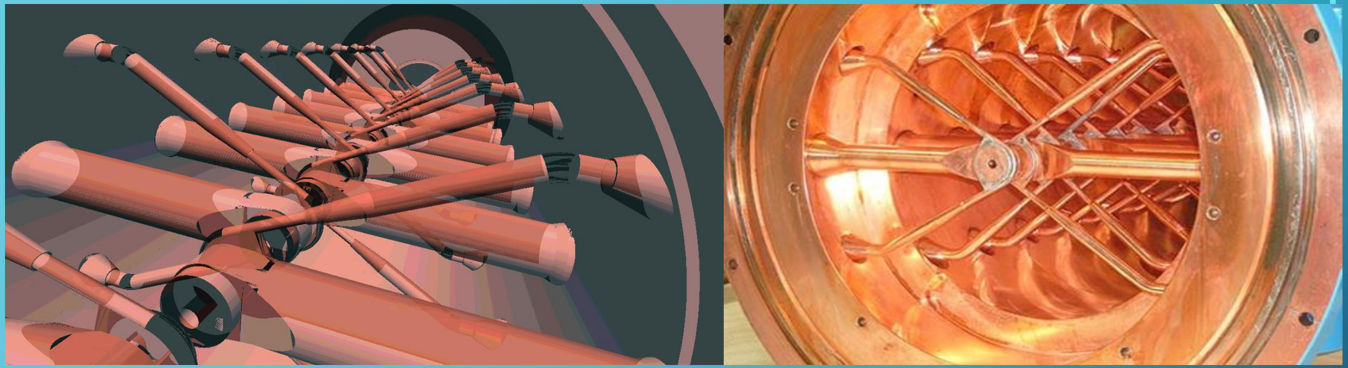

VSim10 model and actual picture of the ILS linac accelerating RF structure.

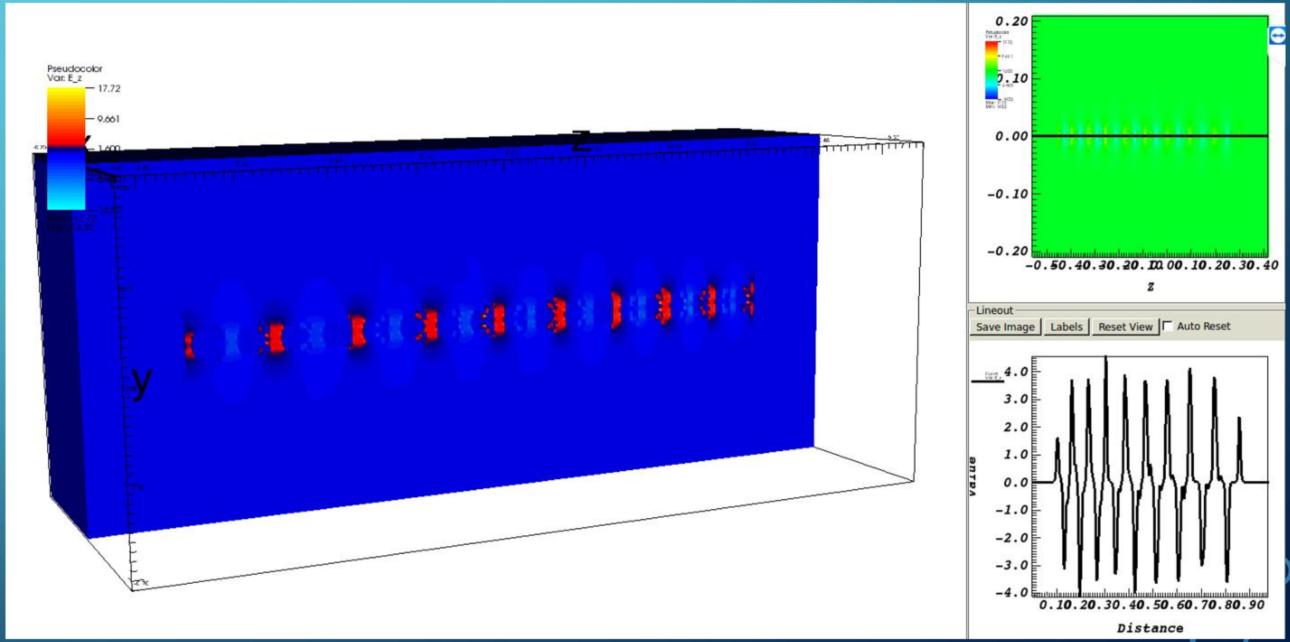

VSim10 simulation of the longitudinal E-field along the axis of the RF structure.

Tech-X's VSim engineering and physics models courtesy of Salvador Sosa

FOR LLRF AND RESONANCE CONTROL SYSTEMS

- LLRF and Resonance control is an important element to achieve desired beam quality.
- Phenomena affecting the field stability include:
  - Microphonics
  - Lorentz force detuning
  - Thermal acoustic oscillations
  - …
  - **All need to be compensated by the LLRF!**
- Effects are both non-linear and difficult to characterize, making **Neural Networks** one of several promising methods to improve existing control.

Courtesy Jorge Alberto Diaz Cruz

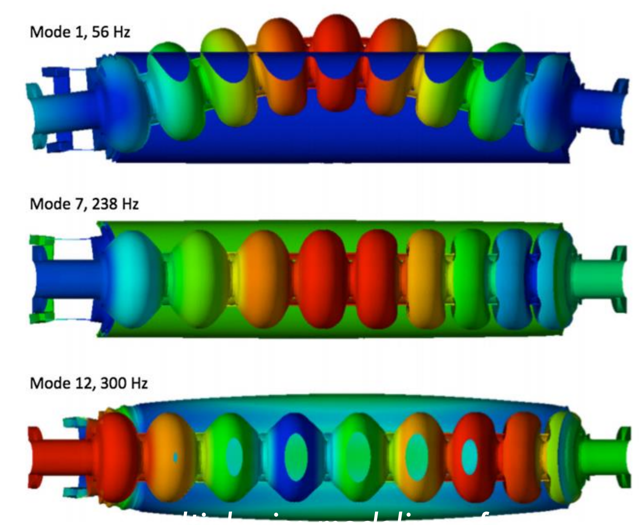

O. Kononenko, *3D multiphysics modeling of superconducting cavities with a massively parallel simulation suite.* PRAB, 2017

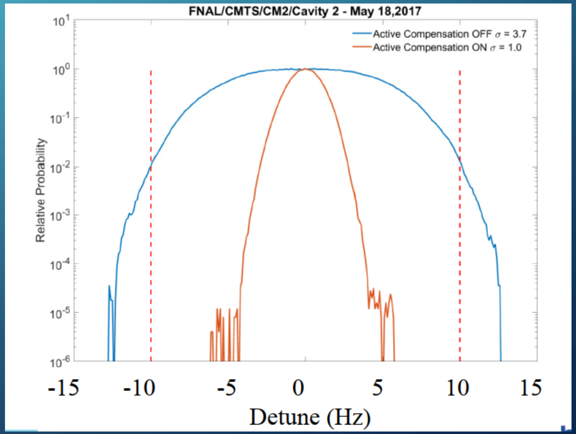

Active Resonance Control developed at Fermilab for TESLA cavities

FOR LLRF AND RESONANCE CONTROL SYSTEMS

- We are exploring intelligent techniques as a tool **to improve performance of existing LLRF** and resonance control systems.
- **Simulations** are performed to create data for training of Neural Networks. (Eventually to merge with experimental data.)
- Gaussian processes, genetic algorithms and Neural networks are all part of the artificial intelligence framework under test.

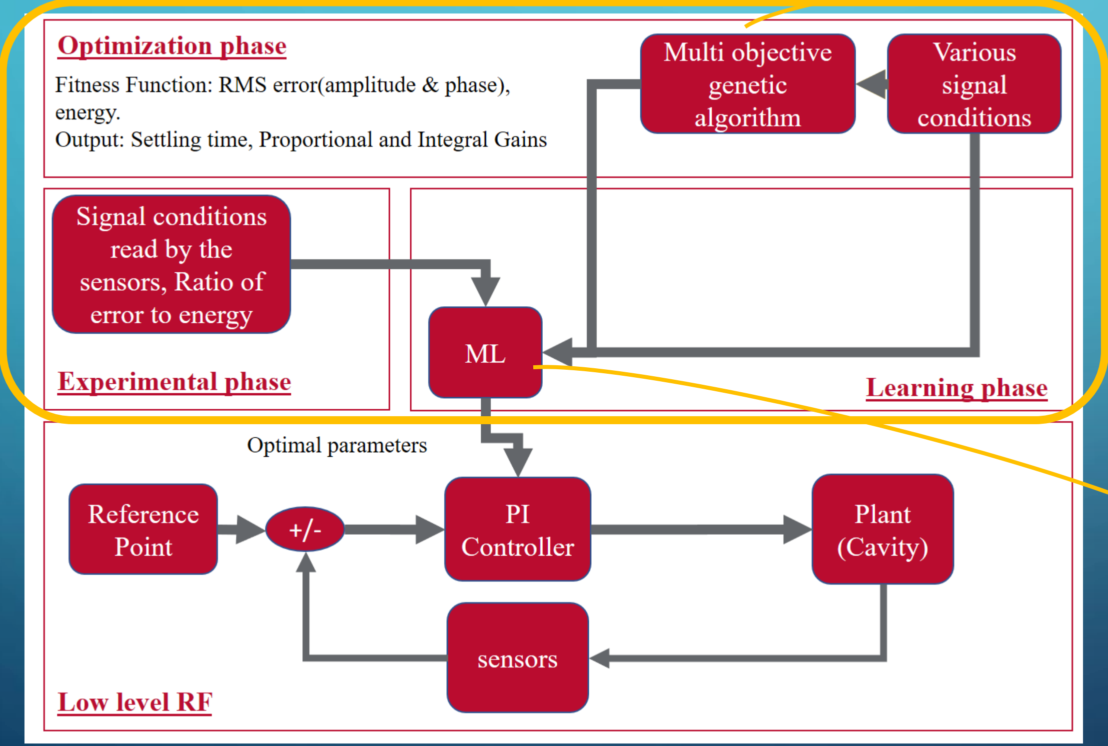

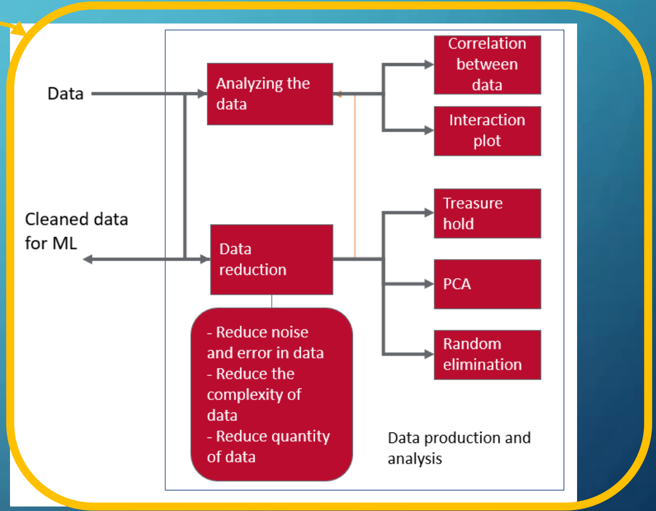

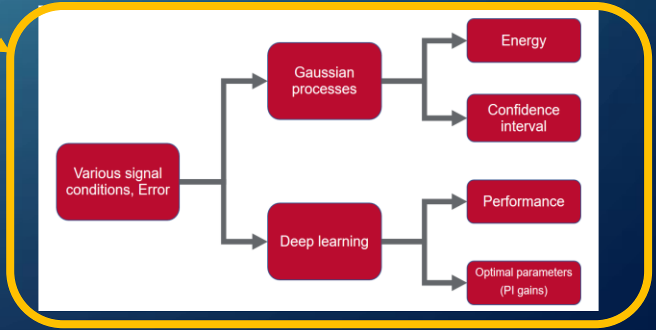

Courtesy Jorge Alberto Diaz Cruz; NAPAC2019, LLRF, ICALEPS

VIRTUAL TELESCOPE FOR X-RAY OBSERVATIONS (VTXO)

- ***A telescope made up of 2 separate satellites in 2 different orbits***
- ***Accuracy requirements are sub-arcsecond relative attitude accuracy and 1 mm relative position accuracy***
- ***Approximately 1 hour observing the Crab Nebula***

Applications

- 1) Investigating the nature of space (Like dark matter)
- 2) Covering the Sun (To see planets and stars behind it)
- 3) Understanding more about black holes
- 4) Detecting space phenomena like star collisions
- 5) Finding other life forms

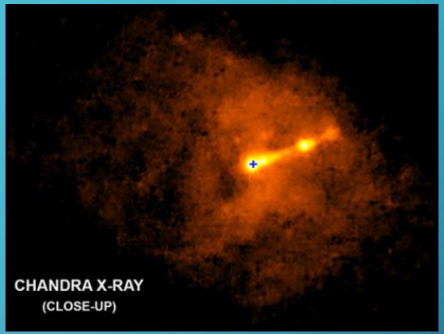

An example of X-Ray observation

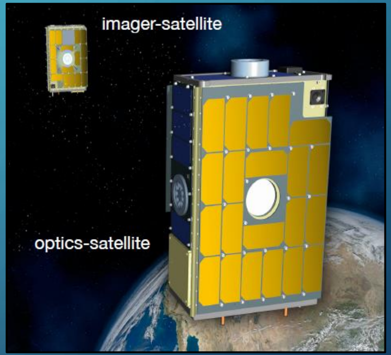

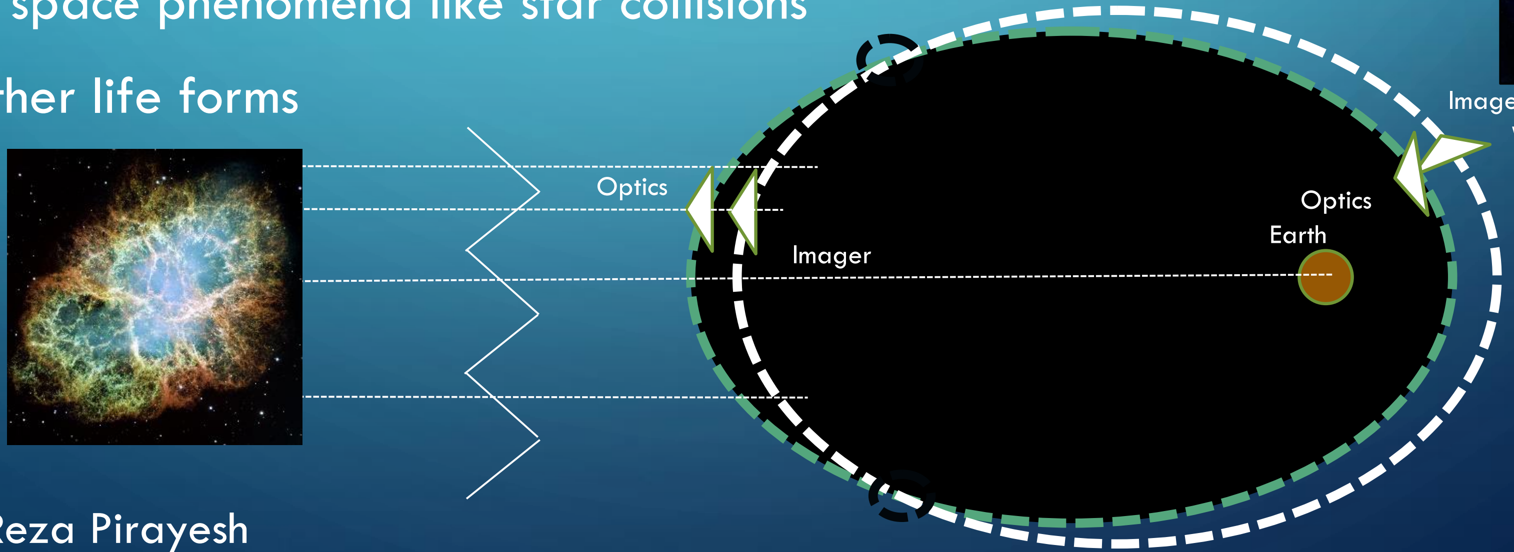

Virtual Telescope

Courtesy Reza Pirayesh

VIRTUAL TELESCOPE FOR X-RAY OBSERVATIONS(VTXO)

- Applying more sensors and sensor fusion algorithm
- Develop artificial intelligence algorithms to decrease the error in pointing to the targets in space
- Combining relative position control and attitude control algorithms to increase the performance of the formation flying control system

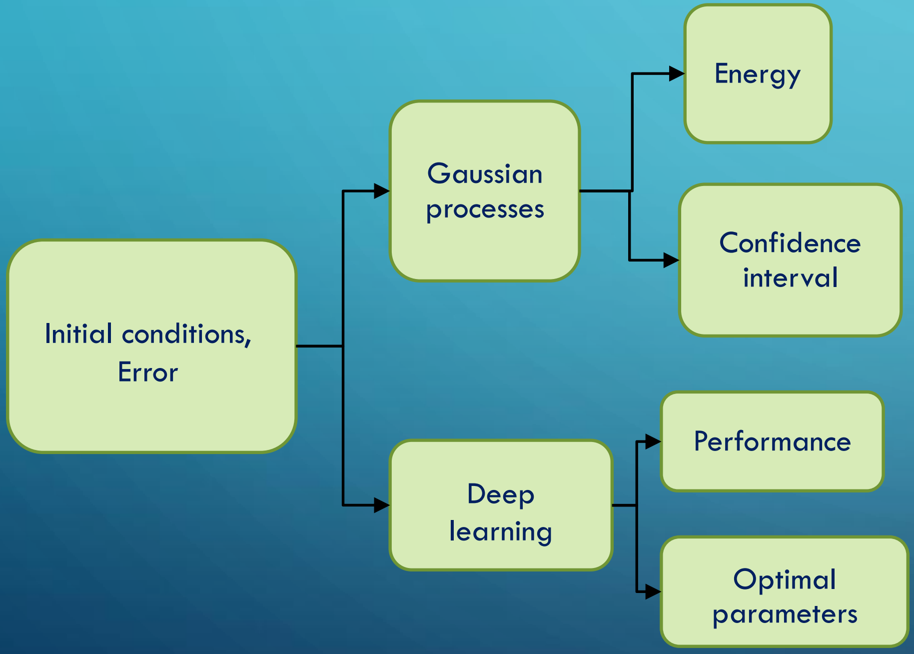

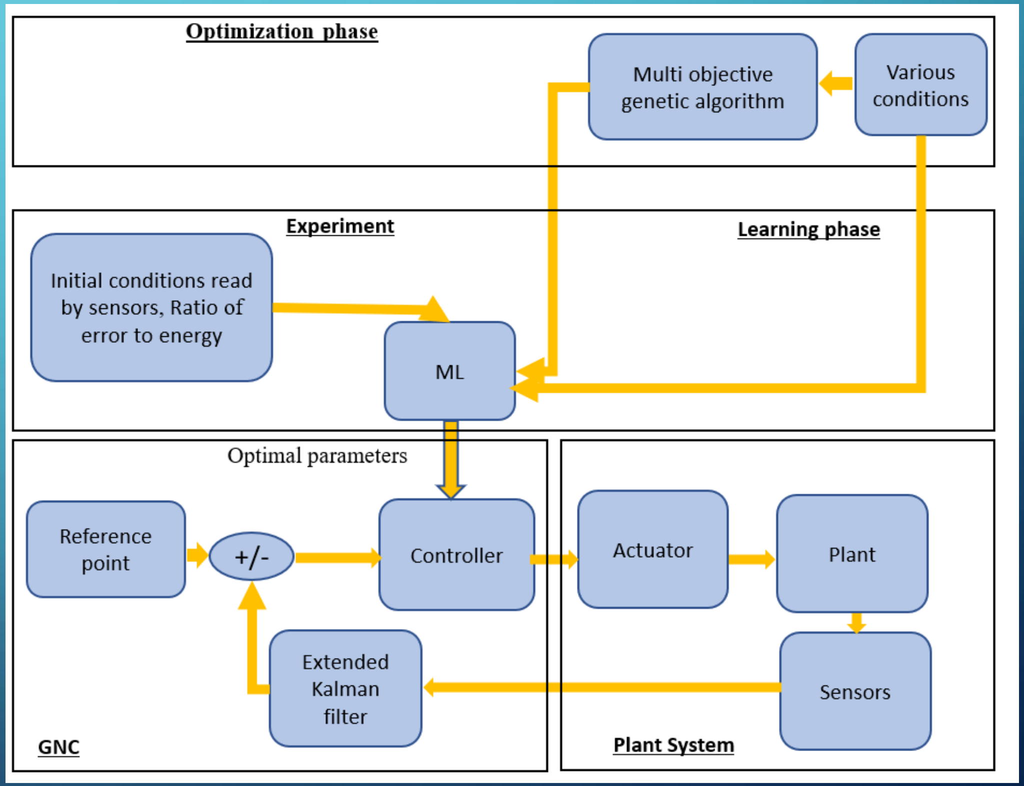

Courtesy Reza Pirayesh; Small Sat, ICALEPS

COHERENT CONTROL OF PLASMA DYNAMICS IN A LASER WAKEFIELD ELECTRON ACCELERATION EXPERIMENT

- The **first input** of the genetic algorithm under open-loop conditions and the corresponding **final output** of the genetic algorithm will be used as input and relevant output respectively to train the neural network.

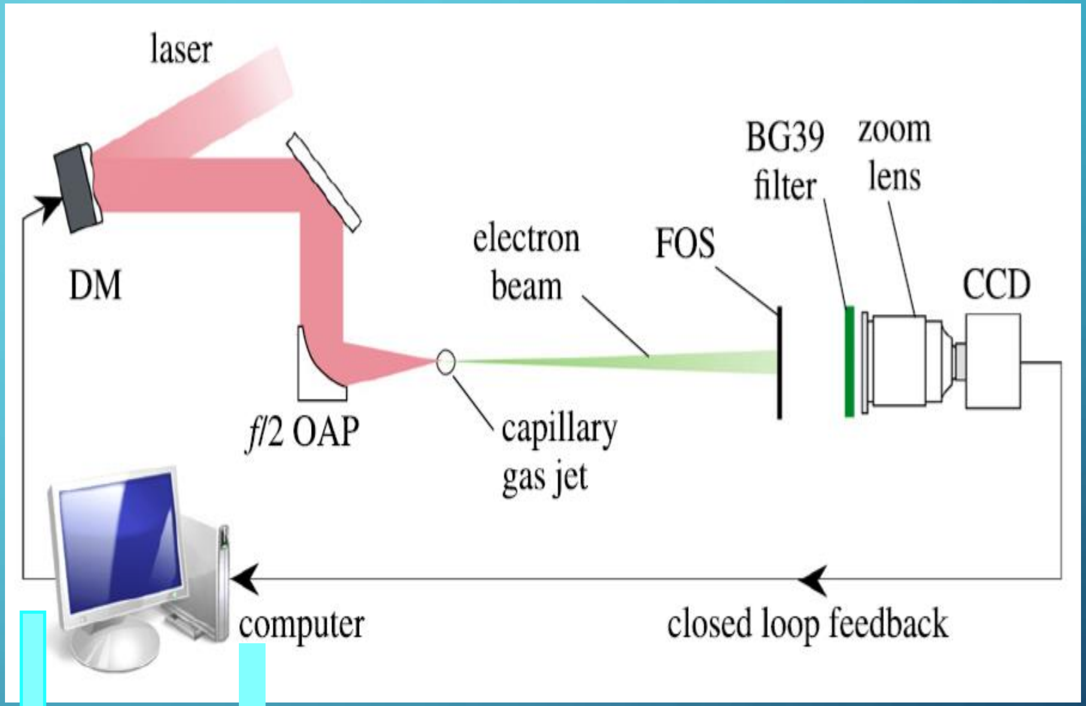

Output for training

Input for training

CNN Architecture

- **Convolutional Neural Networks** take **advantage** of local spatial coherence in the input (often images), which allow them to have fewer weights as some parameters are shared.
- This process, **taking** the form of convolutions, makes them especially well suited to extract relevant information at a low computational cost.

Courtesy A. Aslam, research with the University of Michigan (Thomas, Krushelnick)

COHERENT CONTROL OF PLASMA DYNAMICS IN A LASER WAKEFIELD ELECTRON ACCELERATION EXPERIMENT

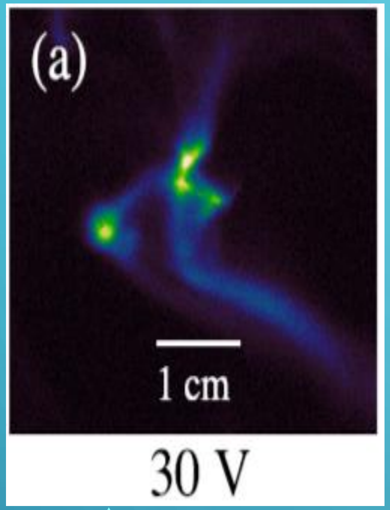

CNNTraining

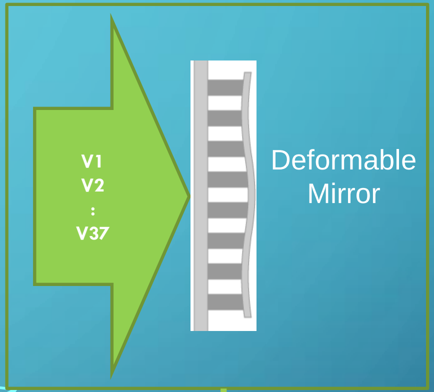

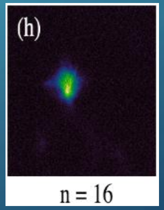

The **first input** of the genetic algorithm under open-loop conditions and the corresponding **final output** of the genetic algorithm will be used as input and relevant output respectively to train the neural network.

Courtesy A. Aslam, research with the University of Michigan (Thomas, Krushelnick)

COHERENT CONTROL OF PLASMA DYNAMICS IN A LASER WAKEFIELD ELECTRON ACCELERATION EXPERIMENT

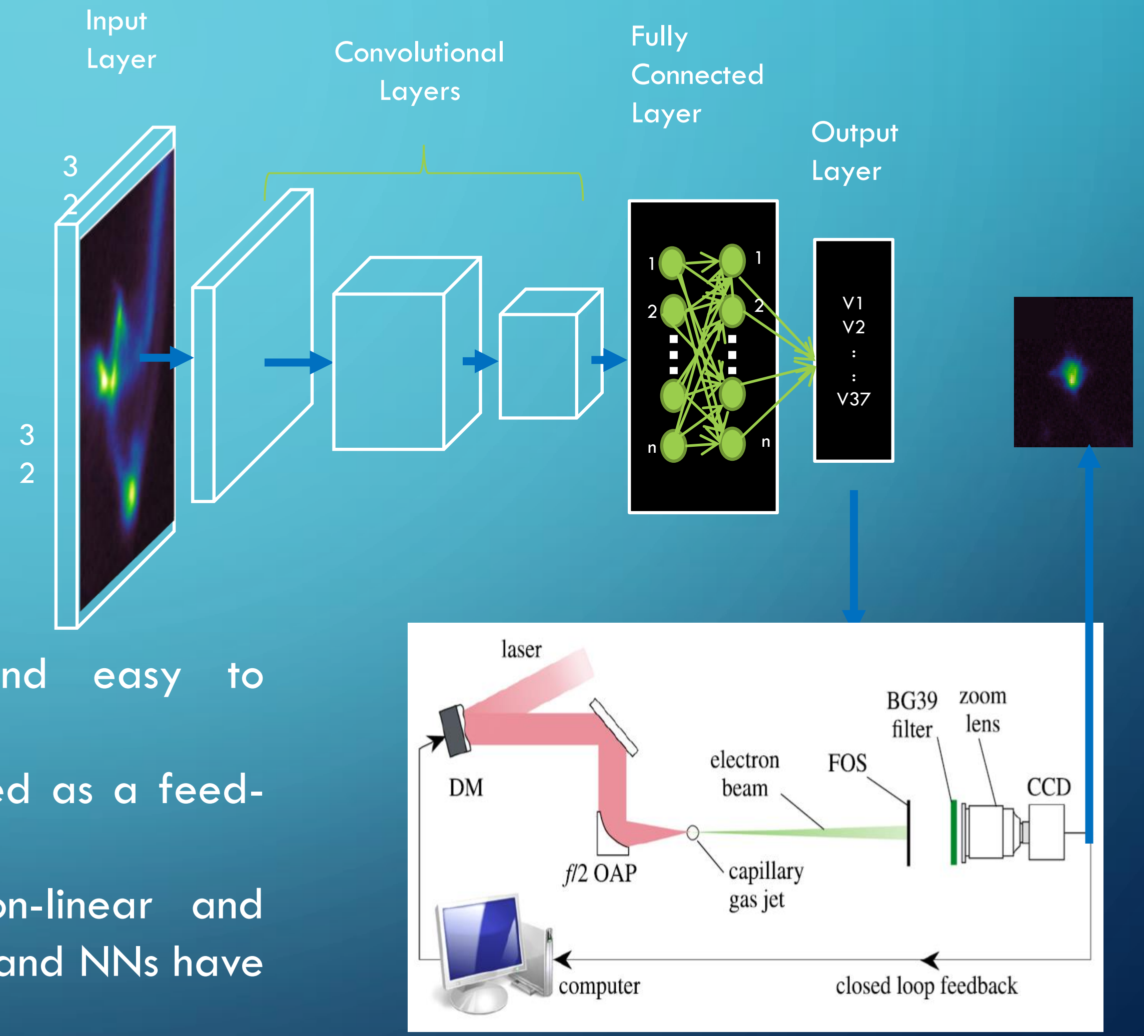

- Once trained successfully, much faster, and easy to implement.
- Can be improved evolutionary while being used as a feed-forward control.
- Most of the real-world problems have non-linear and complex relationship between input and output and NNs have ability to learn and model such relationship.

Courtesy A. Aslam, research with the University of Michigan (Thomas, Krushelnick)

PACMAN: Particle Accelerators & Machine Learning

Co-PIs:	Jochem Snuverink (PSI), Tatiana Pieloni (EPFL),
Partner(s):	Andreas Adelmann, Markus Janousch, Davide Reggiani (PSI)
	Olivier Schneider (EPFL), Anastasia Pentina (SDSC)
PostDocs & PhDs:	Jaime Coello, Loïc Coyle, Sichen Li, Michael Schenk

Problem :
Utilising ML to operate particle accelerators *safely* (LHC and HIPA):

A. Minimise beam losses
B. Better control of accelerator parameters
C. Prevent unnecessary machine interruptions
D. Improve beam dynamics modelling

Solution: Two Tier Model

1. Construct surrogate models from data and sim.
2. Apply surrogate models to online operation and compare to predictions

Several lines of attack:

A. Safe Bayesian Optimisation
B. Surrogate models to enhance performance
C. Prognostics to detect changes in the system
D. Neural networks instead of particle tracking

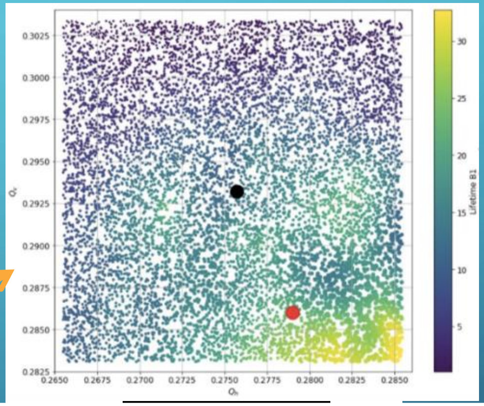

Beam lifetime

● : Nominal
● : Optimised

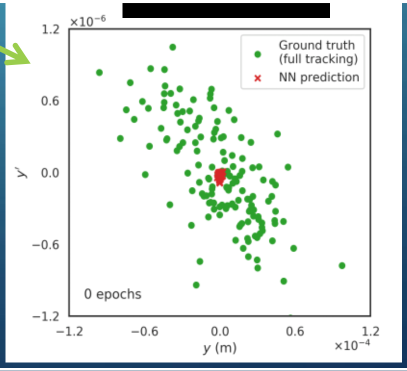

NN tracking training

phase space after 1 turn

speedup factor: 20 x

Courtesy of Jochem Snuverink

Impact:

- Safe and improved operation
- Reduced component damage
- Increased uptime
- Improve or change future accelerator design
- Optimise data storage for ML

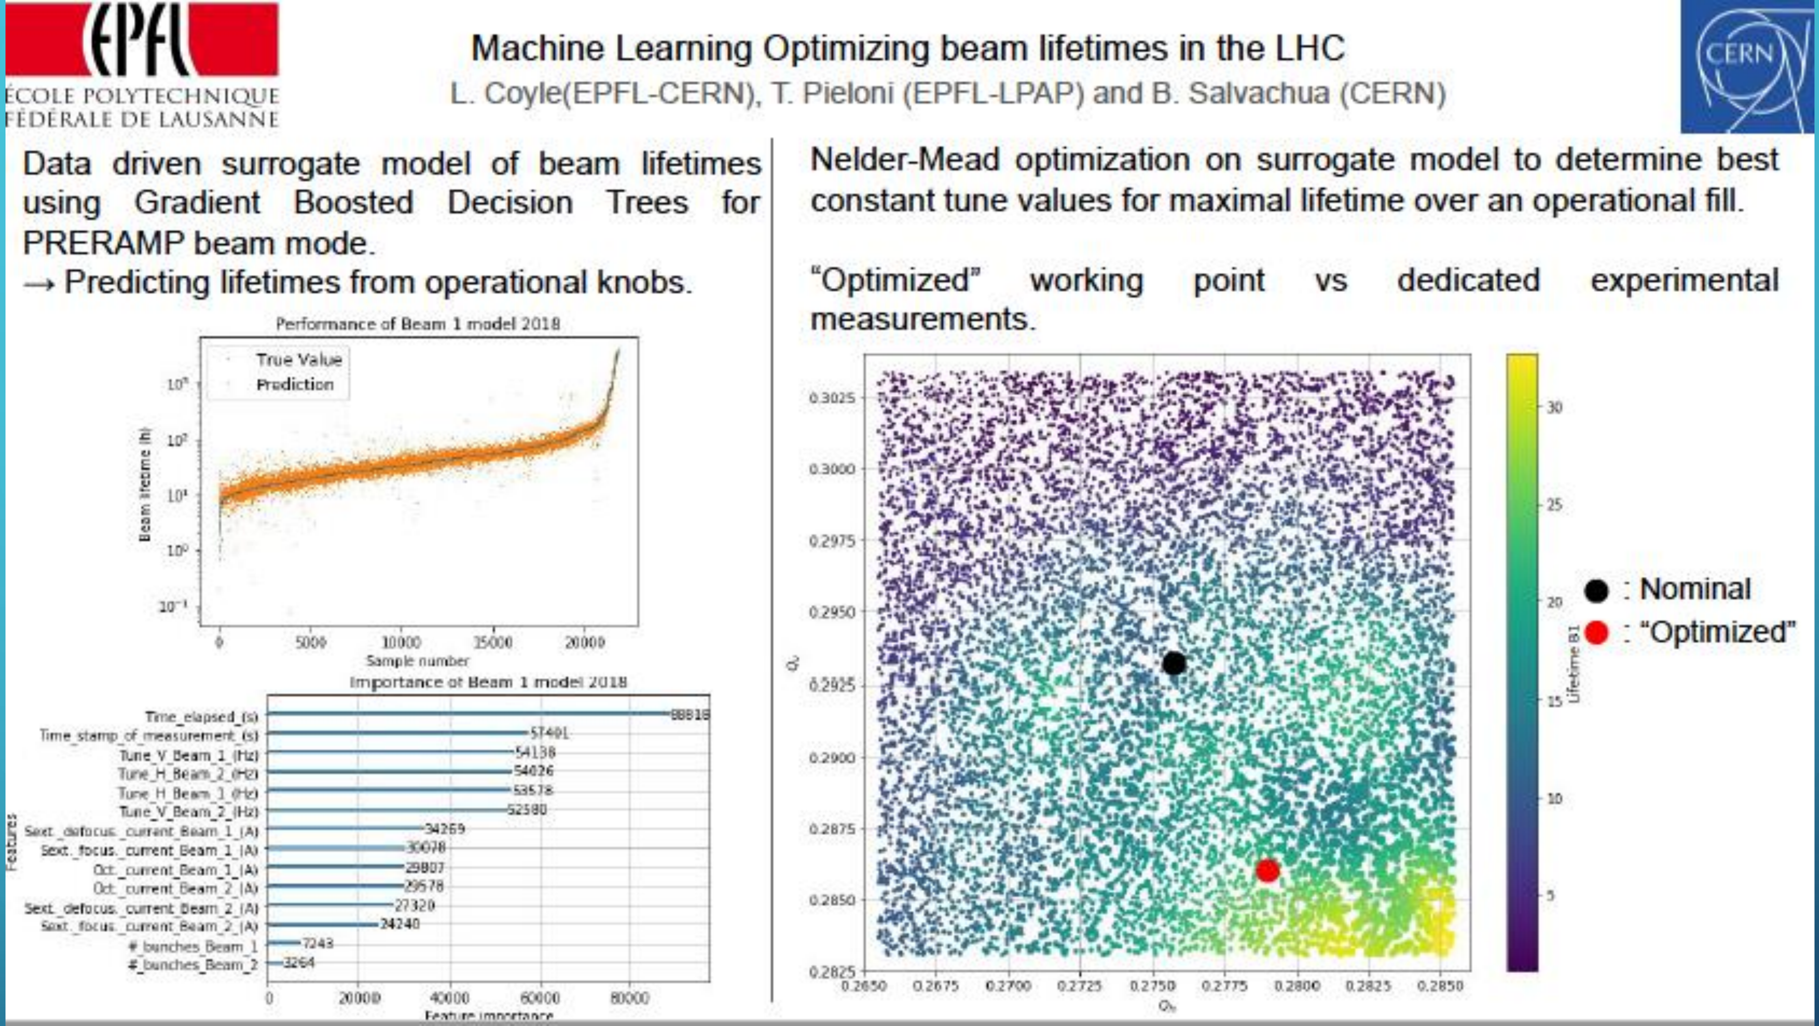

- LHC beam lifetime _> With the LHC physics fills of 2018 year they built a surrogate model of the LHC at injection energy. With that model they determined the best working point (RED DOT). Then they did and experiment scanning the various parameters they had in the model and measured lifetimes (color code points). The yellow area is where the lifetime is the highest and also the prediction of the surrogate model. They believe larger emittances in some cases because of collective instabilities that the trained model did not have (thus a small yellow area).

Courtesy of Jochem Snuverink

RECENT EVENTS HIGHLIGHTING INTELLIGENT TECHNIQUES INCLUDING IN CONTROL

ELEMENT AERO

THESE WILL CERTAINLY LEAD TO OPPORTUNITIES IN OUR CONTROLS COMMUNITY

Invite-only summit to launch the AI office in DOE with members of industry, government, labs, and academia

Four town halls are being conducted to collect input on the needs and visions of the community. Several of my team members and I attended the meetings at Argonne and LBL and tried provided view points of the controls and scientific facilities communities.

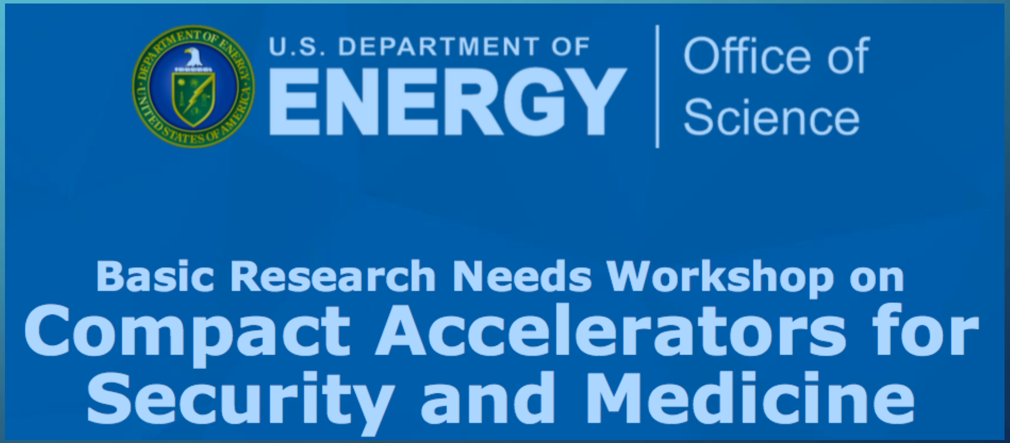

Several of the accelerator facility community were part of the controls, computing mad design cross-cut panel. Report expected in the coming months.

ELEMENT AERO

SERIES OF RECENT AND RELATED WORKSHOPS

In the last years, workshops have been held demonstrating the interest in artificial intelligence techniques **in particle accelerators as well as one for physics in general,** and 18 in high-energy physics

- **"Intelligent Controls for Particle Accelerators,"** was held 30-31 January 2018 at Daresbury Laboratory www.cockcroft.ac.uk/events/ICPA/
- **"ICFA Beam Dynamics Mini-Workshop: Machine Learning Applications for Particle Accelerators,"** was held 27 February-2 March 2018 at SLAC conf.slac.stanford.edu/icfa-ml-2018/
- **"Physics Next: Machine Learning,"** was held by the American Physical Society Physical Review publications group in New York on 8-10 October 2018 https://journals.aps.org/physics-next/2018/
- **"2nd ICFA Workshop on Machine Learning for Charged Particle Accelerators,"** was held from 26 February 2019 to 1 March 2019 at PSI https://indico.psi.ch/event/6698/
- In HEP - 18 Workshops in ML for HEP between 2015-2018 See list in Promise and Challenges of Machine Learning in Particle Physics, Astrophysics, and Cosmology Kyle Cranmer (New York University) https://cdn.journals.aps.org//cf3bab2e-015c-4eac-813a-570bc512c5e9/KyleCranmer.pdf
- ETC

ALSO – LOOKING OVER THE FENCE

- We need to look over the fence at groups such as HEP for guidance with large-scale data acquisition.
- Not all improvements need to be focused around AI. There are many controls algorithms, for instance, that can be adapted from other fields first. (See list of journals in the back-up charts).
- We need to be smart about using intelligence (See background from the IEEE Standards society working to ethical design practices.)

HOW DO WE GET THERE?

- WARNING! – Data science is not a magic black box. **It is not magic.** It is just another tool that like all of our tools, can be powerful if implemented properly. (one way is through AI)

Element Aero

SYSTEMS ENGINEERING IS PART OF THE EQUATION

Building In Intelligence And Data Science From Day 1

- We as a community do not apply systems engineering principles to particle accelerators. (including model based engineering).
- When planning a new or upgrading an old-ish machine, we need to remind ourselves that we need to come out of the 1930s when architecting the machines.
- If we want to use data science tools, then we need to architect this goal into the system (or system upgrades) – real-time data logging, local computing (e.g. GPUs, FPGAs), application-specific integrated circuits, co-location to a powerful cluster or HPC, better diagnostics, better electronics, etc.
- We need to be able to architect in how we can integrate/upgrade with new technologies.
- From our experience, for instance, powerful and smart controls tools cannot be mounted onto systems not architected to handle such data science applications!

If we want to have data science as a priority, the facets of data science must be architected into our system or sub-system.

Element Aero

RESEARCH FOR US TO DO

- ML can look for glitches in our software (simulation software as well as controls software). ML, for instance, can learn where the glitches are, find vulnerabilities, and even help detect cyber breaches (control and data systems). *Think of this as having a few extra (thousand, million, etc.) friends helping manually verify the software and protect your accelerator.*
- Use of specialized computing hardware — such as FPGAs, application-specific integrated circuits, systems-on-a chip are in wider use by our cousins in HEP and are natural stepping stones toward dedicated AI hardware – if we can implement, test and continue to develop on many accelerator systems.
- In the case of the quest for compact, accelerators, the more data we have on a "prototype" the better and cheaper we can make the next versions to be produced and sold. Intelligent methods will be important here.
- We need to generate new intelligent techniques and algorithms specifically to address the underlying physics in an accelerator and peripheral systems (e.g. lasers).
- We have some of the most interesting data sets in the world generated on our machines and by our users. Just in the acquired data thus far, we can be exploring accelerator-specific data science techniques.
- We need to be aware that controls using traditional techniques and/or more intelligent techniques could be improved through better use of data science.

CONCLUSIONS

ELEMENT AERO

- To effectively apply data science to the accelerators, you must have knowledge of the system and physics, engineering, materials, chemistry, etc. behind it.
- Need to choose the right algorithm to tell you about the system, not just about the data interpretation you already know.
- Need to find the right data set for the question(s) you have. This includes the diagnostic (how you got the data).
- We need to realize that we ourselves have to improve our own methods or adapt to a "new job" (e.g. we are now taking something we examined with genetic algorithms and now using a deep belief network algorithm to formulate our control.)
- We need to convince the managers as well as the funding agencies that data science is a low-hanging fruit in scientific systems.
- Better reliability and ease of operation can help make compact accelerators a reality! BUT WE NEED DATA SCIENCE (and systems engineering). E.g. the lifetime of an optical grating of your laser needs to be in your model, inspection/test data, etc.
- The difference between training a horse, for instance, and a machine is that we have control over the systems engineering architecture of the machine itself and the algorithms but not in the case of the horse. [Control starts from day 1.]
- Architect in from day 1 the computational resources. Things will change over time and eventually computing will be some convergence of neurons, bits, and qubits. Be prepared that AI will drive some of these changes and that your system could be architect to "accept" these.
- Don't forget about the end user – their systems need to be part of your model and their systems can also be better automated.
- Lots of work yet to do.

The only working model of a universe is a universe.

ACKNOWLEDGEMENTS

- Thanks to Jim Amundson, Josh Stein, John Cary, Richard Farnsworth, Steve Lidia, John Petillo, Massimo Dal Forno, Mark Curtin, Manel Martinez-Ramon, Stephen Milton, Jorge Alberto Diaz Cruz, Aasma Aslam, Trudy Bolin, Salvador Issac Guitron Sosa, Reza Pirayesh, and to ONR, DOE, Sincrotrone Trieste, and NASA for support that inspired this work either directly or indirectly.

BACK UP MATERIALS

CONTACTS AT THE IEEE

- Alpesh Shah, alpesh.shah@ieee.org and Sam Sciacca, s.sciacca@ieee.org
- https://ethicsinaction.ieee.org/standards/ p7000/

WORKING TOGETHER TOWARD A HUMAN-CENTRIC AI?

Socio-political Awareness

Socio-political Awareness

Legislative actions

Regulative Actions/ Codes

Engagement of civil society

Does the top-down meet the bottom-up?

Internal Boards (Google)

External Boards (SAP)

Consortia (Partnership on AI)

Open Standardization (IEEE,BSI)
Code building (EAD IEEE)

Industry self-regulation

LOTS OF PEOPLE SUGGEST “AI PRINCIPLES”

- Several lists, some original, some “cut and paste” from others
- Reflecting various interests and agendas (industry lobbies, activist NGOs, proxy-religions, existential fears, etc.)
- Confusing categorization, mixing of political, cultural and technical levels
- Partly overlapping, even conflicting aspects

- For this OECD group to have a chance to succeed, we will need to establish rather fast a conceptual framework to address the categorization problem.
- A possibility to conceptualize at a higher level such principles could be to follow the categorization implicit in slide 2 of Ms Doaa "Multi-disciplinary approaches to Ethical AI"

FROM "INTERDISCIPLINARY APPROACHES" TO "PRINCIPLES CATEGORIES" ?

- Philosophical Approaches → *Values principles/aspirational intentions*

  Which societal, moral, cultural, political values do we **intend** to promote through intelligent systems?
  Self-determination, political autonomy/democracy, dignity, privacy, social fairness,….

- Legal Approaches → *Legal principles/obligations*

  Which current or emerging legal frameworks **must** we take into account?
  Human rights, children rights, personal data protection, liability laws, ...

- Computational Approaches → *Systems design principles/ technical implementations*

  -
  Are the systems **dependable**, ie are they doing what they are supposed to do?
  How to measure and verify?
  The 3 principles on Slide 5 of Ms Doaa fall under this category.

IEEE'S PATH TOWARD A HUMAN-CENTRIC AI

- Established in 2016 the **IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems**
- It is an open, global and inclusive (regionally, culturally, gender etc.) community of experts and interested persons from technology and human science to address this question
- Within two years, a community of ~ 1500 people from all continents, ~ 40 % women
- Its mission is to **ensure that every stakeholder involved in the design and development of autonomous and intelligent systems is educated, trained, and empowered to prioritize ethical considerations so that these technologies are advanced for the benefit of humanity**
  - There seems to be a great affinity between this and OECD's mission regarding a human-centric AI.

Other adjacent IEEE (SA) Initiatives:
Global Council on Extended Intelligence
Open Community for Ethics in Autonomous and Intelligent Systems (OCEANIS)

WORKS OF THE IEEE GLOBAL INITIATIVE ON ETHICS OF AUTONOMOUS AND INTELLIGENT SYSTEMS

- **Ethically Aligned Design**
  An iterative process over 3 years with open and massive participation toward a better understanding of the fundamental issues, combined with a process for suggesting high-level principles and actions to address them. No corporate or geopolitical or activists agendas seem to dominate the process. **Crowd-sourced, collective, unbiased intelligence at its best**
  - EAD V1 (2016) and V2 (2017) received 100s of pages of public feedback
  - Currently, 13+ Committees are creating content for V3 (early 2019)
  - **The most comprehensive, crowd-sourced global treatise regarding the ethics of Autonomous and Intelligent Systems available today**
  - Provides an open platform for thought leadership and action to prioritize values-driven, ethically-aligned design for autonomous and intelligent systems.
- **P7000 series of standardization projects**
  Where we believe that - beyond suggesting high-level principles - we may become more concrete and achieve consensus in a meaningful time frame.
- **Education material/curricula**
  Establishment of EAD University Consortium (in 2019).

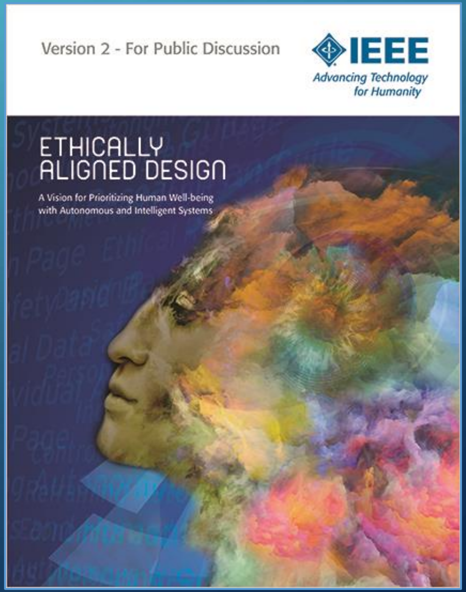

Global Council on Extended Intelligence

- IEEE-SA and the MIT Media Lab announced the launch of the global Council on Extended Intelligence (CXI) on 22 June 2018
- The goals of CXI are to build a new narrative for autonomous and intelligent technologies inspired by principles of systems dynamics and design.
- CXI will promote the deals of responsible participant design, data agency and metrics of economic wellbeing prioritizing people and planet over exponential growth, including these projects:
    - **Extended Intelligence - Awareness and Action –** will focus on creating an introduction to Extended Intelligence and Participatory Design
    - **Digital Identity - Democracy by Design –** will focus on creation of Data Policy template for governments and organizations to utilize in helping individuals and society reclaim their digital identity in the algorithmic age
    - **Enlightened Indicators - Measuring What's Good Versus Simply Growth –** will focus on creating a Wellbeing Indicator template, a metric governments and organizations can utilize in genuinely measuring prosperity in terms of benefit to all

Open Community for Ethics in Autonomous and Intelligent Systems (OCEANIS)

- IEEE-SA jointly launched and became a founding member on 25 July 2018
- OCEANIS is a high level global forum for discussion, debate and collaboration among organizations interested in the development and use and development of standards in autonomous and intelligent systems.
- Community will address needs for coordination and collaboration related to the unprecedented challenges faced by those working in ICT standards and related spaces.
- OCEANIS is open to all interested organizations.
  In addition to IEEE-SA, Founding Members include the:

-African Regional Organization for Standardisation (ARSO)
-Austrian Electrotechnical Association (OVE)
-Austrian Standards International (A.S.I.)
-British Standards Institution (BSI)
-China Electronic Standardizations Institute (CESI)
-CIO Strategy Council (Canada)
-International Electrotechnical Commission (IEC)
-Ecuadorian Service for Standardization (INEN)
-National Standards Authority of Ireland (NSAI)
-Turkish Standards Institute (TSE)
-Verband und Deutsche Kommission Elektrotechnik Elektronik Informationstechnik

IEEE Initiatives & activities around autonomous/intelligent

<table>
<tr><th>Standards</th><th>Other Activities</th></tr>
<tr><td rowspan="3">IEEE 7000 Series Standards:
• 7000 - Model Process for Addressing Ethical Concerns During System Design
• 7001 - Transparency of Autonomous Systems
• 7002 - Data Privacy Process
• 7003 - Algorithmic Bias Considerations
• 7004 - Standard for Child and Student Data Governance
• 7005 - Standard for Transparent Employer Data Governance
• 7006 - Standard for Personal Data Artificial Intelligence (AI) Agent
• 7007 - Ontological Standard for Ethically Driven Robotics and Automation Systems
• 7008 - Standard for Ethically Driven Nudging for Robotic, Intelligent and Autonomous Systems
• 7009 - Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems
• 7010 - Wellbeing Metrics Standard for Ethical Artificial Intelligence and Autonomous Systems
• 7011 - Standard for the Process of Identifying and Rating the Trustworthiness of News Sources
• 7012 - Standard for Machine Readable Personal Privacy Terms
• 7013 - Inclusion and Application Standards for Automated Facial Analysis Technology</td><td>Industry Connection Group:
• The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems
• Launch of the NeuroTech Industry Connections Group</td></tr>
<tr><td>IEEE Brain Initiative Workshop on Advanced NeuroTechnologies</td></tr>
<tr><td>IEEE Brain Sensors Workshop</td></tr>
</table>

OTHER RELATED IEEE ACTIVITIES

Journals-examples

- IEEE Transactions on Neural Networks and Learning Systems
- IEEE Intelligent Systems
- IEEE Transactions on Fuzzy Systems
- IEEE Transactions on Nuclear Science
- IEEE Transactions on Human-Machine Systems
- IEEE Transactions on Cybernetics
- EEE Transactions on Control Systems Technology
- IEEE Transactions on Automation Science and Engineering
- IEEE Transactions on Industrial Informatics
- IEEE Transactions on Systems, Man, and Cybernetics: Systems
- IEEE Transactions on Circuits and Systems
- IEEE Transactions on Pattern Analysis and Machine Intelligence
- IEEE Transactions on Cognitive and Developmental Systems
- IEEE Transactions on Image Processing
- IEEE Robotics and Automation Letters
- IEEE Transactions on Emerging Topics in Computational Intelligence

OTHER RELATED IEEE ACTIVITIES

Journals- special issues- examples

- Computer – Sept 2018 – The future of AI
- IEEE JSAC (Journal on Special areas in Communications) Special Issue on Artificial Intelligence and Machine Learning for Networking and Communications
- IEEE Transactions on Engineering Management - Special Issue: Services Computing Management for Artificial Intelligence and Machine Learning
- IEEE Transactions on Network Science and Engineering-Special Issue on Big Data and Artificial Intelligence for Network Technologies
- Journal of Biomedical and Health Informatics – Special issue -Pervasive Sensing and Machine Learning for Mental Health
- IEEE Internet of Things Journal - RRCPS: Reliable and Resilient Cyber-Physical Systems

OTHER RELATED IEEE ACTIVITIES

Artificial Intelligence (34)-conferences

- 2018 International Conference on Advanced Computation and Telecommunication (ICACAT), 28 Dec - 29 Dec 2018
- 2018 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR), 10 Dec - 12 Dec 2018
- 2018 3rd International Conference on Pattern Analysis and Intelligent Systems (PAIS), 24 Oct - 25 Oct 2018
- 2018 IEEE 3rd Advanced Information Technology, Electronic and Automation Control Conference (IAEAC), 12 Oct - 14 Oct 2018
- 2018 First International Conference on Secure Cyber Computing and Communication (ICSCCC), 11 Oct - 13 Oct 2018
- 2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC), 07 Oct - 10 Oct 2018
- 2018 First International Conference on Artificial Intelligence for Industries (AI4I), 26 Sep - 28 Sep 2018
- 2018 Joint IEEE 8th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob), 17 Sep - 20 Sep 2018
- 2018 Intelligent Systems Conference (IntelliSys), 06 Sep - 07 Sep 2018
- 2018 IEEE International Conference on Autonomic Computing (ICAC), 03 Sep - 07 Sep 2018
- 2018 IEEE Conference on Computational Intelligence and Games (CIG), 14 Aug - 17 Aug 2018
- 2018 International Conference on Computer and Applications (ICCA), 06 Aug - 07 Aug 2018
- 2018 IEEE International Work Conference on Bioinspired Intelligence (IWOBI), 18 Jul - 20 Jul 2018
- 2018 33rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), 09 Jul - 12 Jul 2018
- 2018 10th International Conference on Modelling, Identification and Control (ICMIC), 02 Jul - 04 Jul 2018
- 2018 10th International Conference on Electronics, Computers and Artificial Intelligence (ECAI), 28 Jun - 30 Jun 2018
- 2018 19th IEEE/ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD), 27 Jun - 29 Jun 2018
- 2018 IEEE International Conference on Smart Computing (SMARTCOMP), 18 Jun - 20 Jun 2018
- 2018 Second International Conference on Intelligent Computing and Control Systems (ICICCS), 14 Jun - 15 Jun 2018
- 2018 Global Internet of Things Summit (GIoTS), 04 Jun - 07 Jun 2018

OTHER RELATED IEEE ACTIVITIES

Artificial Intelligence (34)-conferences
• 2018 International Conference on Optical Network Design and Modeling (ONDM), 14 May - 17 May 2018 • 2018 International Symposium on Consumer Technologies (ISCT), 11 May - 12 May 2018 • 2018 9th International Conference on Information and Communication Systems (ICICS), 03 Apr - 05 Apr 2018 • 2018 International Conference on Control, Automation and Diagnosis (ICCAD), 19 Mar - 21 Mar 2018 • 2018 13th ACM/IEEE International Conference on Human-Robot Interaction (HRI), 05 Mar - 08 Mar 2018 • 2018 International Conference on Intelligent Autonomous Systems (ICoIAS), 01 Mar - 03 Mar 2018 • 2018 International Conference on Emerging Technologies in Data Mining and Information Security (IEMIS), 23 Feb - 25 Feb 2018 • 2018 Second International Conference on Computing Methodologies and Communication (ICCMC), 15 Feb - 16 Feb 2018 • 2018 International Conference on Communication information and Computing Technology (ICCICT), 02 Feb - 03 Feb 2018 • 2018 International Conference on High Performance Computing in Asia-Pacific Region (HPC Asia), 29 Jan - 31 Jan 2018 • 2018 IEEE International Conference on Future IoT Technologies (Future IoT), 18 Jan - 19 Jan 2018 • 2018 10th International Conference on Agents and Artificial Intelligence (ICAART), 16 Jan - 18 Jan 2018 • 2018 IEEE International Conference on Big Data and Smart Computing (BigComp), 15 Jan - 17 Jan 2018 • 2018 8th International Conference on Cloud Computing, Data Science & Engineering (Confluence), 11 Jan - 12 Jan 2018

WHAT NOT TO DO

Resist to dogmata

Do not oversimplify

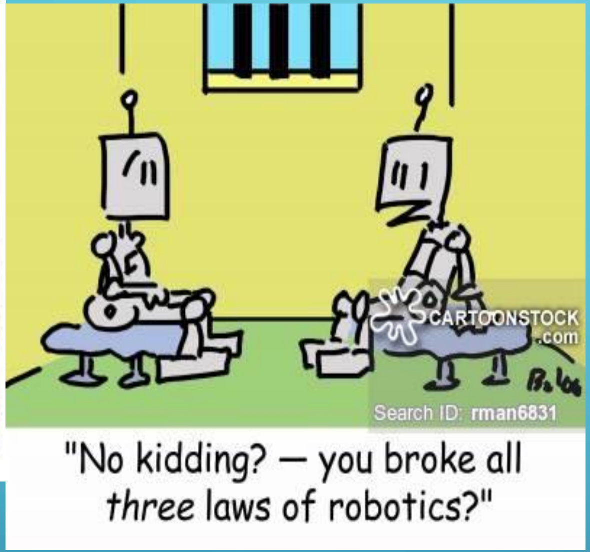

Don't confuse science with marketing

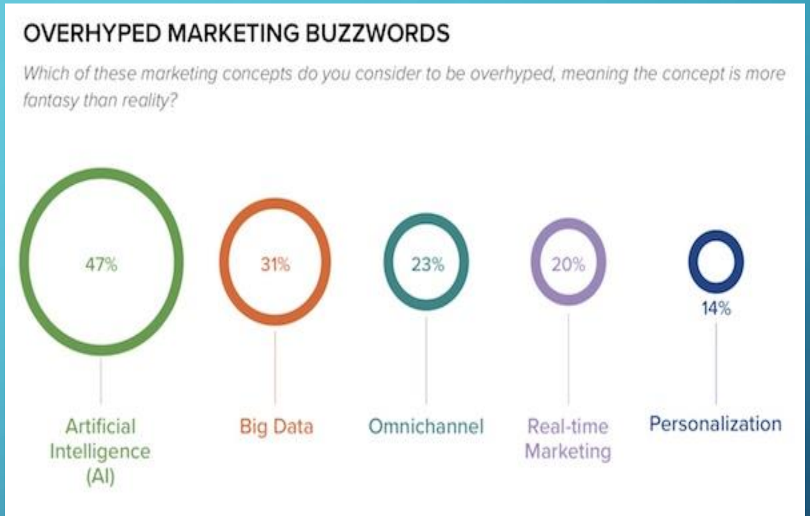

STEP 1. GET THE DATA

ELEMENT AERO

- We need to gather the data from our challenging accelerator or accelerator sub-system(s).
- This can be difficult as in especially challenging cases, you might not collect the correct data OR maybe you think something is outside the possibility of influencing a sub-system (so data not logged or not even monitored from this source).
- Make certain you collect (or generate if simulation only) enough data.

STEP 2. DATA PREPARATION

- May not need if you are doing more traditional data analytics not involving AI/machine learning.

REDUCE THE NUMBER OF FEATURES

A dataset can contain a large amount of features associated with it.

Reduction in features essentially summarizes the data, playing an important role in being able to apply techniques of data science.

Why do we need to do this? EXAMPLES

- Help reduce the computation or training time by algorithms.
- Preserve limited storage space.
- Permit use of some algorithms that do not perform well with large dimensions.

- Removes redundant features (reduce multicollinearity).
- Reductions helps in visualizing data (as reducing to 2D or 3D may permit plotting and observing patterns more clearly.

DIMENSIONALITY REDUCTION

- Formally - Dimensionality reduction is the process of reducing the number of random variables under consideration. Through reduction, one obtains a set of principal variables.
- Algorithms like Principal Component Analysis, Linear Discriminant Analysis, Random Forest, Single Value Decomposition, Missing Value Ratio, Low Variance Filter, High Correlation Filter, Backward Feature Elimination, Forward Feature Selection, Independent Component Analysis, Factor Analysis, Methods Based on Projections, Uniform Manifold Approximation and Projection (UMAP), t-Distributed Stochastic Neighbor Embedding (t-SNE) can be used to reduce the dimensionality.
- Again, such reduction will improve the results of, for instance, classification with a neural network.

SPLITTING THE DATA

- Need to split the data into two separate pots - training and evaluation data.
- You want to train your model on the training data and then evaluate the model's effectiveness through the evaluation data.

STEP 3. CHOOSE A LEARNING METHOD

- Choose a learning (training) method that is appropriate from the many that exist already suited to your data (image, text, numbers, etc…)
- Again – this is not a black box. Many algorithms and infinite combinations

ELEMENT AERO

STATISTICAL EXAMINATION OF THE DATA

We need to realize that not all the data we collect makes sense in terms of volume, type, or if there are errors.

- Mean average - determines overall trend of a data set
- Standard deviation - the spread of data around mean
- Regression - examines relation between variables
- Outlier detection & treatment - used where noise is present
- Support Vector Machines (SVM) based anomaly detection
- Clustering-based anomaly detection
- Hypothesis Testing

ELEMENT AERO

STEP 4. ALLOW THE MACHINE TO LEARN

- Build a candidate model from the data using the selected learning method.
- This can take many iterations.
- NOTE unsupervised ML models have the ability to self-learn patterns to deliver answers even when input data is unlabeled and has unknown outcomes.

“All models are wrong, but some are useful” George E. P. Box

STEP 5. EVALUATE

- Test your model on the evaluation data.

Step 6. Refine the model – tuning step(s)

7. Prediction or Use Step!

Element Aero

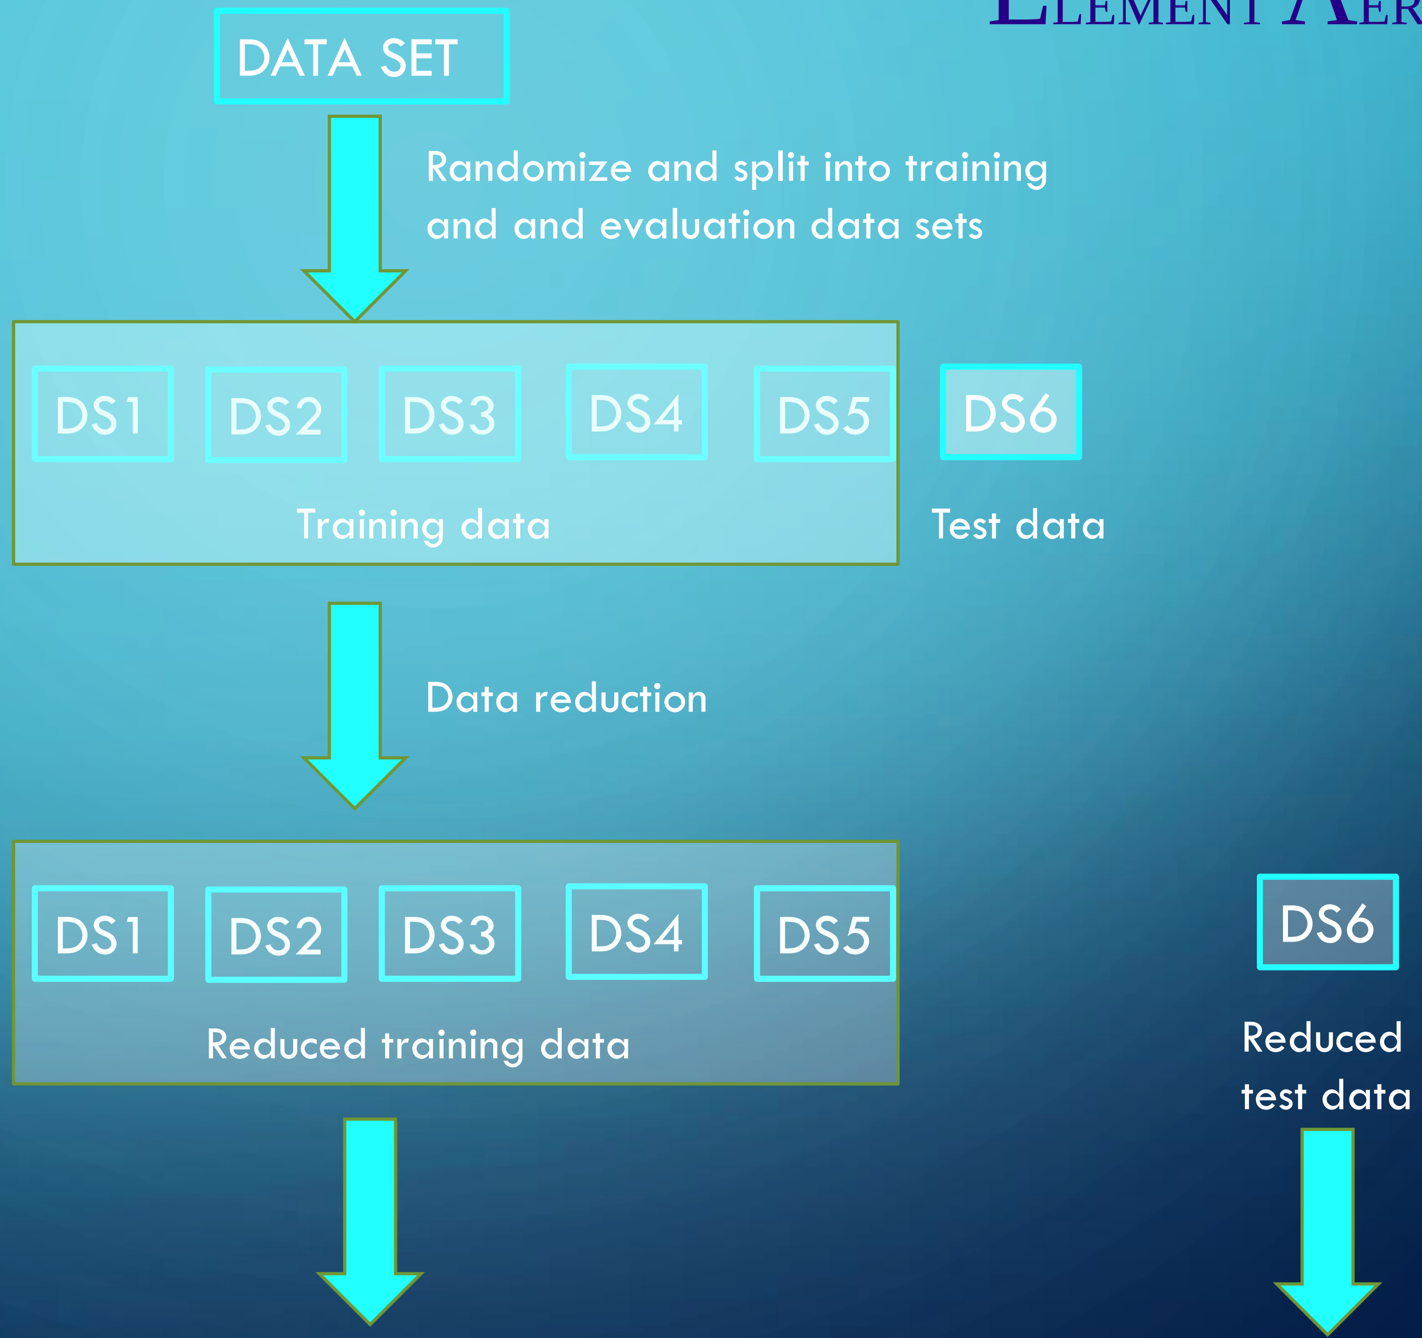

ELEMENT AERO

Select informative voxels and determine any optimal parameters

DS1 DS2 DS3 DS4 DS5

Ready training data

DS6

Reduced test data

Model building (e.g. Decision Trees, Neutral Networks, Linear regression, Bayesian networks, Nearest Neighbor, Support Vector Machine)

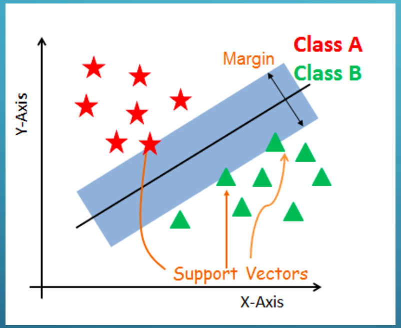

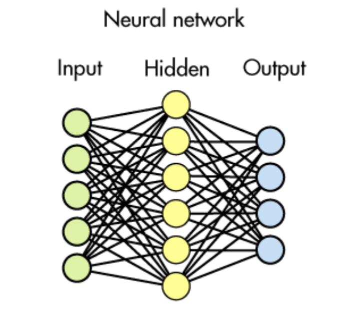

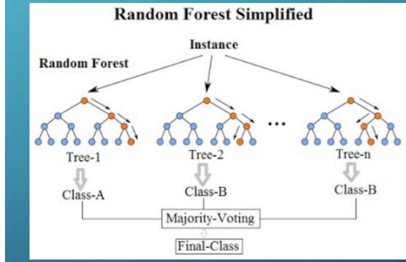

Test model

Evaluation

Repeat as necessary

A FEW ADDITIONAL EXAMPLES

- A few examples recently exploiting intelligent methods and HPC applied to particle accelerators.

Simulations and HPC are central to the APS Upgrade (M. Borland)

- Massively parallel multi-objective genetic optimizer to develop physics design
- Design of complex multi-function magnets using 3D codes
- Simulation of collective behavior of high-charge, multi-bunch electron beams
- Simulation of vacuum system performance in response to synchrotron radiation
- Simulation of commissioning and operational performance
- Understanding and pushing intensity limits in the injector complex
- Prediction of radiation levels and shielding requirements

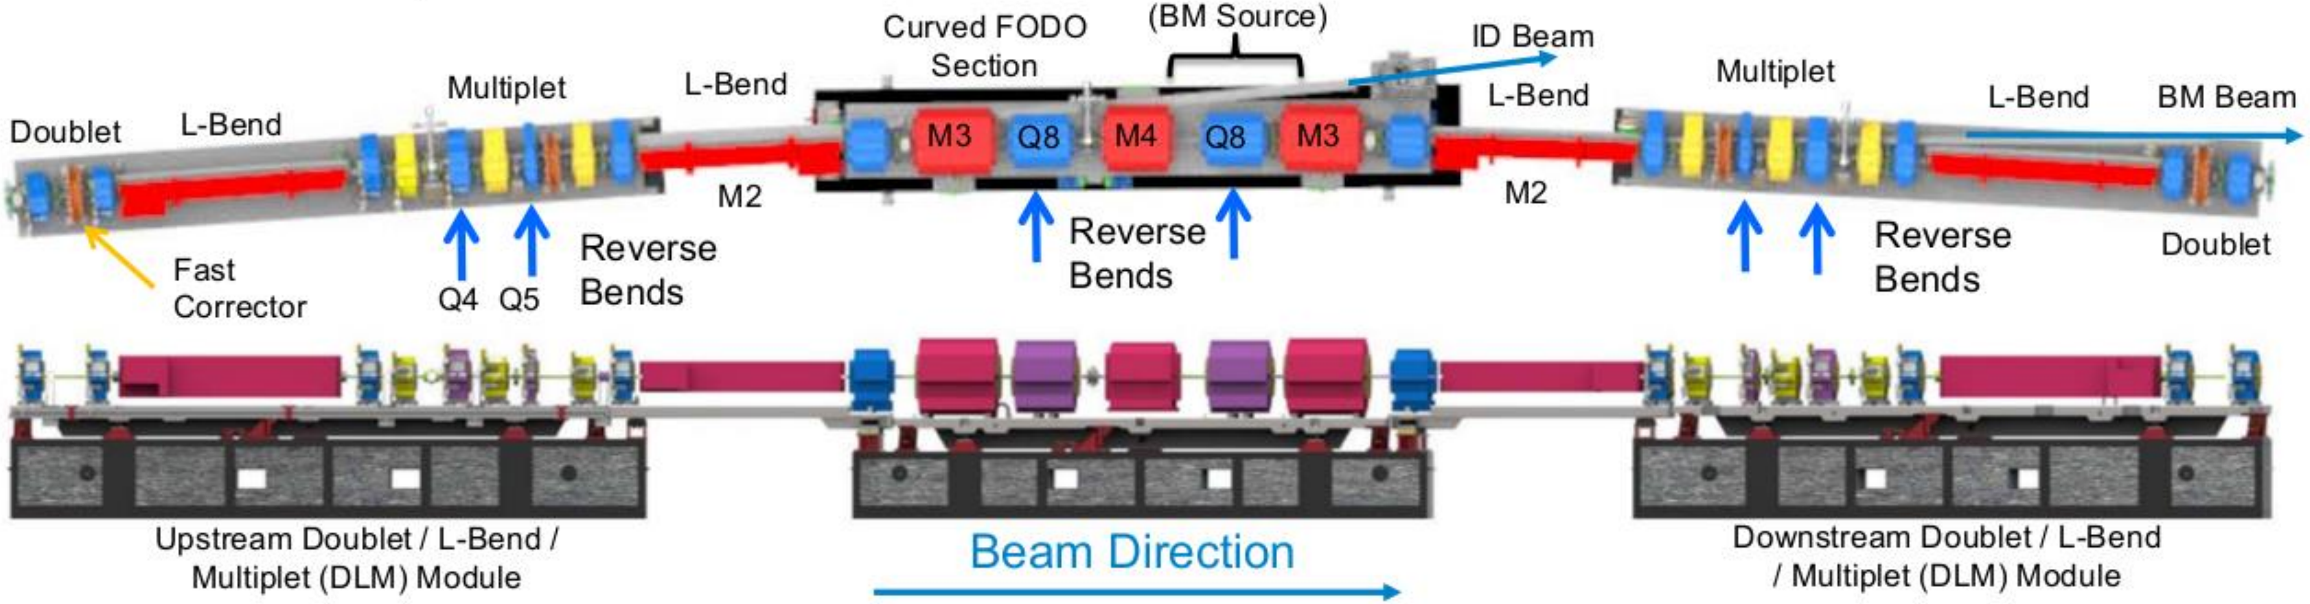

Image courtesy G. Decker, APS-U

Argonne NATIONAL LABORATORY

Courtesy M. Borland Argonne

HPC saves money in the end, but often underutilized

- To reduce design costs, HPC is essential
  - Provide high-fidelity models with fewer compromises, leading to fewer (often no) iterations with hardware prototypes
  - Make best use of expensive engineering and physics staff
  - Provide robust statistical data on expected performance
- Common reasons HPC is underutilized
  - Staff not trained to use HPC systems and software
  - Software chosen for familiarity and convenience, not HPC capabilities
  - Management undervalues computation, fails to provide access to HPC systems and support
- Example: APS upgrade beam physics effort provided with ~60M core hours per year on ANL internal computing clusters, excellent support
  - In early years, also given access to ALCF resources (not needed now)
  - Similar resources being devoted to ALS Upgrade

Argonne NATIONAL LABORATORY

Courtesy Jean-Luc Vay, LBL

Extreme computing enables detailed beam & accelerator science

Parallel Luminosity Optimization in conventional accelerator

Code: BeamBeam3D + optimizer

Time-to-solution:

- **serial code & optimizer:** ~35 years
- **Parallel code & optimizer:** ~1 day

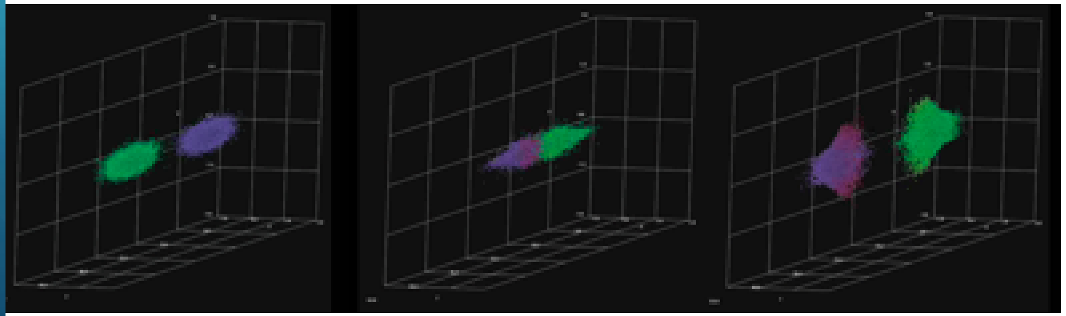

Study of pulse front tilt in laser-plasma accelerators

Code: WarpX

Time-to-solution (largest run):

- **serial code in lab frame: ~500 years**
- **Parallel code in boosted frame: ~4 h**

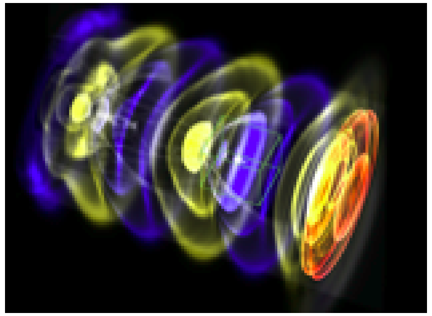

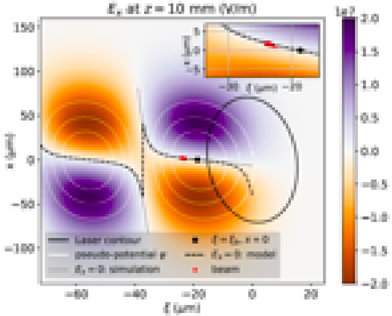

NERSC **Computer:** Cori at NERSC – 9,688 nodes, 658,784 cores.

ECP EXASCALE COMPUTING PROJECT

MINIMIZING HIGH-FREQUENCY JITTER

- In order to compensate for the PID algorithm deficiencies, we wanted to build a system that
  - Acted in a feed forward way, for correction of high-frequency jitter
  - Increased the bandwidth in order to avoid re-tuning when conditions change.
- We recorded data from the machine in order to train a neural network to predict a future deviation based on past records.
- We complemented the system with an optional PI algorithm to further decrease remaining deviations.

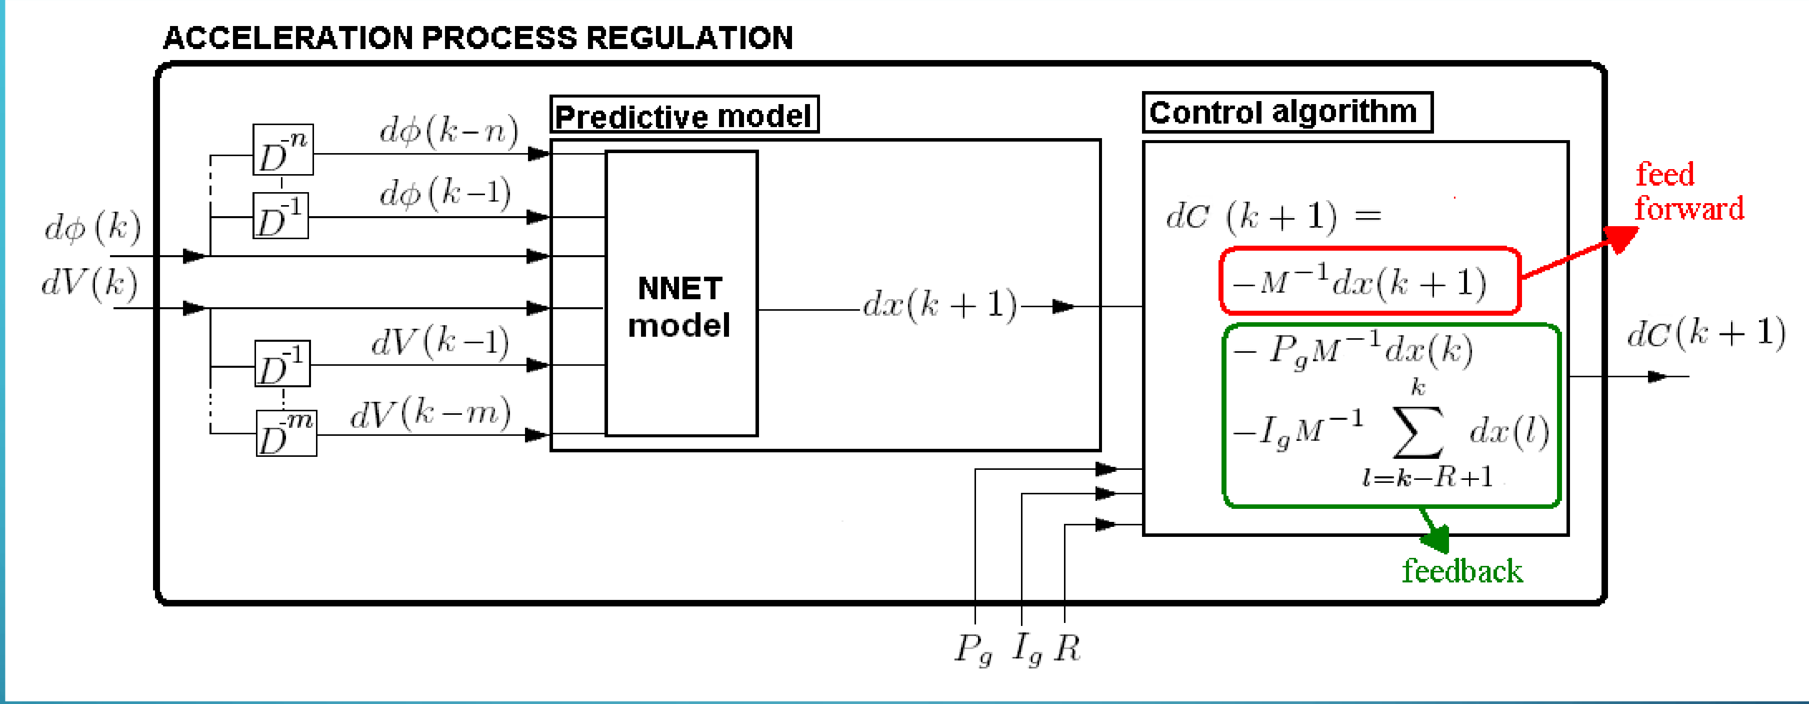

- **Predictive model (left block):** The Neural Network (NNET) receives lagged values of the perturbed klystron phase and voltage. It gives a prediction of the next pulse position deviation *dx(k+1)* used to compute the feed forward correction.
- **Control algorithm (right block):** The algorithm is composed of a feed forward (first term) augmented by PI control terms (second and third terms) to compute the correction to apply *dC(k+1)* .

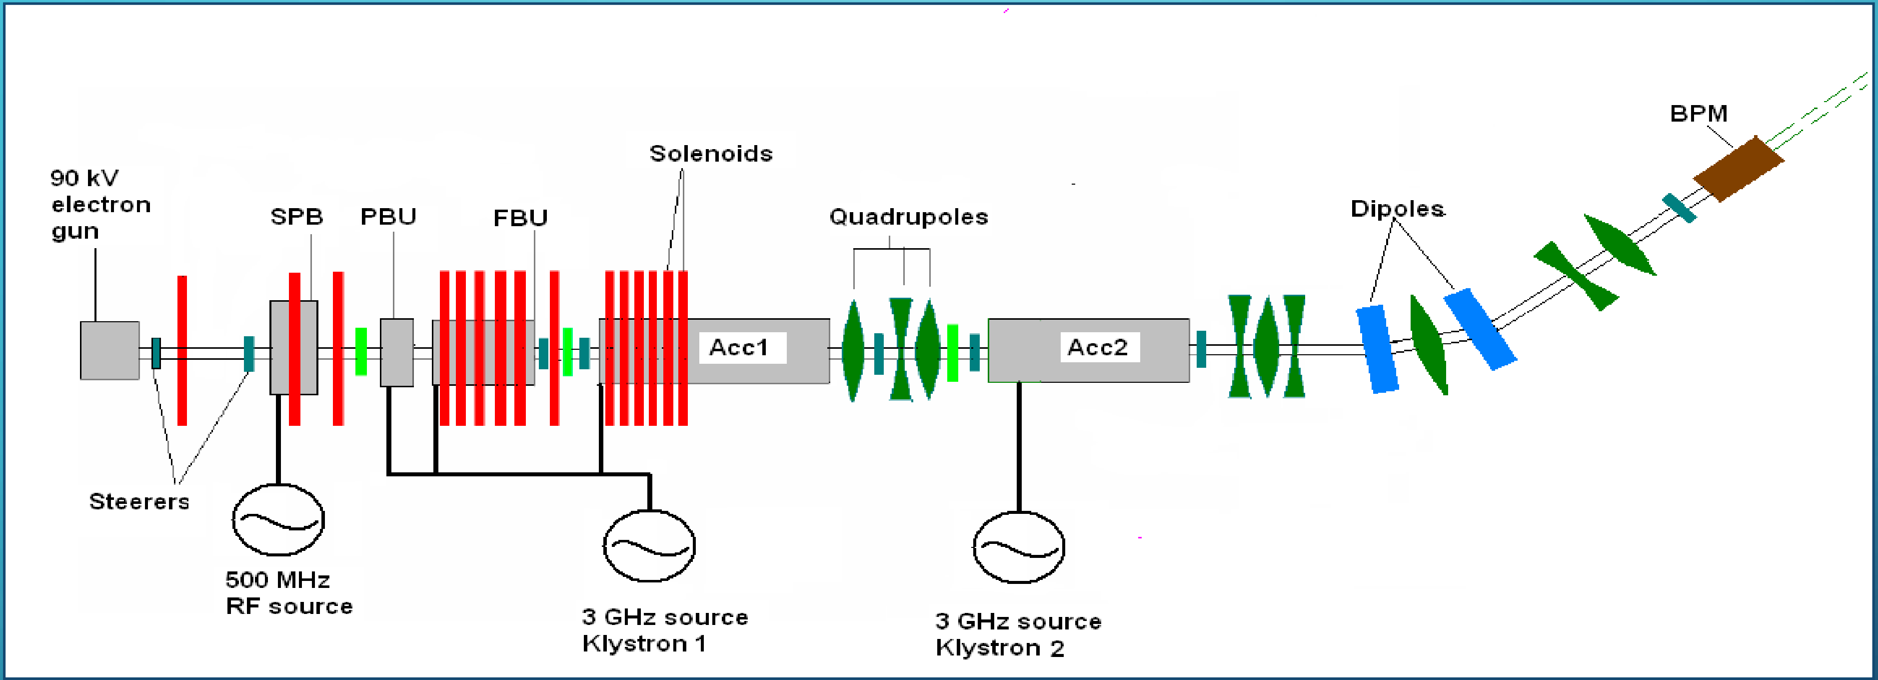

- ACC1 and ACC2 provide 100 MeV beam @ 1-10Hz.
- BPM resolution ~50µm , white noise level ~0.11 mm rms
- Jitter are induced in klystron 1 phase and voltage and corrected using klystron 2 voltage
- Slow actuator response limits the experiment to max. ~0.075 Hz

REAL TIME ENERGY CONTROL

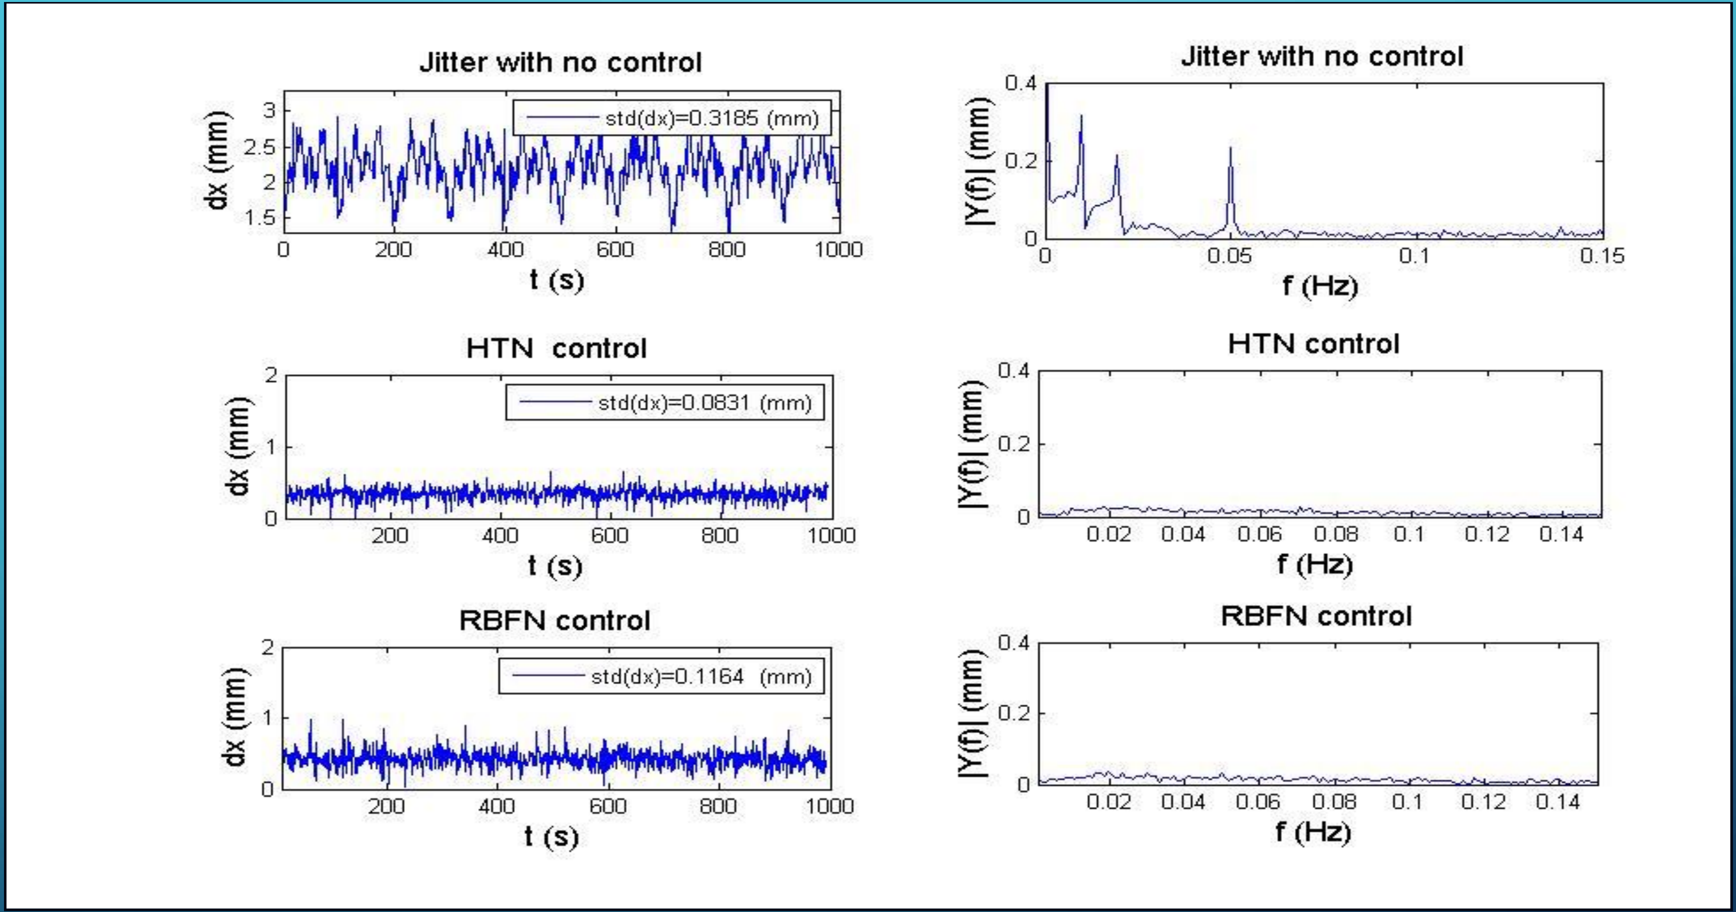

Both the hyperbolic tangent network (HTN) and radial basis function network (RBFN) received 7 lagged values of V_1 and 5 lagged values of ϕ_1. The HTN and RBFN had 7 and 76 hidden neurons, respectively.

Real Time energy feed-forward feedback control.

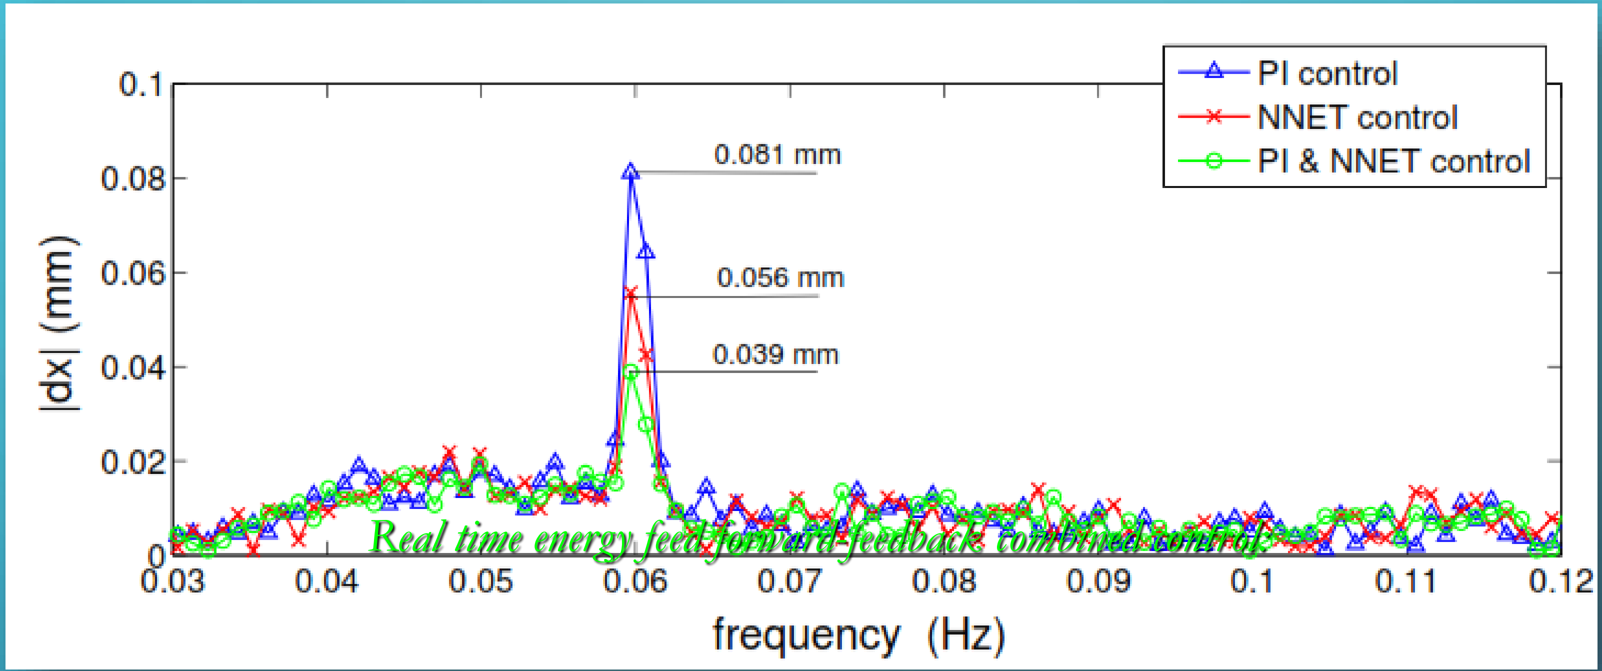

- The PI controller was tuned for a 0.4 Hz perturbation.
- The NNet was trained to correct the same 0.4 Hz perturbation .
- A shift in frequency to 0.6 Hz shows that the NNet operates better than the PI, and that the combined controlled provides further decrease of the perturbation amplitude.

E. Meier, M. Morgan, S.G. Biedron, G. LeBlanc, J. Wu, 2009, "Development of a Combined Feed Forward- Feedback System for an Electron Linac," Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 609, 79-88.