Predicting preventable Interruptions to Operations at NSLS-II

In collaboration with SLAC, ANL

Reid Smith : Speaker

Feng Bai : Postdoc researcher

Mar.07.2024

IFCA Beam Dynamics Workshop - MLAPA

@BrookhavenLab

Contents

NSLS-II Accelerator Overview

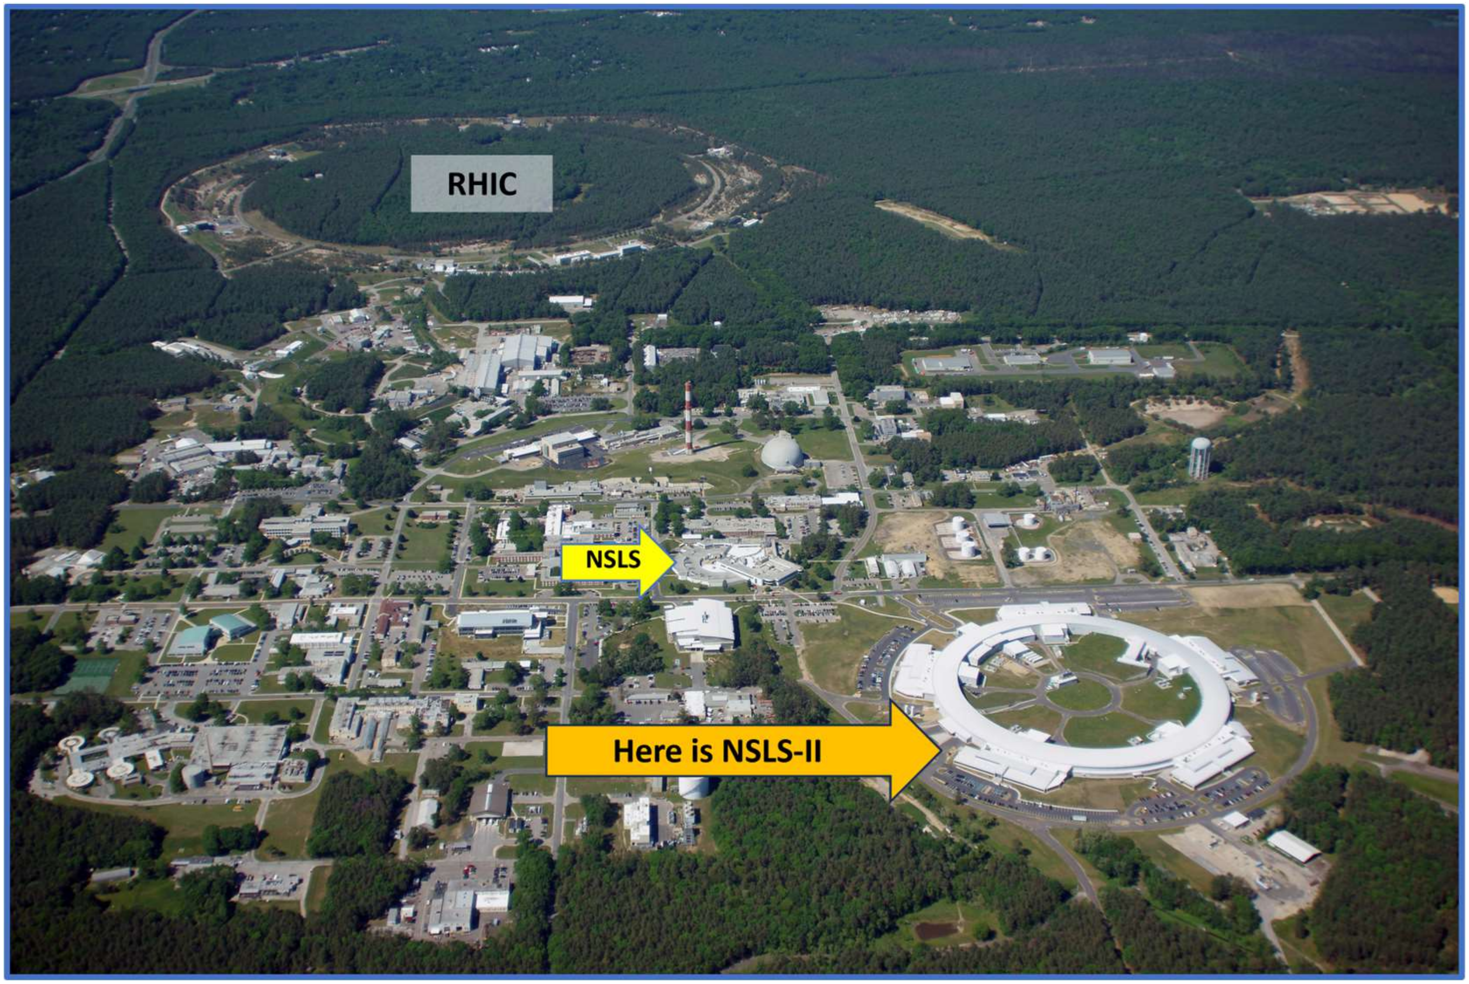

Brookhaven National Laboratory

NSLS-II Accelerator Overview

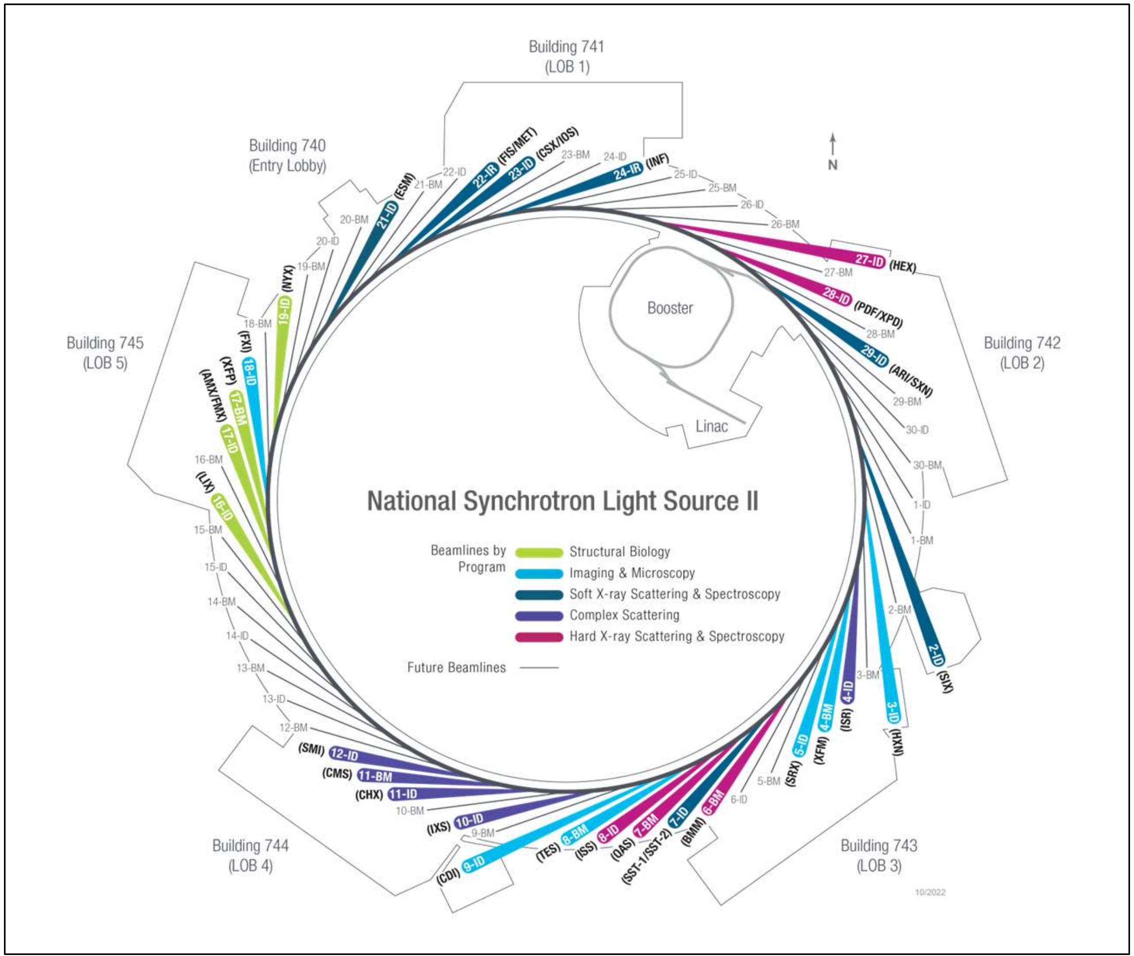

STORAGE RING PARAMETERS

Ring circumference		792 m (.5 mile)
Ring Energy		3 Gev
Ring Current	500 mA	
# Cells		30
Cell design type		Double-Bend Achromat
Vertical emittance		.008 nm-rad
Horizontal emittance		.55 nm-rad
Time between bunches		2 ns
Revolution period		2.64 us
RF frequency	500 mHz	
#RF Buckets	1320	
# Bunches		1056
Active beamlines		31

- Complex scattering 5
- Hard X-ray scattering & spectroscopy 6
- Soft X-Ray scattering & spectroscopy 12
- Structural biology 5
- Imaging & Microscopy 6

Brookhaven National Laboratory

Accelerator Reliability

Reliability is % uptime when in Scheduled Ops.

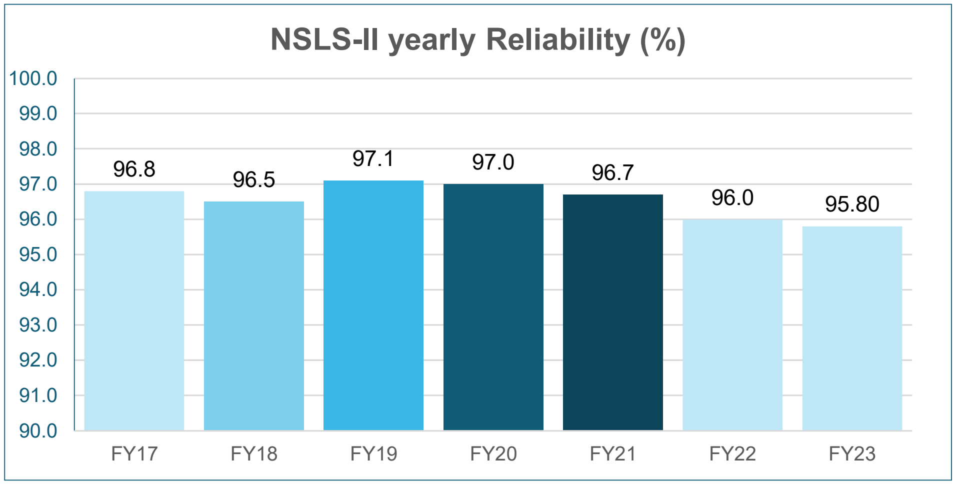

- Accelerator Reliability goal defined by DOE as 90%
- NSLS-II Internal goal of 95%
- NSLS2 regularly runs between 95-97% reliability

…but 5% of 4800 hrs is still 240 hrs of downtime

- Even in a reliable machine, downtime is a tangible burden on users.

Where can we reasonably gain in reliability?:

- Fast Faults -
  - Power dips, power supply trips, network outages, vacuum spikes, RF trips, cryo quenches, etc.
  - **Not especially predictable – actionable by system improvement, or reducing trip-recovery**
- Slow Faults
  - Temperature, pressure, flow rate, ground current, etc.
  - Trip usually initiated by Equipment Protection systems
  - **Actionable when caught early enough**

How early?

One 2-day maintenance & one 1-day interlocks period w weekend studies, repeating every 3 weeks.

10 days minimum, 20 days preferred

Magnet Temps & Cooling water system

Issue Summary: local clogs in magnet cooling-water system sporadically cause heating issues with magnets. Seen most in QM magnets.

- Normal temps ~35C; Klixons will dump beam between 70-80C unless overridden; Experts will abort at 80C.
- Overheating can cause permanent damage to magnets – leading long & costly replacement.
- Not the first time we've dealt with this issue

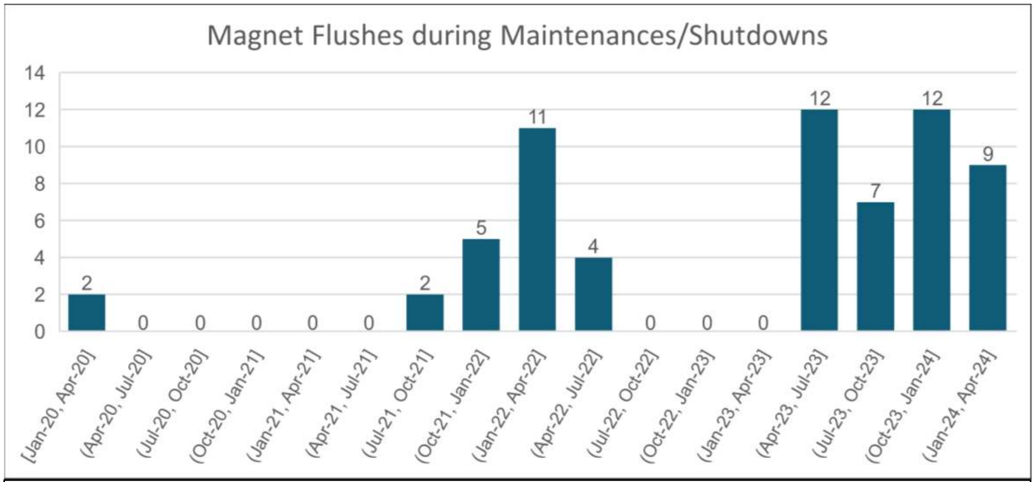

# of magnet cooling-water flushes, 3-month bins

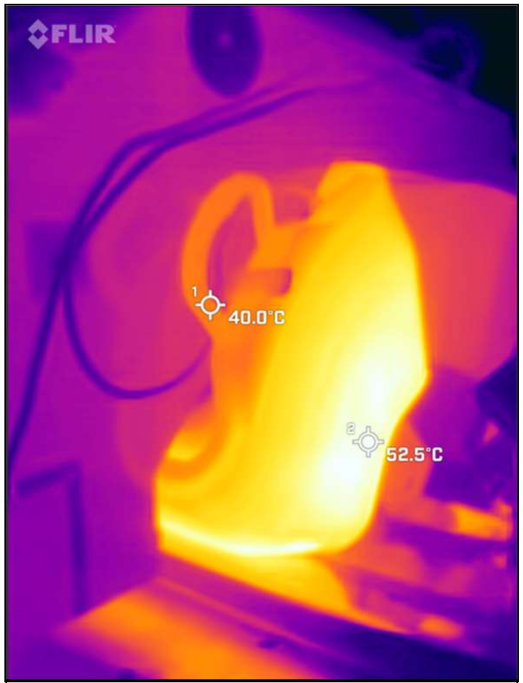

Thermal camera image of a hot magnet.

Magnet Temps: the run-up to Predictive Programs

Monitors:

1. Plots (visual inspection)
2. Alarms
3. Calculating change over time (drift)
4. Programmatic prediction…

1

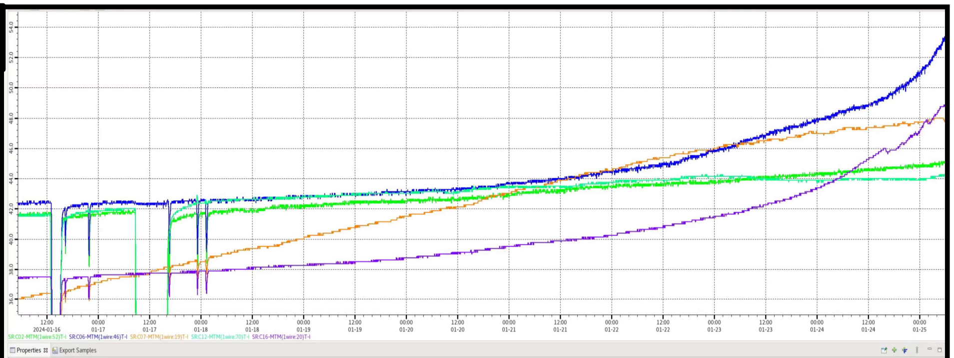

Plot of quadrupole magnet temperatures in several ‘bad’ spots around the ring; rising over 10 days in January 2024.

2

C01	C02	C03	C04	C05	C06	C07	C08	C09	C10
36.8	45.0	38.8	37.1	39.2	53.5	48.0	38.2	42.4	39.6
C11	**C12**	**C13**	**C14**	**C15**	**C16**	**C17**	**C18**	**C19**	**C20**
37.8	44.2	37.8	38.5	37.5	48.9	37.9	41.2	38.1	37.5
C21	**C22**	**C23**	**C24**	**C25**	**C26**	**C27**	**C28**	**C29**	**C30**
40.1	39.4	37.0	37.0	36.8	35.6	39.5	37.1	38.4	41.0

Summary-alarm values of Magnet RTD temperatures. Alarms sound in the accelerator Control Room, email experts, etc.

3

	C16	C17	C18	C19	C20	C21	C22	C23	C24	C25	C26	C27	C28	C29	C30
1	-0.1	-23.4	-0.3			-0.6		-16.1	-57.7				-10.7		-3.8
2	0.0	-28.5	-0.4			-11.6		-17.2	-45.4				-10.2		-11.4
3	-10.2	-7.2	-9.0			-3.5		-1.5	-18.6				-7.4		-3.3
4	-10.9	-15.6	-18.5			-11.4		-0.4	-3.7				-19.5		-3.7
5	-4.1	-14.7	-13.6			-8.9		-31.5	-14.1				-13.1		-5.5
6	-4.0	-10.5	-9.5			-11.9		-24.6	-21.9				-1.3		-2.9
7	-9.9	-0.5	-8.7	-23.0	-24.7	-7.3	-16.0	-2.9	-23.1	-14.8	-9.1	-11.9	-22.4	-20.0	-11.3
8	-9.9	-0.5	-7.9	-26.0	-23.8	-0.3	-16.5	-0.9	-24.9	-16.0	-9.3	-12.1	-22.7	-20.8	-5.4
9	-1.8	-13.7	-0.4	-30.9	-1.8	-3.7	-2.6	-1.9	-1.5	-1.3	-10.3	-1.5	0.1	-9.8	-1.7
10	-2.2	-16.7	-4.7	-37.0	-4.4	-3.4	-3.8	-0.2	-3.8	-3.6	-8.5	-3.1	-4.2	-4.8	-4.7
11	-12.4	-16.9	-0.5	-35.1	-2.1	-1.7	-1.9	-3.4	-2.9	-2.9	-10.5	-2.5	-3.0	-5.4	-2.1
12	-0.2	-25.2	-8.2	-10.7	-3.6	-7.5	-1.1	0.0	-3.1	-2.8	-8.5	-3.0	-12.1	-12.4	-3.6

Cell 01 – RTD 07

Live Temp	**25.5**	**Drift**	**-5.8 deg C**	(Ref Beam I)	399.44 mA
Warning Temp	80.0	Ref Temp	**31.3 deg C**	Save New Reference	
Trip Temp	100.0	(Ref Date)	09/17/2018		

Drift monitor shows difference between reference temp & current temp – colors indicate size of delta

Brookhaven National Laboratory

Development Philosophy

Machine Learning Operations

The name MLOps incorporates machine learning (ML), software development (DEV) and operations (OPS).

MLOps = ML + DEV + OPS

Machine Learning Operation is designed to incorporate machine learning engineering and DevOps to seek for data governance and deployment in production with reliability and scalability.

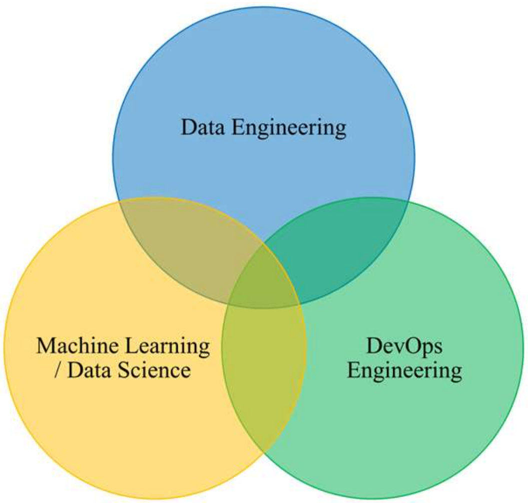

Brookhaven National Laboratory

Slide content provided by Feng Bai

Data Structure for QM magnets

Data for all QM magnets around the ring handled in 4 layers:

1. Temperature data as a list of HDF5 files
2. 30 Cells in NSLS-II storage ring.
3. 4 sensors/magnet * 4-5 QM/cell = 20-24 PVs per cell
4. Temp data time series (288 meas/day @ 5 minute sampling rate) for each sensor

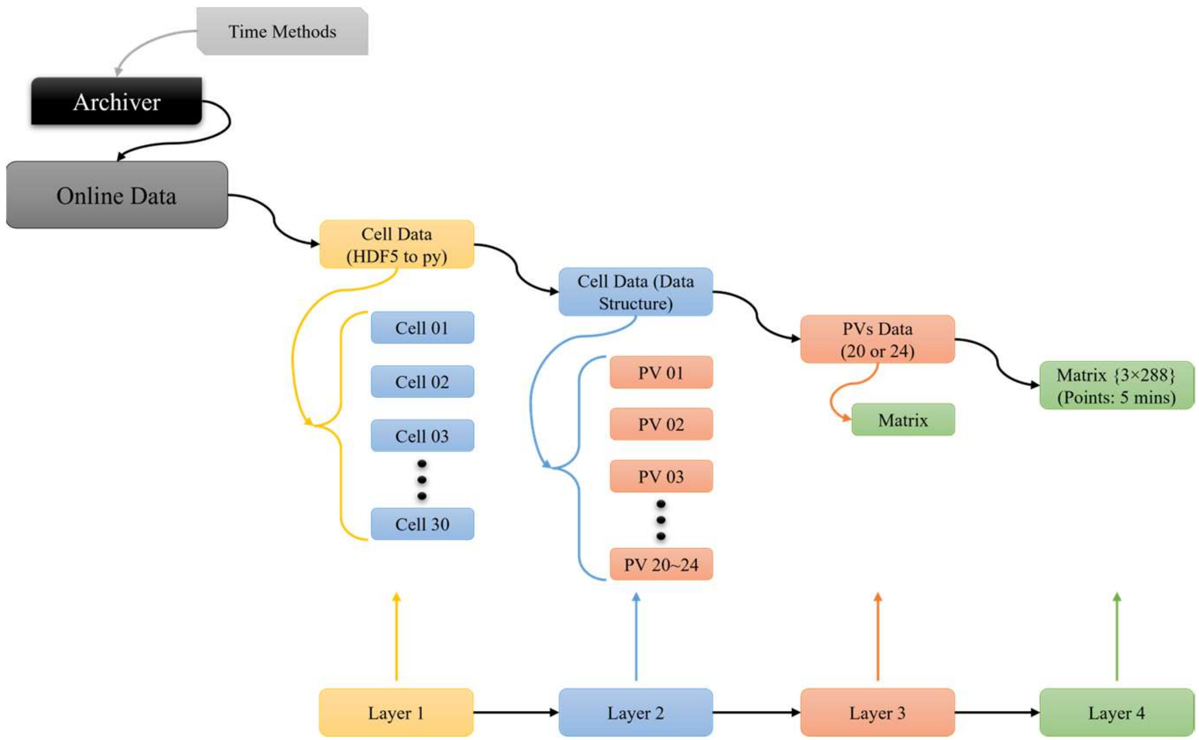

Slide content provided by Feng Bai

Data Structure for QM magnets

Data Engineering and Storage Systems

Q: What to do about asynchronous triggered data?

A: data sampled from archiver at fixed interval (5 min by default)

Data Engineering includes:

1.) Data acquisition based on the network paths from online sensors.

2.) Raw data written into HDF5 files.

3.) Data transfer by reorganizing the raw data into self-structured data in .py files (we reorganize the sensor data into 5 minutes per point in time series)

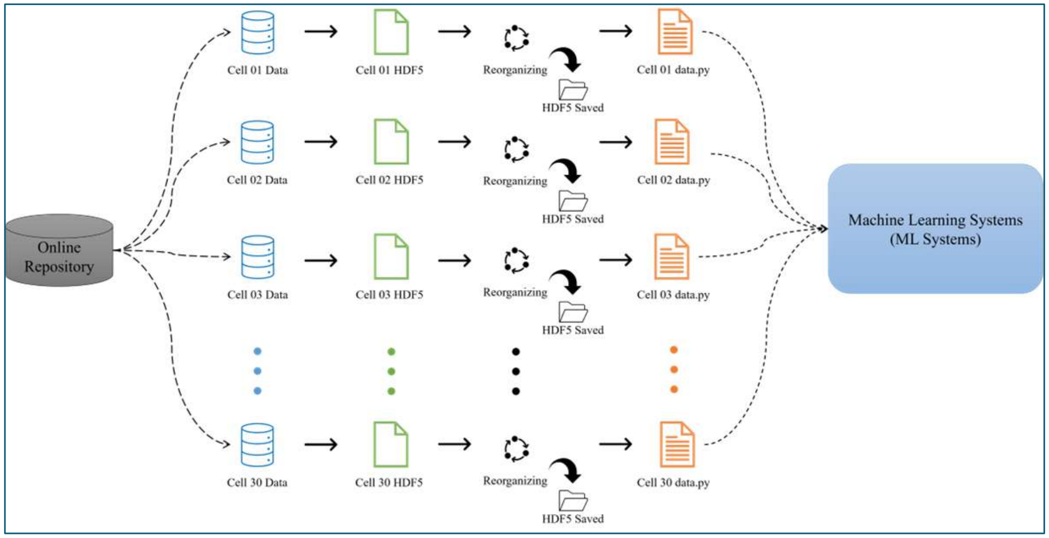

Slide content provided by Feng Bai

Data Structure for QM magnets

Machine Learning Systems (ML systems)

The input temperature data is retrieved from the archiver appliance, organized, processed, and the outputs are written to files and PVs.

There are three components in the ML systems that are parallelized:

(1) Data Cleaning
(2) Model/Learning
(3) Prediction

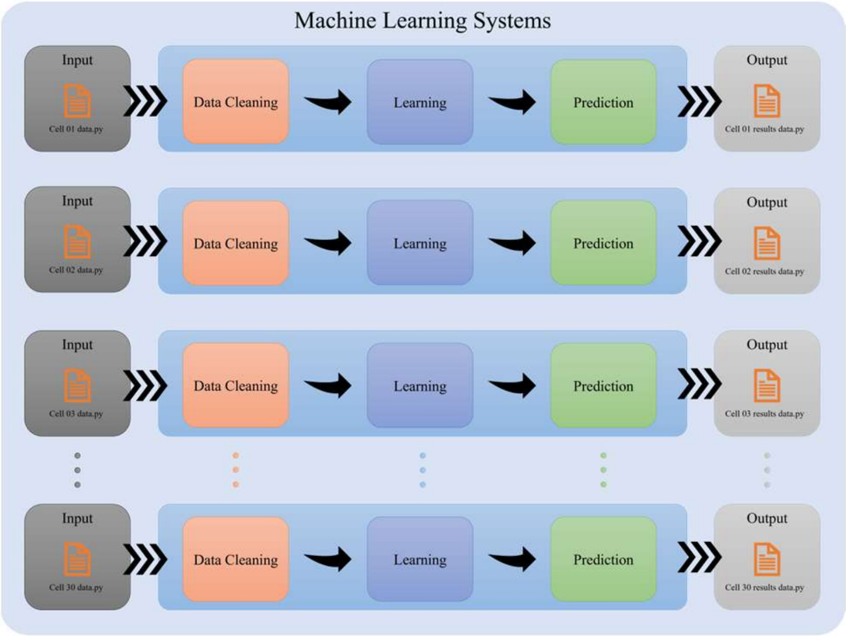

Brookhaven National Laboratory

Slide content provided by Feng Bai

1. Data Cleaning

Filtering out outliers from sensor faults using isolation forest algorithm (over Euclid norms)

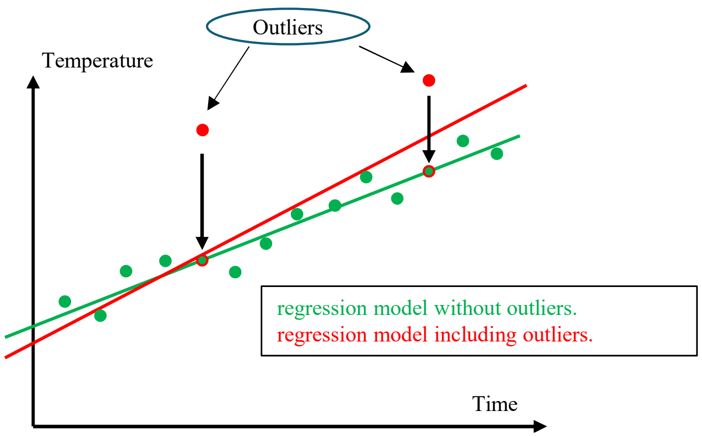

Outlier Detection (Linear regression case)

Adjusting sensitivity to account for jitter “created by” sensor resolution
(0.125 C for our 1-wire sensors)

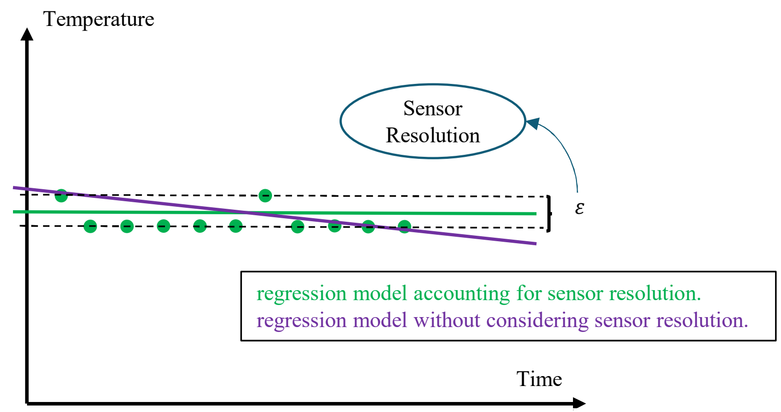

Sensor resolution (Linear regression case)

Slide content provided by Feng Bai

Methods

2. Modeling & Regression

We use time-series data per PV make temperature prediction in the pipelines. We tried several regression methods and deep learning methods for online predictions.

1.) Linear Regression;

2.) Quadratic Regression;

3.) Nonlinear Exponential Regression;

4.) LSTM.

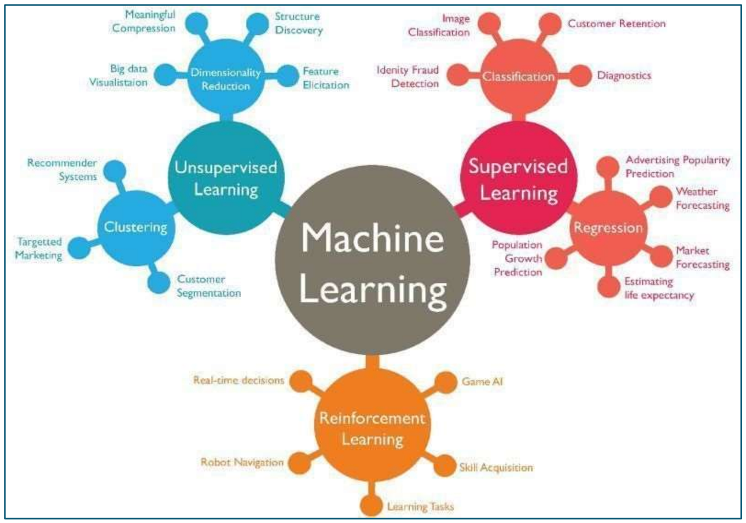

3. Prediction

The newly trained ML models are used to model future behavior and predict temperatures at points in future.

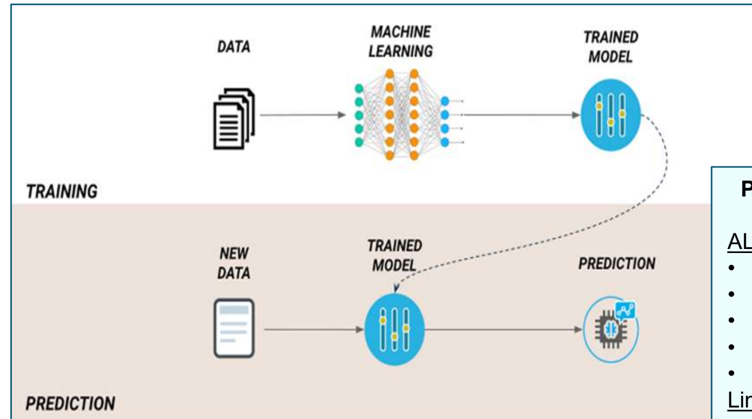

Predictions/Parameters

ALL models:
- Days to 40 C
- T 3 days
- T 7 days
- T 14 days
- T 30 days

Linear Fit:
- Slope
- Least sq error

Quadratic Fit:
- Square coeff
- Linear coeff
- Least sq error

LSTM:
- <shape parameters>
- <error parameter(s)>

Brookhaven National Laboratory

Slide content provided by Feng Bai

(Tested but Unused Prediction Models)

Piecewise Linear Model

Dataset of temperature for every day:

$$D \coloneqq \{(t_1, T_1), \dots, (t_N, T_N)\}$$
$$N = 12 * 24 = 288$$

D: daily dataset;
N: # of daily data inputs (5 mins);
n: number of days;

The regression done every day with N data inputs;

The piecewise linear models can capture the daily parameter θ_n (slopes)

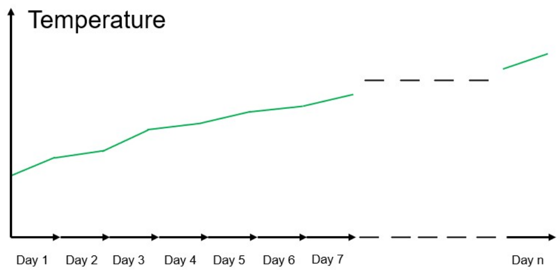

Overlap Piecewise Linear Model

Dataset of temperature for every day:

$$D \coloneqq \{(t_1, T_1), \dots, (t_N, T_N)\}$$
$$N = 12 * 24 = 288$$

D: daily dataset;
N: # of daily data inputs (5 mins);
n: number of days;

The regression done every 12hs with N data inputs (including the dash lines);

The piecewise linear models can capture the daily parameter $\theta_{2n\,-\,1}$ (slopes)

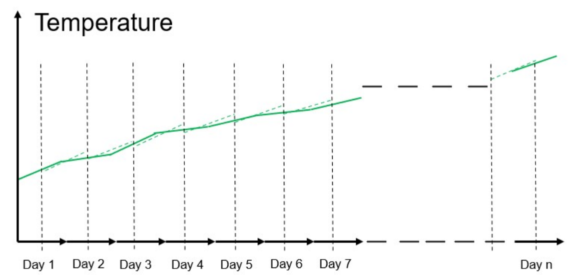

Linear Models, Lasso/ Ridge

1. Lasso $\min_{\theta \in \Re} \frac{1}{N} ||T - X\theta|| + \lambda ||\theta||_1$ → Strong Regularization
   L1 Norm

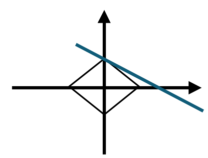

2. Ridge $\min_{\theta \in \Re} \frac{1}{N} ||T - X\theta|| + \lambda ||\theta||_2$ → Smooth Regularization
   L2 Norm

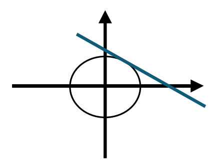

Strong Regularization can feature the small parameters with the L1 norm distance to 0. However, Smooth Regularization just feature the characteristics of parameters gradually.
X: data inputs (daily time series temps), θ: output parameter

Brookhaven National Laboratory

Slide content provided by Feng Bai

(Tested but Unused Prediction Models)

Machine Learning Methods: LSTM

The data structures in the LSTM are shown in distributed pipelines per PV, per cell.
The data in the training are displayed in tensor networks with the given period to form the “time window”.

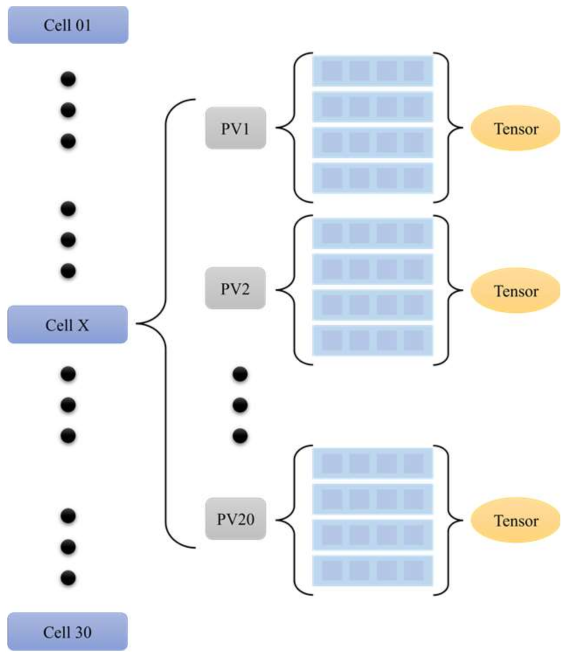

Brookhaven National Laboratory

Slide content provided by Feng Bai

Results of Fitting

- Shown: results of Linear/least-squares and Lasso models over a 10-day span.
- In production: training and prediction periods are independent.

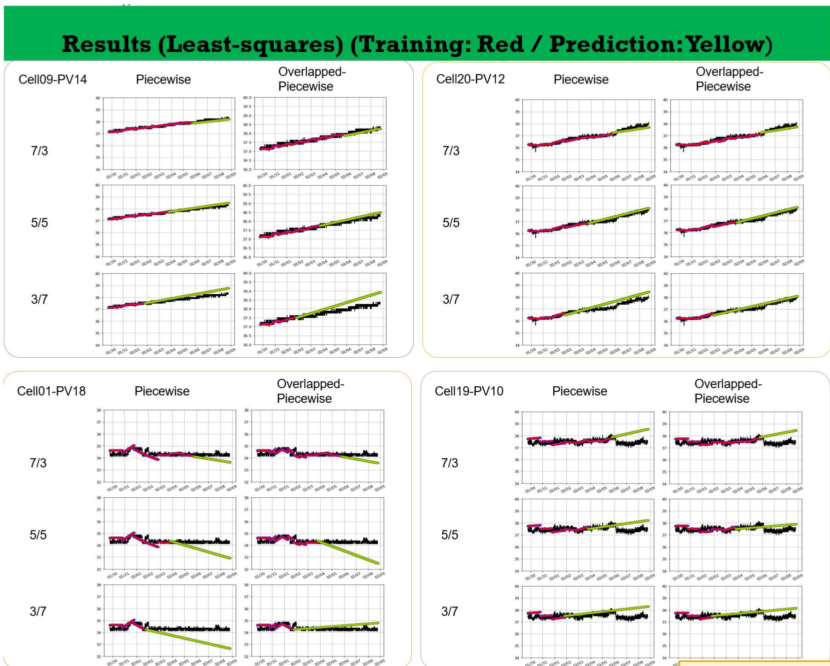

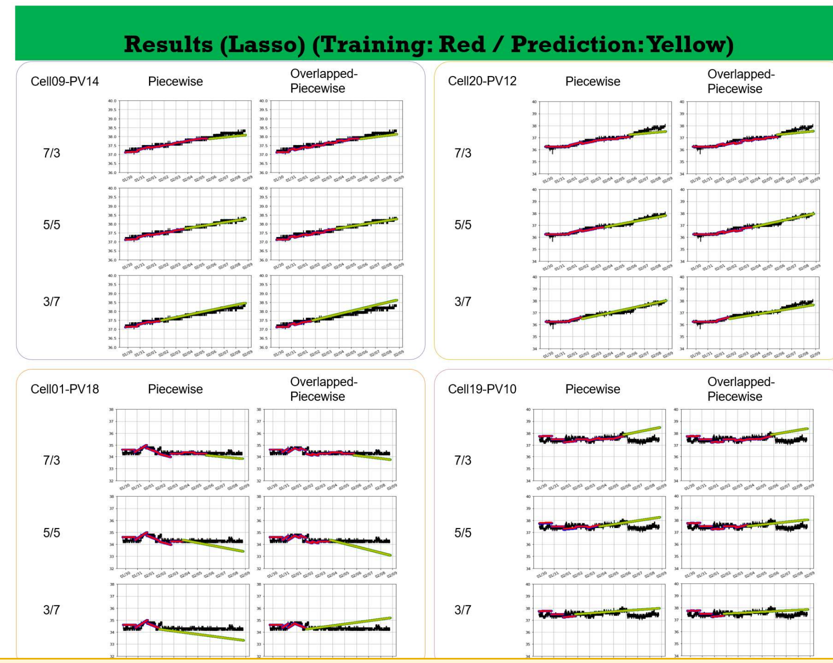

- Better predictions with longer fitting dataset, EXCEPT:
- Not able to deal with abnormalities in training set

Brookhaven National Laboratory

Slide content provided by Feng Bai

Predictions & Displays

Predictions/Parameters

<u>Both models:</u>
- Days to 40 C
- T in 3 days
- T in 7 days
- T in 14 days
- T in 30 days

<u>Linear Fit:</u>
- Slope
- Least sq error

<u>LSTM:</u>
- <shape parameters>
- <error parameter(s)>

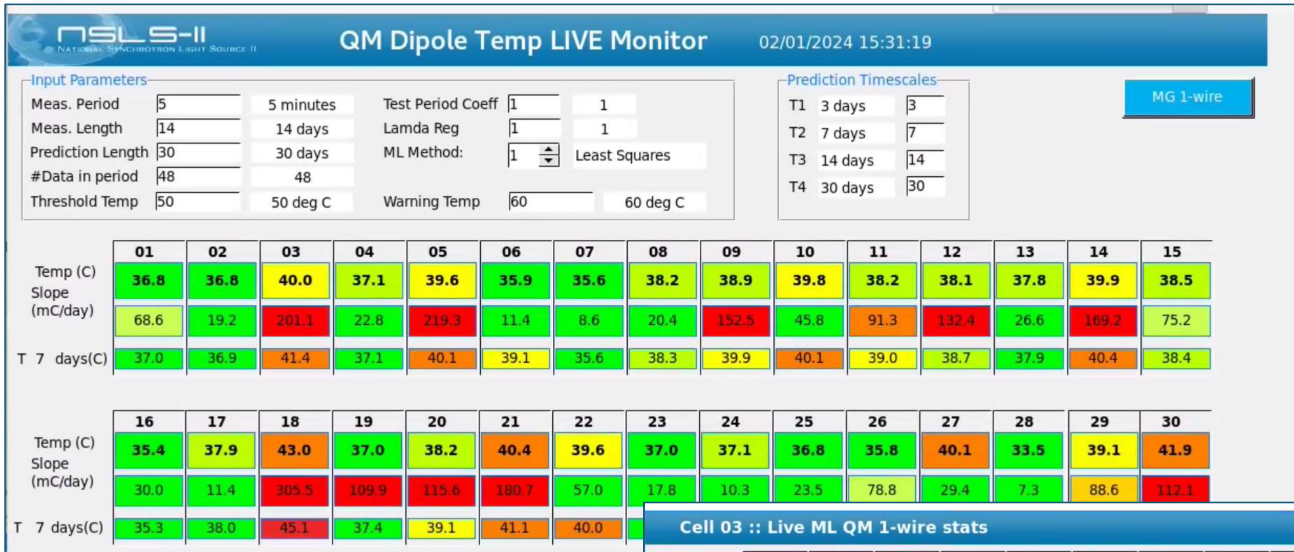

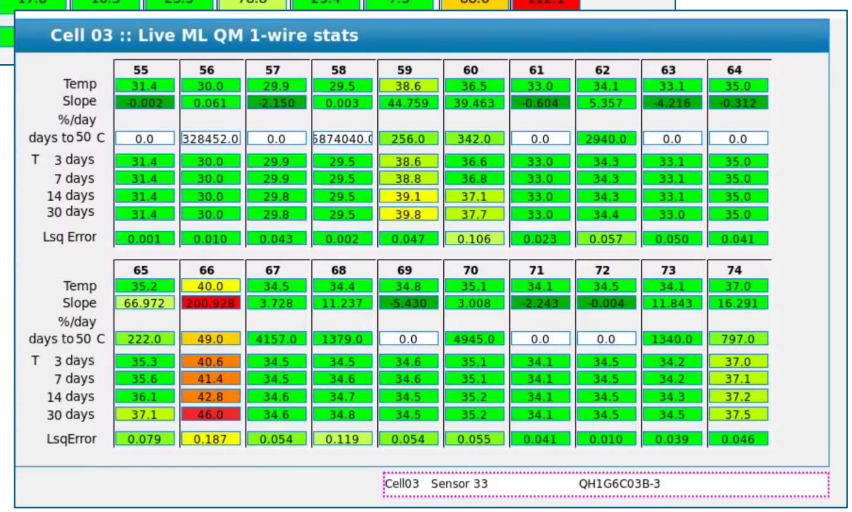

Using linear fit prediction right now.
Still testing learning-model fitting and prediction.

Brookhaven National Laboratory

Use by Experts

Presently, EE & Mechanical groups are watching hot magnets.

Slope, temp predictions, and ‘time to threshold’ PVs are directing and replacing the same calculations being done by hand or in Excel, with much less time committed.

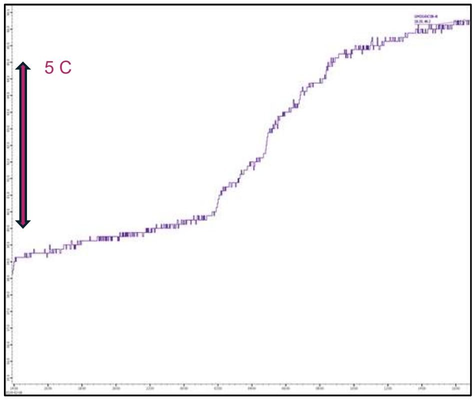

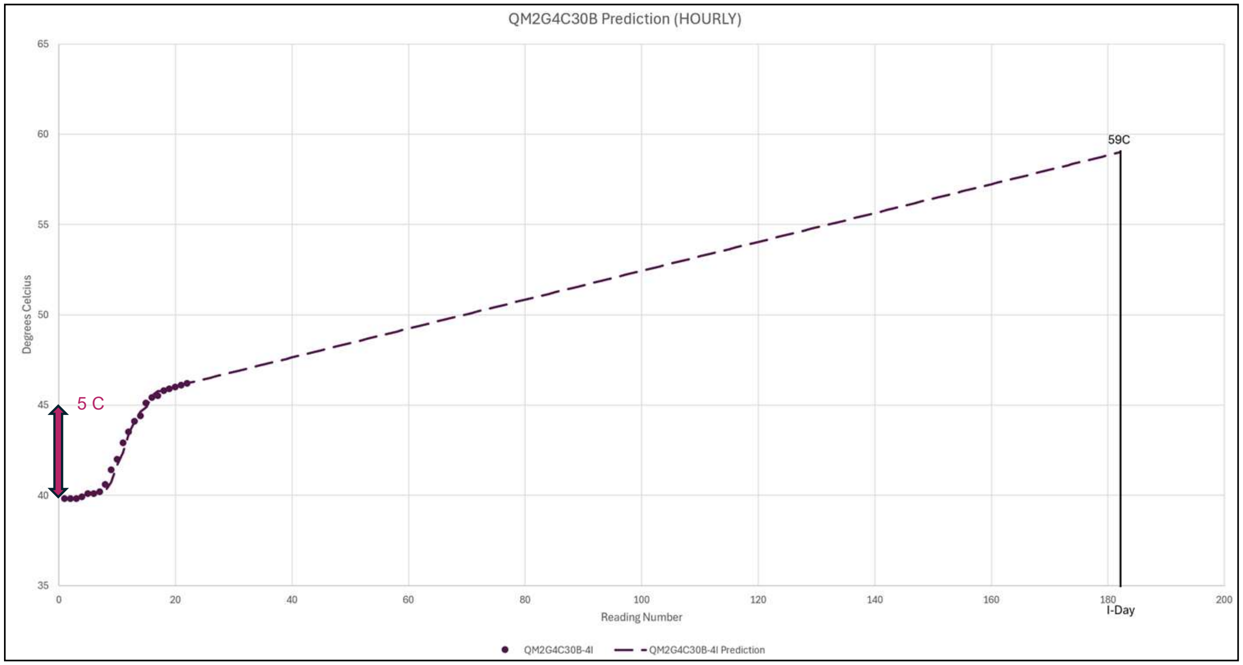

Conclusions & Next Steps

Strengths and limitations

+
- Slope is a VERY helpful early warning.
- Linear prediction catches magnets warming from ‘normal baseline’.
- Changing time windows & thresholds can help adapt to situations.
- Python; able to add features

-
- Fitting models deal poorly with fast changes (repairs, clogs);
- Even with perfect modeling, ‘sudden temp jumps’ in modeling time-window will result in bad predictions
- LSTM learning model needs more data (both length & type) to be worth using

Next steps & Future applications

Next Steps
- Expand ML model to use years of history, & other magnet/water readbacks
- Decide if all QM temp datasets can be treated (& trained) interchangeably
- Consider multivariable Transformer vs LSTM
- Can we make a model to deal with sudden temp rise/drop separately?

Applied to QM Magnets
- Need to expand scope to make learning model worthwhile:
    - If ML is wanted, train using GPU cluster (for parallel processing)
- Make the display easier for experts to use/scan:

Future Systems
- Power supply ground currents
- Possible use in RF and/or Cryo

Brookhaven National Laboratory

Thanks to:

BNL: Feng Bai, Guimei Wang, Yoshi Hidaka, Jinhyuk Choi,
Wing Louie, Joseph Gormley, Brian Walsh, Michael Charumaneeroj
SLAC: Daniel Ratner, Finn O'Shea, Anwesha Das
ANL: Louis Emery, Nikita Kuklev, Ihar Lobach

Mar.07.2024

Brookhaven National Laboratory