#### NEI View on Artificial Intelligence

James Slider Technical Advisor, Innovation & Regulatory Affairs

November 20, 2024





#### Industry Adoption of Al

- Operating plants are approaching artificial intelligence cautiously; learning what is real about AI takes research, experimentation, and time
- Pilot projects focus on a clearly defined problem
- Typical pilot seeks to streamline a repetitive process and reduce human drudgery and labor hours in that process
- Costs of building AI infrastructure (data acquisition and curation, model development and training, user interface, maintenance) can be significant
- Many choose on-premises designs to protect company IP

#### NEI Member Interests in Artificial Intelligence



| Use Category                 | Now  | Near    | Future |
|------------------------------|------|---------|--------|
| Work Management              | Few  | Many    | All    |
| Corrective Action Management | Many | Most    | All    |
| Maintenance Management       | Many | Many    | All    |
| Operations Support           | Few  | Several | All    |
| Outage Management            | Few  |         | All    |
| Regulatory Affairs           | Few  | Many    | All    |

#### Standards and Best Practices



- The standards community is challenged to keep up with the pace of AI development
- Many US and international standards pertaining to Al are already available; the NRC's gap analysis report provides a valuable catalog of those standards
- More are in development
- Legislation (e.g., the EU Al Act) and government-wide direction (e.g., Executive Orders) can impose additional requirements
- Corporate policies impose additional requirements on licensee Al projects, such as data governance
- Best practices are in early stages of development

#### Recommendations



- Recognize that AI is evolving more rapidly than regulations can keep up; look for ways to speed up and streamline development of guidance
- Rely on the existing regulatory framework as much as possible
- Interact through the use of table tops and pilot tests in the revision or creation of requirements for AI (e.g., Reactor Oversight Process)
- Strive to keep the regulatory touch commensurate with impact on safety to allow for continued rapid innovation
- Allow flexibility and credit for use of industry-developed guidance

#### Final Thoughts

- Education of the public about AI use in nuclear facilities will be important to maintain public trust.
- We appreciate the NRC's engagement thus far.
- We look forward to working with the NRC to find an appropriate balance between regulatory oversight and creative freedom to enable NRC and industry to innovate with AI tools to improve safety, efficiency, and quality of work life.



### Model-Agnostic Explainability Methods for Trustworthy AI in Nuclear Reactor Applications

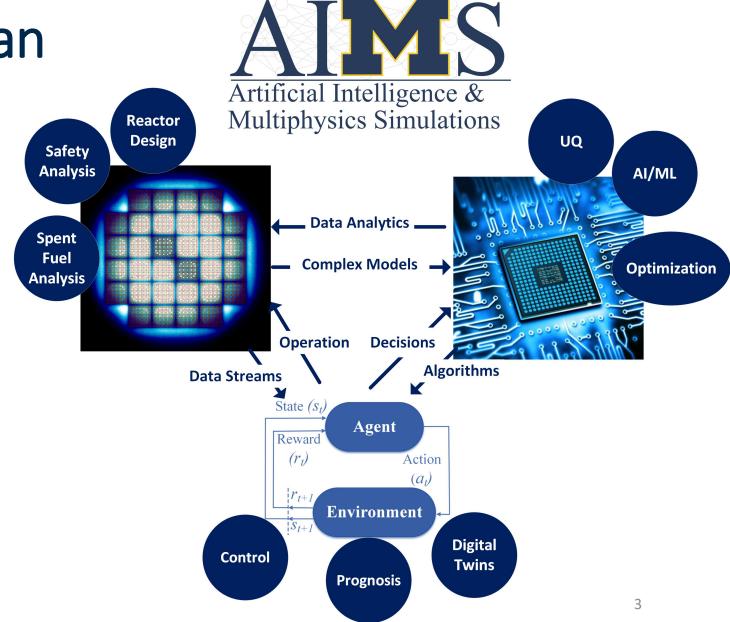
U. S. NRC Advisory Committee on Reactor Safeguards
Joint Human Factors Reliability & PRA, and Digital I&C Subcommittee Meeting

PI: Prof. Majdi I. Radaideh (Assistant Professor, <u>radaideh@umich.edu</u>) Student: Nataly Panczyk (PhD Student, University of Michigan)



## AIMS @ U of Michigan

- Established in January 2023.
- Students' background: Nuclear engineering, Computer Science, Mathematics/statistics, Software engineering.
- Fall 2024 roster (22 members): 2 postdocs, 9 PhD students, 5 Masters students, 6 undergraduate students.
- PI: Majdi Radaideh (Assistant Professor)



## **AIMS Ongoing Projects**



NLP and LLM for sentiment analysis and public engagement

Autonomous control in microreactors: Classical control vs Reinforcement Learning

Mathematical Foundations and Computational Implementation for Nuclear Digital Twins

Tensor train methods for accelerated neutron transport solutions

Low-to-high fidelity mapping with ML for accelerated CFD simulations for advanced reactor

Robotics and drones for autonomous inspection and radiation mapping

Fast multi-objective optimization algorithms to enable reactor design in expensive simulations.

Computer vision for stress-corrosion cracking detection and tracking in LWR fuels

AI/ML framework development, regulation and licensing in nuclear power

















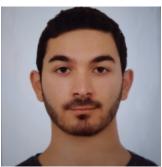


## Acknowledgement

- Work done mostly by the AIMS
   Reactors sub-group (Nataly Panczyk,
   Patrick Myers, Omer Erdem, and Leo
   Tunkle).
- Most of the topics presented were sponsored by the NRC University R&D Grant Program.
- Some of the topics presented were sponsored by the DOE Distinguished Early Career Program.















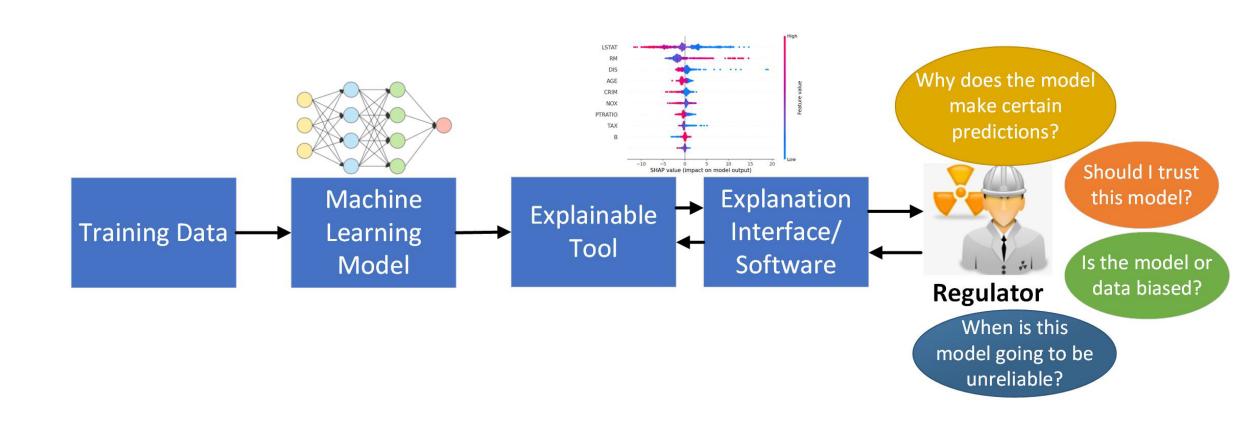
#### Content

- Background
- What is explainable AI?
- Our proposed AI regulatory framework?
- Preliminary results
- Group discussion

## Al Regulatory Gap

- If some vendor came in and said they want to use some nontraditional model to control/model a process in a nuclear facility
  - e.g., equipment degradation, predictive maintenance, control rod movement, non-safety process
- They give the NRC the model they built, the routine they used to build the model, and a set of testing scenarios of the model under different conditions
  - Here we assume the data are not shared with the NRC.
- "As a regulator, what do you want to know about this model to license it?"

## Simplified Flowchart



#### Content

- Project Background
- What is explainable AI?
- Al Applications (without explainability)
- Our proposed AI regulatory framework?
- Preliminary results
- Group discussion

## Explainability vs Interpretability in Al

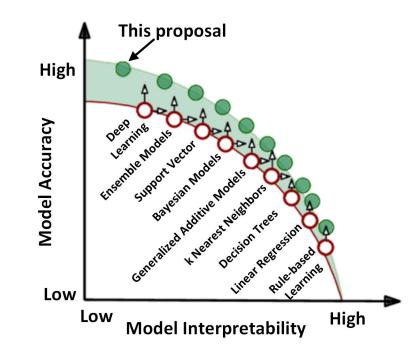
- Interpretability and explainability are often used interchangeably, but they have nuanced differences.
- Interpretability refers to the ability to understand how a model works internally. It means that the model itself is simple enough for a human to understand how it makes decisions.
- Explainability refers to the ability to explain the decision-making process of a model, even if the model itself is complex and not inherently interpretable.

| Aspect              | Interpretability  | Explainability  |
|---------------------|---|---|
| Model<br>Complexity | Typically simple models (e.g., linear models, decision trees) | Can be applied to complex models (e.g., neural networks, ensembles)         |
| Understanding       | Direct understanding of how the model works                   | Requires additional tools or<br>methods to explain the model's<br>decisions |
| Transparency        | High (the model itself is transparent)                        | Variable (depends on the method used for explanation)                       |
| Examples            | Linear regression, decision trees                             | Deep neural networks, random forests with LIME or SHAP explanations         |

## Model Interpretability curve

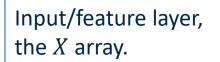
- Rule-based learning refers to the class of methods where you encode every rule/action in the system in if else if ... statements, where you know every possible action/prediction the model could do.
- Classical models include linear regression, decision trees, and K-nearest neighbors.
- Black-box models involve deep neural networks and transformers (i.e., the core of generative AI models).





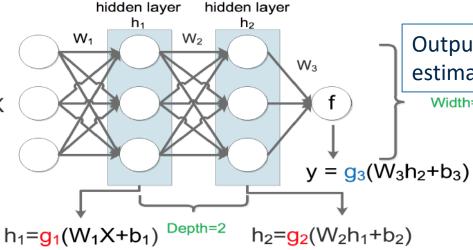
#### Feedforward Neural Networks

$$f^*(x) \approx f(h_2(h_1(x)))$$



input X

 $h_i(.)$ , where i=1,2. Called hidden layers.



Output/label layer, the *Y* array, estimated by function f.

Width=3

 $g_i(.)$ , where i=1,2,3. Called non-linear transformers or activation functions.

#### Weight matrices W might also be called **kernels**

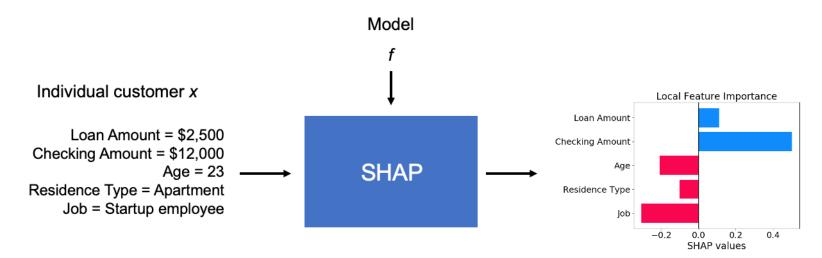
 $W_i$ , where i=1,2,3. Are Weight matrices (to be determined).

 $\theta = (W, b)$ Are usually grouped into one variable called **model** parameters

 $b_i$ , where i=1,2,3. Are network biases (to be determined).

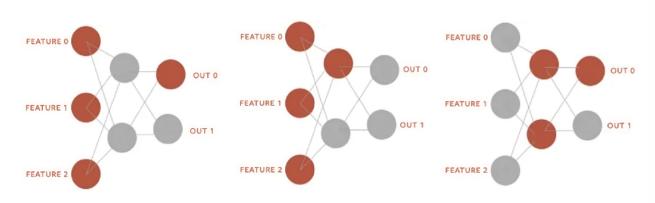
## SHAP (SHapley Additive Explanations)

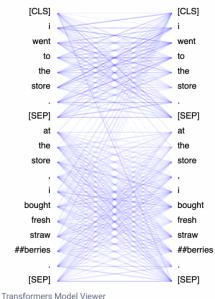
- SHAP assigns each feature (x) a value that represents its contribution to the prediction (y), based on the Shapley value calculation.
  - Theoretically, it considers all possible combinations of features to measure their individual impact on the prediction.
- It is model-agnostic and can be used with neural and non-neural network models.
- Features with positive SHAP values positively impact the prediction, while those with negative values have a negative impact on the prediction. The magnitude is a measure of how strong the effect is.



## End goal: Layer and node attributions

- A more microscopic look on what each node and layer in a complex neural network is doing to the model prediction.
- Offering global input/output model view, layer view, and neuron view.
- Classical methods like SHAP can be repurposed with visualization tools to allow for neuron and layer views.

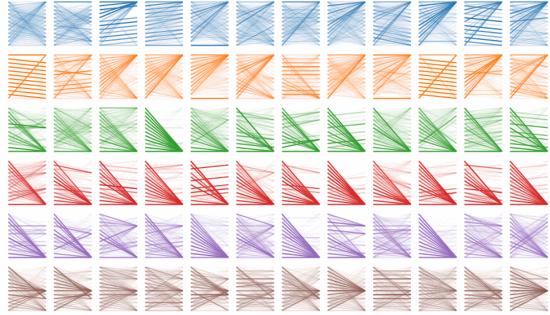




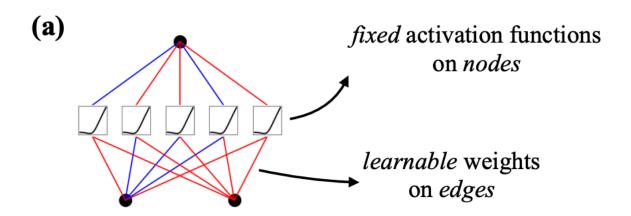
idiisioiiileis wodel viewel

Experiment: ad70f74e3b1349caa67b78c57c519ad4 v

Asset: attn-view-sentiment-Many people dislike Steve Jobs, while acknowledging his genius-POSITIVE-0.98.json >



## Kolmogorov-Arnold Networks (KANs) as a path for interpretable Al



learnable activation functions on edges

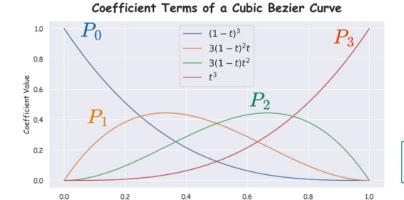
sum operation on nodes

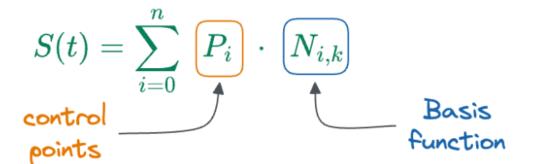
In a neural network, activation functions are fixed across the entire network, and they exist on each **node**.

In KAN, however, activation functions exist on edges (which correspond to the weights of a traditional neural network), and every edge has a different activation function.

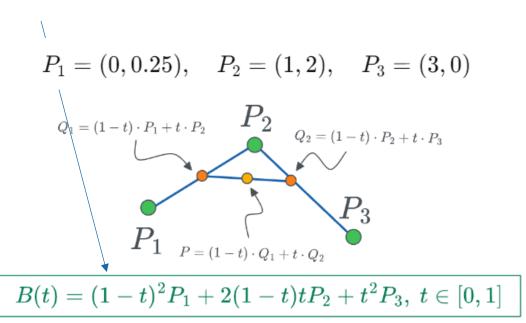
### Control points

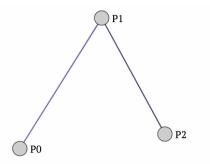
- KAN builds on the concepts of Bezier curves and Bsplines.
- KAN uses the positions of control points learnable in the activation function so that the model is free to learn any arbitrary shape activation function that fits the data best.



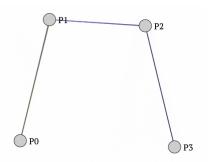


- $P_i$ : Control points that define the shape of the curve.
- $N_{i,k}(t)$ : B-spline basis functions of degree k associated with each control point  $P_i$ ,



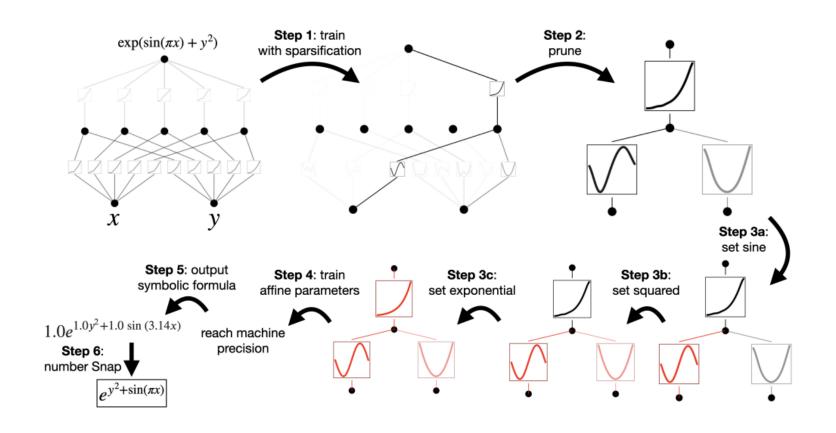


t = 0.00



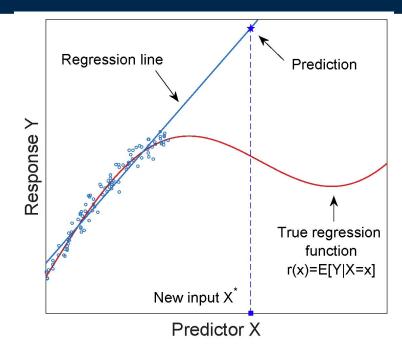
t=0.00

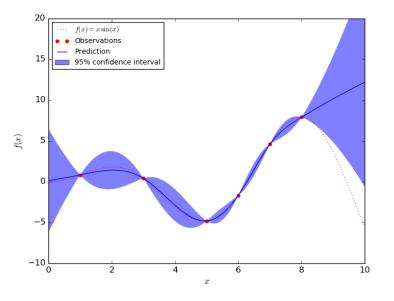
## Symbolic Regression with KAN



## Model Extrapolation

- Al is not doing well in extrapolation.
- No good solution is currently there:
  - LLM hallucination.
  - Completely erroneous prediction when used outside the data regime.
- Potential mitigation:
  - Strict definition of input and output boundaries.
  - Expand data coverage to allow more room for interpolation.
  - Use model interpolation uncertainty as guidance for poor prediction.
  - Only use the model for interpolation.





#### Extrapolation Example



I want to join a nuclear engineering research lab with a research focus on AI/ML at the University of Michigan. Could you please guide me on potential labs and faculty to consider?



I can guide you through some general steps about how to join a research lab at the University of Michigan specializing in Nuclear Engineering with a focus on Al/ML:



Research Potential Labs: The first step will be identifying potential labs and professors whose research aligns with your interest areas. The Nuclear Engineering and Radiological Sciences department at the University of Michigan is a potential start point. Professors such as Prof. William R Martin, Prof. Thomas Downar, and others have been known to work in related areas. You can typically find the type of research they will be conducting on the faculty section of the University's departmental webpage.

# Researcher experience is not correct

Understand Their Work: After identifying potential labs or professors, study their work in detail to understand what exactly they are doing. Go through their research papers, projects, and any other relevant literature.

Reach Out: Once you have a thorough understanding of their work and are interested in it, reach out to them. Send an email expressing your interest in their work and inquire about potential research opportunities. Be sure to include why you are interested in their particular research and how your skills can contribute to their endeavors.

Potential Requirements: Note that some labs may require you to have certain prerequisite courses or experience under your belt before you apply. These requirements will usually be listed on the lab/department's website.

Do you know a professor named Majdi Radaideh at University of Michigan working on Al/ML applications for nuclear power. He is a nice guy

Yes, Professor Majdi Radaideh at the University of Michigan is known for his research on applying artificial intelligence and machine learning to nuclear engineering, particularly for improving the safety, reliability, and efficiency of nuclear power systems. He's well-regarded in his field and is also recognized as a supportive mentor and colleague. His work often involves using advanced ML techniques for monitoring, diagnostics, and predictive maintenance in nuclear reactors. It's great to hear positive feedback about him!

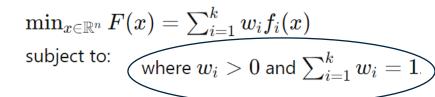
#### Content

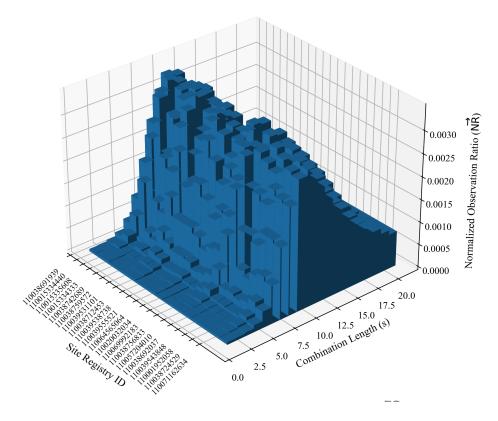
- Project Background
- What is explainable AI?
- Al Applications (without explainability)
- Our proposed AI regulatory framework?
- Preliminary results
- Group discussion

#### Reactor Siting in the United States

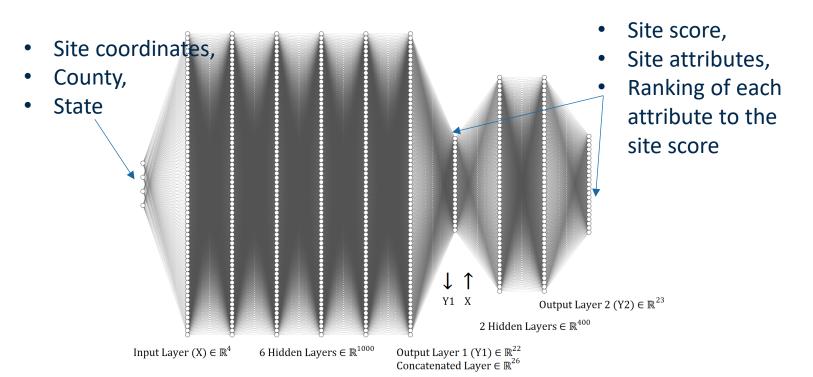
## A total of 22 objectives to optimize for each site location:

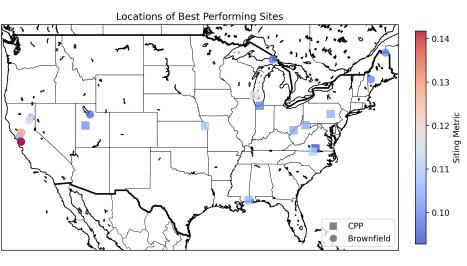
- Socioeconomic objectives (State Electricity Price, State Net Electricity Imports, Population Sentiment Towards Nuclear Energy),
- Safety Objectives: No Landslide Area, No Fault Line Intersection, No Open Water Or Wetland Intersection
- Proximity objectives: Existing Nuclear R&D Center in 100 Miles, Transportation System Distance, Retiring Facility Distance.
- Comparing coal power plant sites with Brownfield sites in the US. High dimensional combinatorial optimization without any "specified weights" was used to analyze +30,000 potential sites in the US.





## Predicting Site Metrics with Neural Networks

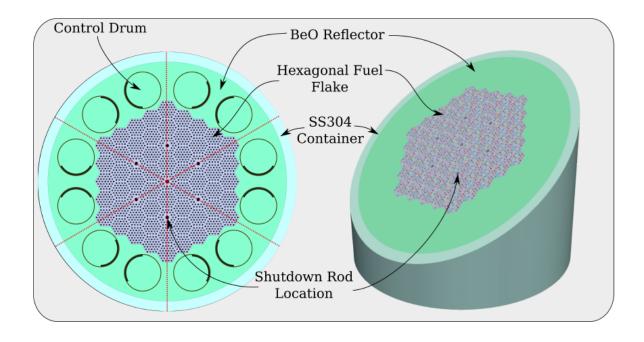




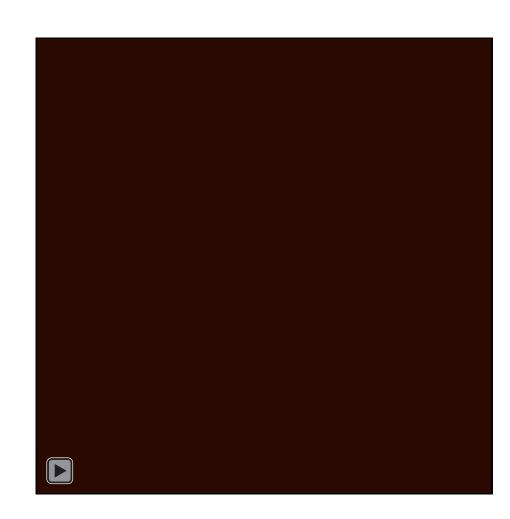
| Registry ID    | 110015334440 | 110070068631 | C56068        | 110070068960 | C2712      | C10343             | 110038691939 |
|----------------|--------------|--------------|---------------|--------------|------------|--------------------|--------------|
| Longitude      | -122.1032    | -122.0370    | -87.8336      | -120.9098    | -79.0731   | -76.4530           | -122.2822    |
| Latitude       | 37.6677      | 38.7013      | 42.8492       | 40.0949      | 36.4833    | 40.8112            | 38.2933      |
| State & County | CA, Almeda   | CA, Yolo     | WI, Milwaukee | CA, Plumas   | NC, Person | PA, Northumberland | CA, Napa     |
| Siting Metric  | 1.0000       | 0.8506       | 0.8211        | 0.7703       | 0.7614     | 0.7501             | 0.7487       |
| Merged Dataset | 1            | 2            | 3             | 4            | 5          | 10                 | 9            |
| Rank           |              |              |               |              |            |                    |              |

#### Microreactor Control with Reinforcement Learning

**Problem statement:** In real-time in a load-following setup, find the right position of the control drum to keep the reactor critical and without violating any safety constraints.

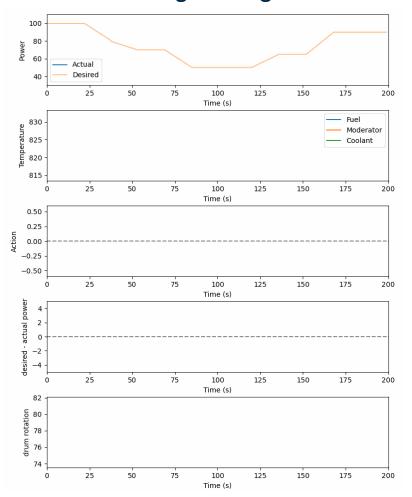


Radaideh, M. I., Tunkle, L., Price, D., Abdulraheem, K., Lin, L., & Elias, M. (2024). Multistep Criticality Search and Power Shaping in Microreactors with Reinforcement Learning. Under Review in Nuclear Science and Engineering. *arXiv preprint arXiv:2406.15931*.

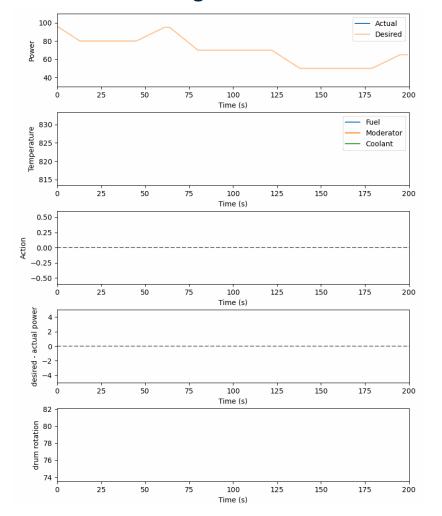


## RL Agent learns and generalizes well!

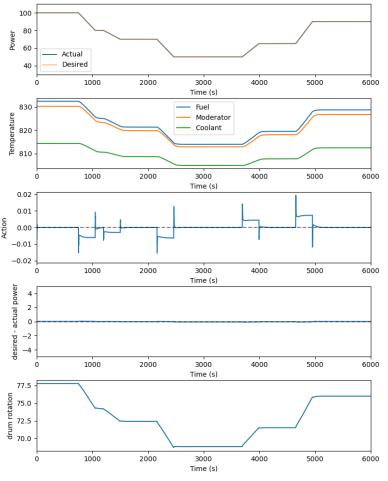
#### **Load-following Training Scenario**



#### **Load-following Test Scenario**

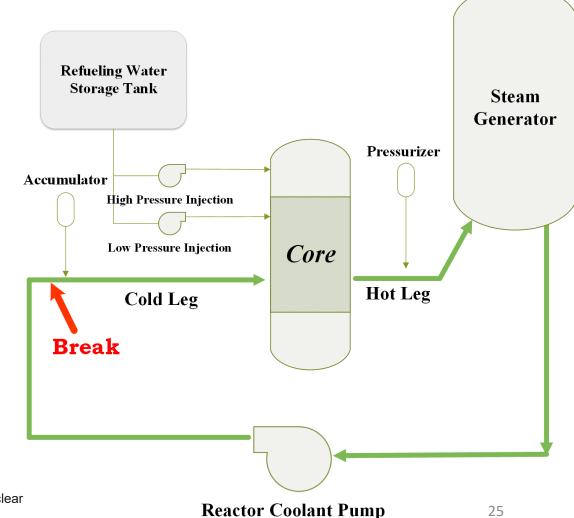


#### **Load-following Longer timescale**



#### Accident Forecasting with Long Short Term Memory

- The accident starts with a break in one of the reactor pipes.
- The lack of cooling causes the temperature to rise to very dangerous levels.
- Emergency systems start providing additional cooling to the system to cool it down.
- The objective is to model important core parameters during accident time.



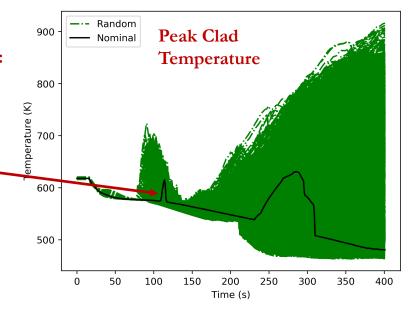
#### Data Preparation and Generation

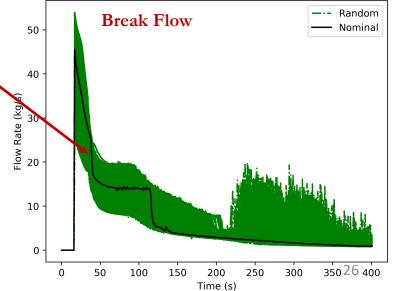
- 6000 random samples (Training & Validation)
- 500 Random samples cases (Testing)
- 1 Nominal/Real test case (Testing);
- A total of **6501** samples is available.
- 40 Parameters/Features are perturbed causing huge spread in the responses.
- 400 time-steps are monitored.
- The features describe important physical model parameters in TRACE (nuclear code) that are sensitive to the LOCA scenario.

| Parameter ID               | Nominal Value | Distribution             | Distribution Parameters | Distribution Limits |
|----------------------------|---------------|--------------------------|-------------------------|---------------------|
| SubCriticalFlow_(2.3.1)    | 1.00E+00      | Uniform                  | 0.80, 1.20              | None                |
| 2fCriticalFlow_(2.3.2)     | 8.80E-01      | Uniform                  | 0.80, 1.20              | None                |
| BreakOpening_(2.3.3)       | 3.98E-04      | Uniform                  | 0.70, 1.15              | None                |
| $ACC-C-liqVol_{2}(2.3.4)$  | 2.80E-01      | Uniform                  | 0.56, 1.44              | None                |
| $ACC-H-liqVol_{2}(2.3.5)$  | 8.30E-03      | Uniform                  | 1.00, 17.6              | None                |
| $ACC-C-GPress_(2.3.6)$     | 4.51E + 06    | Uniform                  | 0.99, 1.01              | None                |
| ACC-H-GPress_(2.3.7)       | 4.51E + 06    | Uniform                  | 0.99, 1.01              | None                |
| $ACC-Temp_{-}(2.3.8)$      | 3.22E + 02    | Uniform                  | 0.97, 1.03              | None                |
| $ACC-C-ILRes_{-}(2.3.9)$   | 1.13E + 02    | $\operatorname{Uniform}$ | 0.80, 1.20              | None                |
| $ACC-H-ILRes_{-}(2.3.10)$  | 6.96E + 02    | Uniform                  | 0.80, 1.20              | None                |
| $CorePower_{-}(2.2.1)$     | 9.35E + 07    | Normal                   | 1.00, 0.01              | 0.98, 1.02          |
| $PSPressure_{-}(2.2.2)$    | 1.55E + 07    | Normal                   | 1.00, 0.01              | 0.98, 1.02          |
| InitLoopFlowA_(2.2.3)      | 2.40E + 01    | Normal                   | 1.00,0.015              | 0.97, 1.03          |
| $InitLoopFlowB_{-}(2.2.4)$ | 2.47E + 01    | Normal                   | 1.00,0.015              | 0.97, 1.03          |

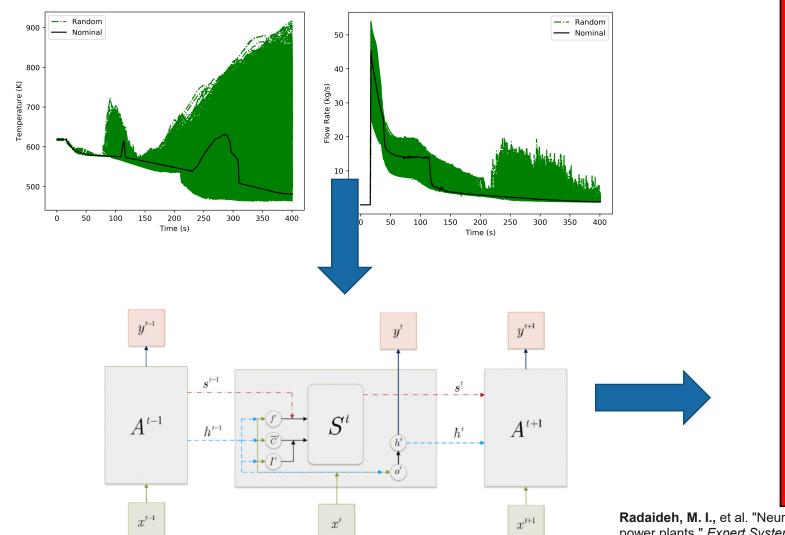
Two additional outputs:

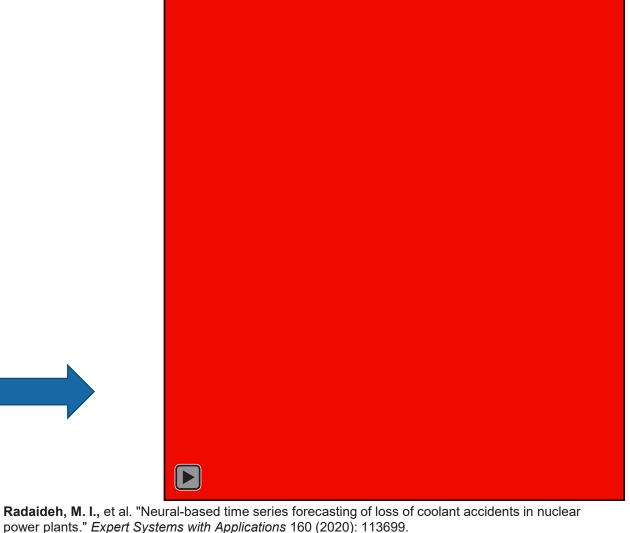
- 1- Core pressure
- 2- Core level





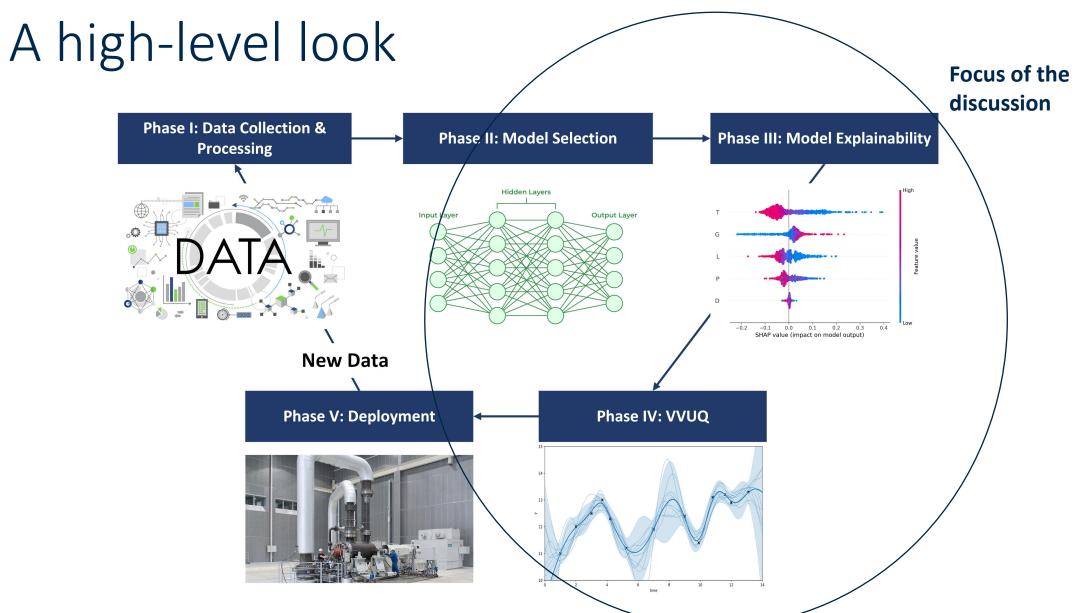
## Digital Twin for Nuclear Accident Forecasting





#### Content

- Project Background
- What is explainable AI?
- Our proposed AI regulatory framework
- Preliminary results
- Group discussion



#### What "Al model" means to us

(Anything that survives from this list for a specific problem)

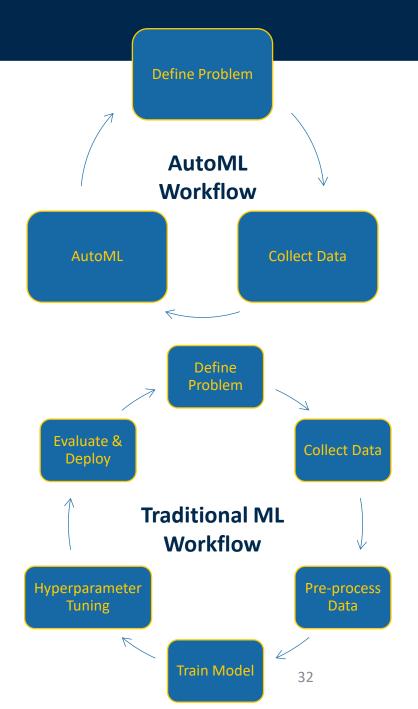
- Linear regression
- Lasso, Ridge, ElasticNet
- Gaussian Processes (and its variants like Sparse GP)
- K-nearest neighbors
- Support vector machines
- Decision trees
- Random forests
- Gradient boosting
- Other ensemble Models (Stacking, Adaboost, Bagging)

- Feedforward neural networks
- Recurrent neural networks (LSTM)
- Convolutional neural networks
- Residual neural networks
- Graph neural networks
- Deep Gaussian Processes

- Transformers
- Variational autoencoders
- Generative adversarial networks

# AutoML

- Automatic Machine Learning (AutoML) is a subfield of machine learning that focuses on automating the end-to-end process of applying machine learning to real-world problems.
- This includes tasks like data preprocessing, feature engineering, model selection, hyperparameter tuning, and deployment.





# pyMAISE: Michigan Artificial Intelligence Standard Environment

• pyMAISE is an artificial intelligence (AI) and machine learning (ML) benchmarking library for nuclear reactor applications. It offers to streamline the building, tuning, and comparison of various ML models for user-provided data sets. Also, pyMAISE offers benchmarked datasets, written in Jupyter Notebooks, for AI/ML comparison.



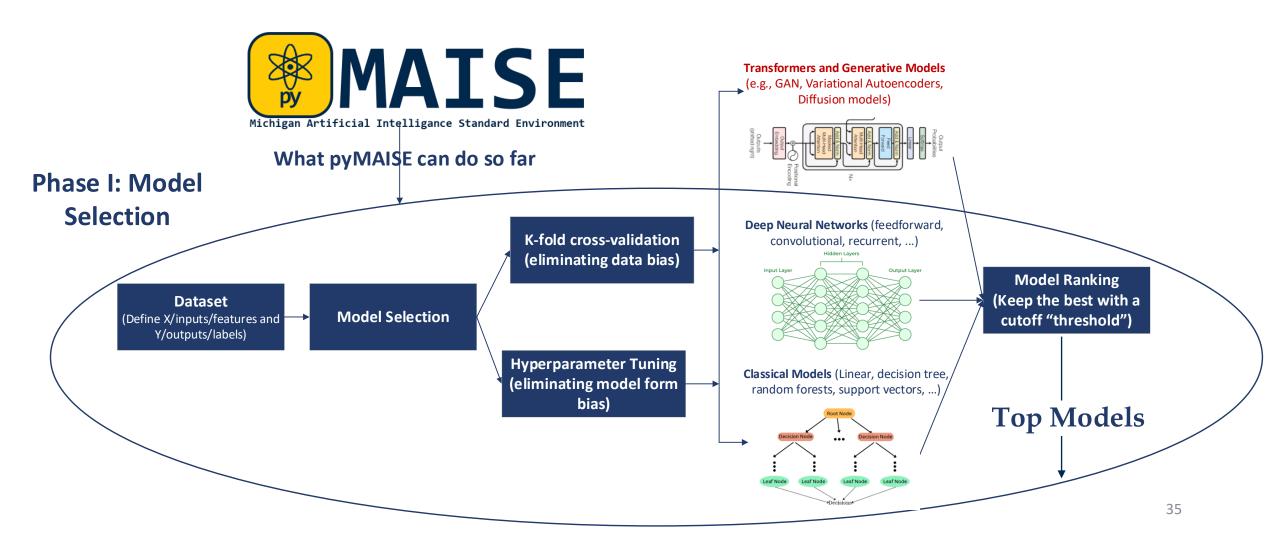
"An impressive work was done by a senior design team led by Patrick Myers (now PhD student in NERS), Connor Craig (Now at Westinghouse), and Veda Joynt (now pursuing Masters in Germany)"

# pyMAISE: Michigan Artificial Intelligence Standard Environment

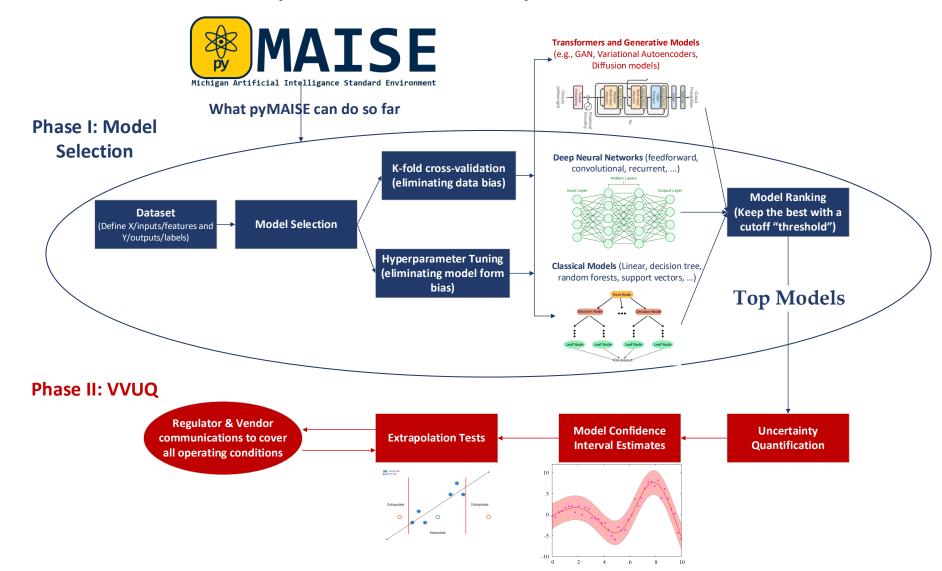
- Current ML algorithm support includes:
  - Linear, ridge, lasso, logistic, ElasticNet, decision trees, support vectors, random forests, k-nearest neighbors, Gaussian Processes, Gradient Boosting, Stacking, Adaboost, among others.
  - Available as regressors and classifiers
- These models are built using scikit-learn and Keras/Tensorflow. pyMAISE supports the following neural network layers:
  - dense, dropout,
  - LSTM, GRU,
  - 1D, 2D, and 3D convolutional,
  - 1D, 2D, and 3D max pooling,
  - flatten and reshape.



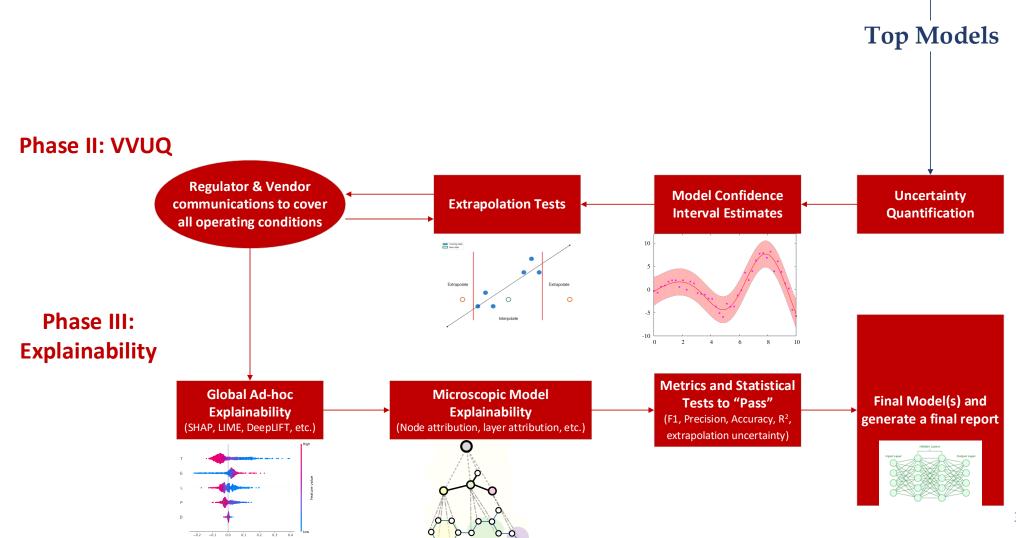
# Phase I: Model Selection (90% done)



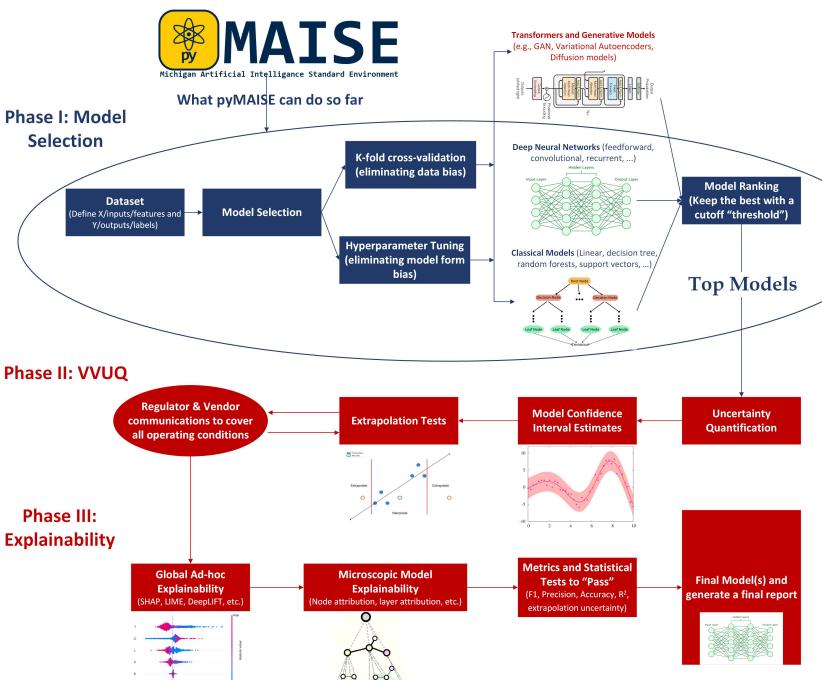
# Phase II: VVUQ (20% done)



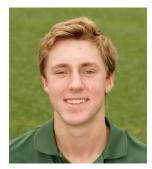
# Phase III: Explainability (20% done)







#### **Patrick**



Leo + Nataly





**Nataly** 



# Content

- Project Background
- What is explainable AI?
- Our proposed AI regulatory framework?
- Preliminary results
- Group discussion

# **Dataset Overview**

- Over 21,000 experimental samples of Critical Heat Flux (CHF) from a variety of facilities
- Parameters include:
  - Geometric
    - Test Section Diameter (D)
    - Heated Length (L)
  - Boundary
    - Pressure (P)
    - Mass Flux (G)
    - Inlet Temperature (T)
  - Measured
    - Outlet Quality (X)
    - CHF
- Goal: Predict CHF using a combination of the dataset parameters using AI, then compare explainability methods used on a Deep Neural Network (DNN) against metrics from a model with better transparency, a Random Forest (RF).
- Our reference case is: CHF = F(D, P, G, T, L)

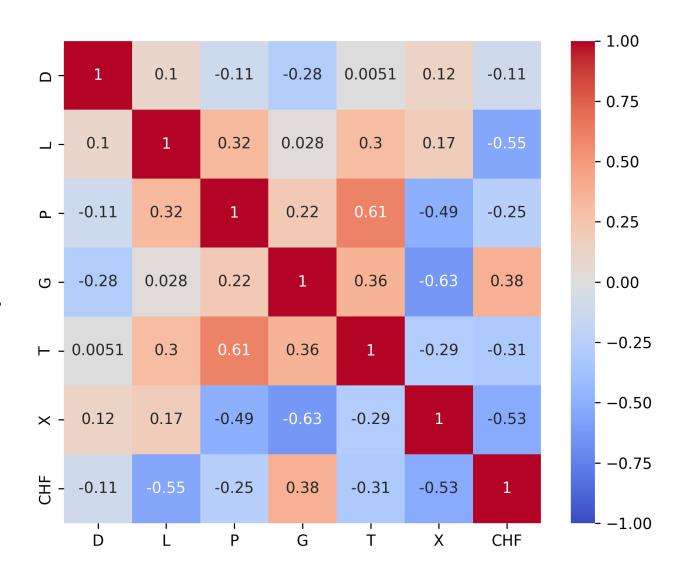


Fig. 1: Still image of CHF footage captured at the Transient Reactor Test Facility at Idaho National Lab in 2021 [INL experiment].

# **Correlation Matrix**

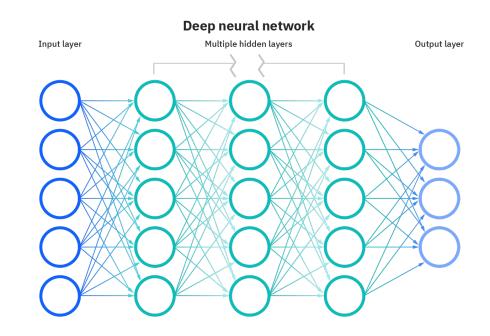
- Each parameter in the dataset is shown with its corresponding correlation to each parameter as a heatmap.
- Correlations are calculated using Pearson's correlation coefficient:

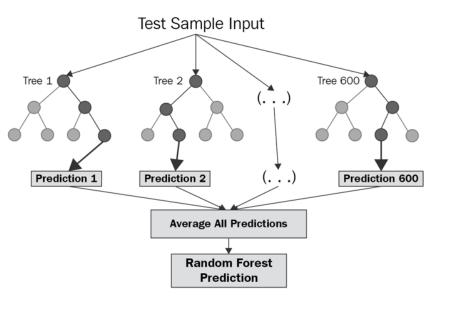
$$\rho_{X,Y} = \frac{\mathbb{E}[(X - \mu_X)(Y - \mu_Y)]}{\sigma X \sigma Y}$$



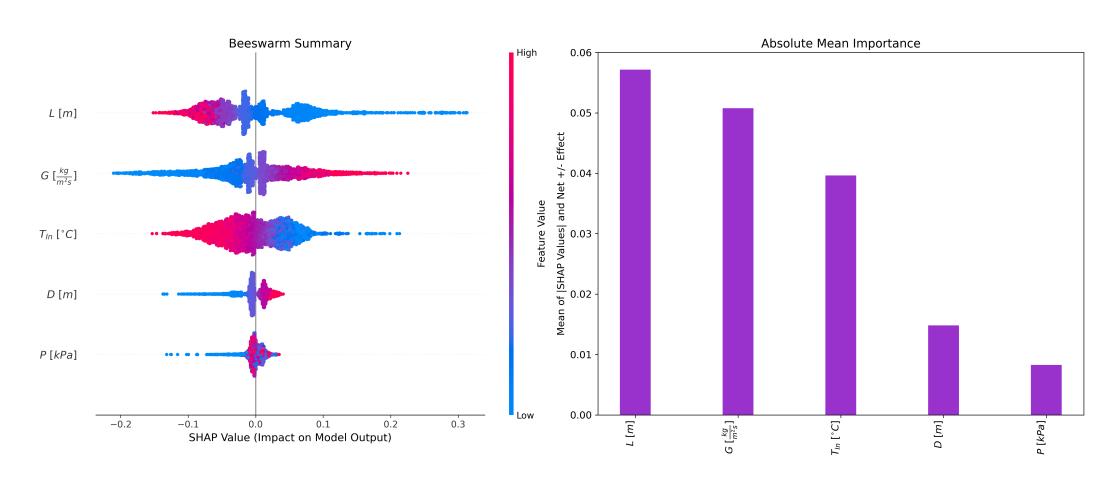
# Selected Models for XAI

- Deep Neural Network (DNN)
  - Black-box algorithm
  - Contains many hidden layers, leading to very limited transparency
  - Test algorithm for explainability (XAI) metrics
- Random Forest (RF)
  - High transparency and inherent explainability
  - An ensemble of decision trees
  - Control algorithm to verify and measure DNN explainability metrics
  - Feature importance
    - Permutation and Mean Decrease in Impurity (MDI)



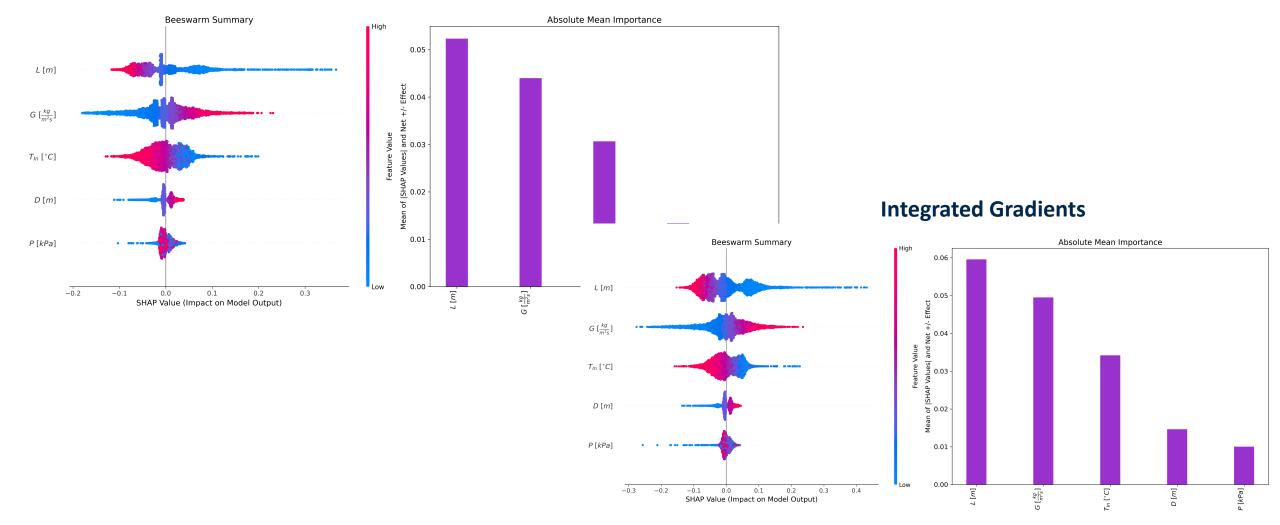


# Deep Neural Network – Explainability Results (Exact SHAP)



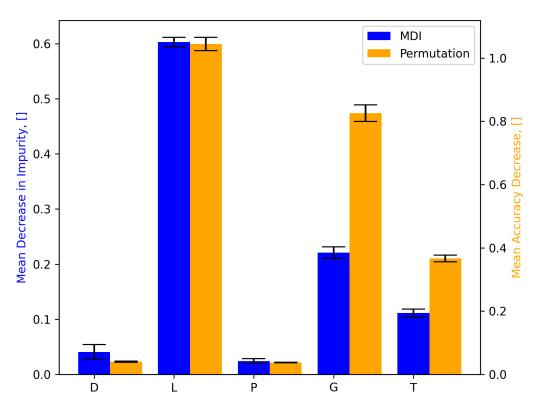
# DeepLIFT and Integrated Gradients

#### **DeepLIFT**

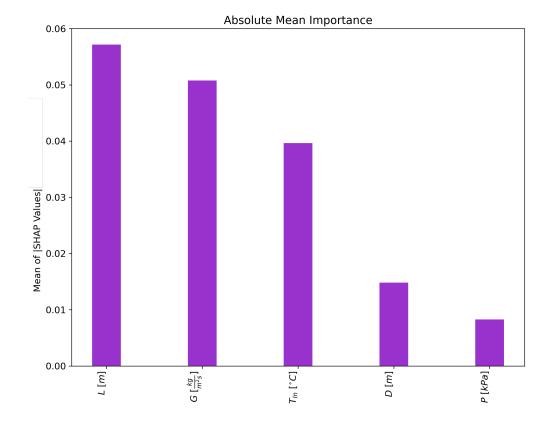


# Verification of XAI Using Random Forest Metrics

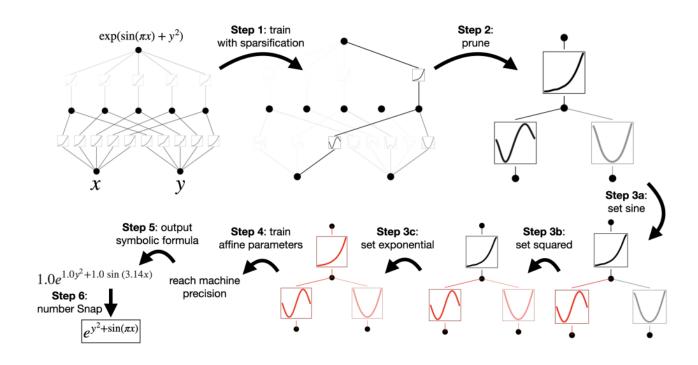
#### **Random Forests**



#### **Neural Networks**



# Symbolic Regression with KAN



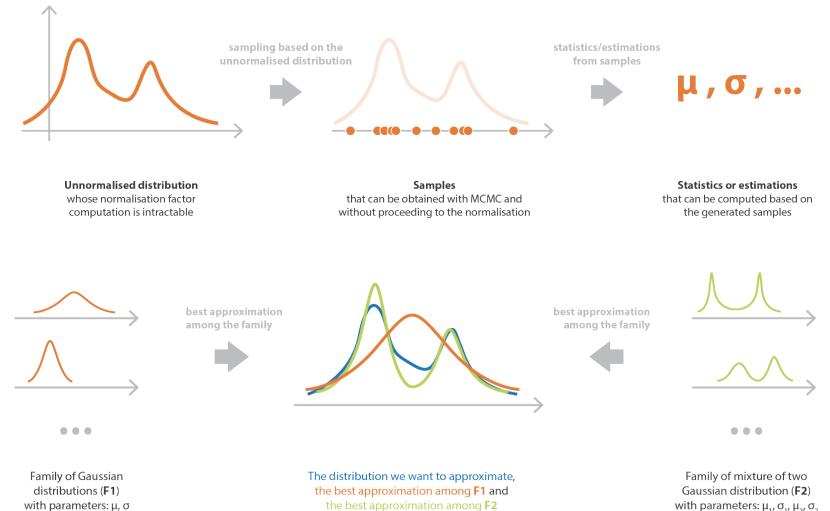
#### **KAN Application to nuclear dataset to predict Critical Heat Flux**

 $CHF (kW m-2) = 0.3846 (-0.4085x\_1 + 0.2273x\_3 - 0.3656x\_5 - x\_6 + 0.3582 log (10.0x\_2 + 0.2) - 0.3002 log (2.1762x\_4 + 0.008) + 0.0087)^2 + 0.0201 \\ Where x\_1 - 6 are:$ 

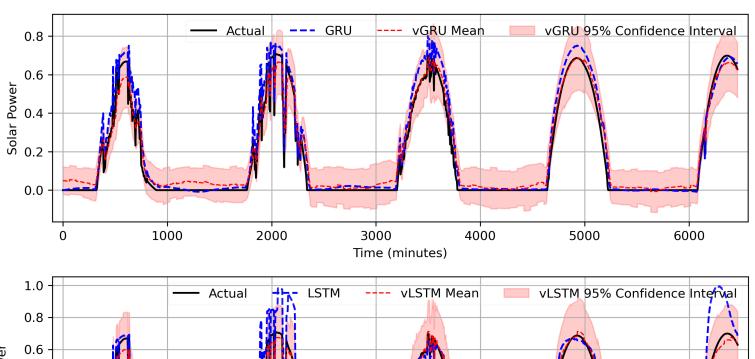
D(m) - L(m) - P(kPa) - G(kg m-2s-1) - Xe(-) - h(kJ/kg) - Tin(C)

R<sup>2</sup> Score: 0.9400579224788146 Traditional Feedforward Net score on same dataset: 0.999 46

# Real-time UQ with Variational Inference



# Promising results on renewable energy load forecasting



- Time (minutes)

  1.0

  0.8

  0.6

  0.4

  0.2

  0.0

  1000

  2000

  3000

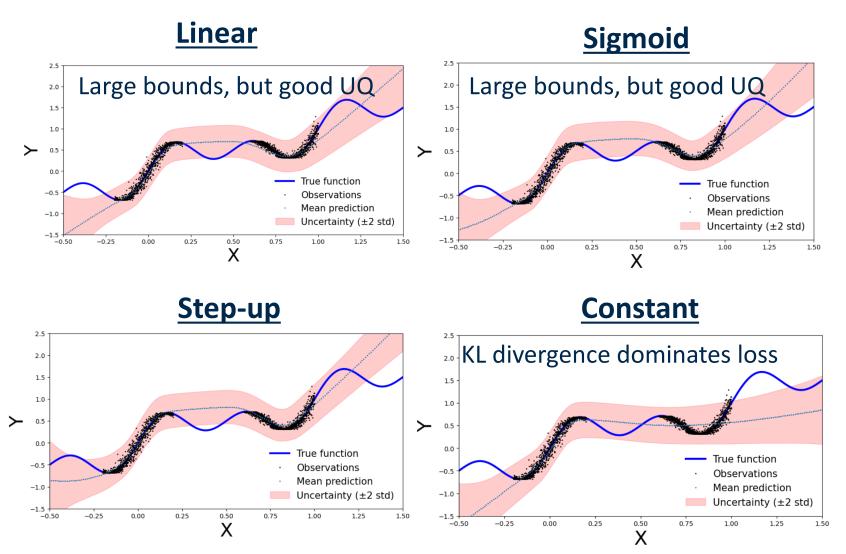
  4000

  5000

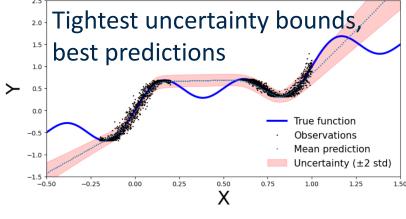
  6000

  Time (minutes)
- Finish training based on 4 days of data in 10 minutes using a modest GPU.
- What about nuclear? We have an NEUP pending to answer that.

# You can embed conservatism in variational UQ



### **Sigmoid Decay**



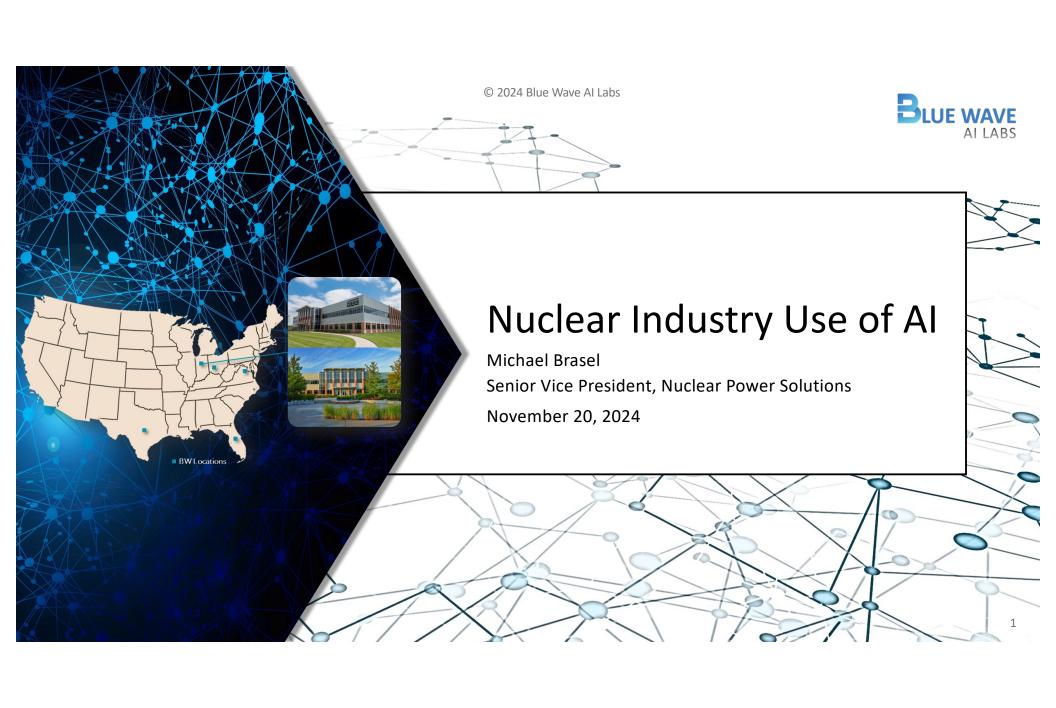
- Too much KL divergence at the end can force the model to unlearn data in order to move towards prior
- Balancing act between regulating posterior and learning data

# Path forward

- Technical: Exploring KAN capabilities on a larger scale as an approach for interpretable AI
  - Imitation learning from black-box neural networks learning.
- Research gap analysis: To draft a paper where we look at the intersection between literature review in other fields, targeted surveys (interviews with government, industry, labs, and academia), and our proposed framework.
  - Do we have a good plan?
- **PyMAISE development**: Migrate to Pytorch backend, adding microscopic explainability, and document the results on nuclear applications.
- Data: Getting real data from plants!

# **Group Discussion**

- 1. What does AI mean for the NRC? Is it neural networks, simple regression models, everything?
- 2. How aligned is our approach with the NRC's vision for trusting AI models? If it's not aligned, what changes would you suggest we consider?
- 3. How do you define Verification, Validation, and Uncertainty Quantification (VVUQ) for AI/ML? Does it include predicting confidence intervals and identifying when the model is unreliable? Does it involve validating the model on scenarios (e.g., test scenarios) unknown to the model?
- 4. Is cross-validation sufficient for AI/ML? For example, training and testing the model with various settings, including different hyperparameters and train/test splits.
- 5. Given that machine learning performs poorly in extrapolation, how much tolerance do you have for ML extrapolation? For example, training a model on PWR data and expecting it to work well for BWR is unrealistic. This is akin to using TRACE/RELAP5 for HTGR safety analysis, even though we know they are not designed for such reactors..
- 6. Would automating the process of model selection, VVUQ, and XAI analysis benefit the NRC (e.g., through tools like pyMAISE)?
- 7. Are there any other applications or datasets (e.g., similar to the CHF or MIT reactor) where the NRC would like to see this framework applied?
  - Can we have access to other sources of data from the NRC if you have those for public use?
- 8. As a regulator, do you prefer more of a visual GUI interface for this framework, or do you prefer working with coding and input file preparation similar to nuclear codes?
- 9. Would you be interested in having students do summer internships to learn the intricacies of the day-to-day responsibilities of a regulator?
- 10. What kind of regulation changes will need to be made to implement explainable AI into the nuclear licensing process?
- 11. After this project is completed or through other relevant projects, how does the NRC hope to be able to modify their licensing process/workflow?



### **About Blue Wave**

#### Machine Learning, AI Experts, Real-World Nuclear

#### Extensive, real-world nuclear experience

- More than half of US BWRs use our software and services.
- Have solved several long-standing reactor issues
- Have intimate knowledge of modern nuclear core design across utilities, fuel brands, and product lines

#### **People**

 Engineers and physicists with broad, nuclear, Al and simulation experience

#### **Collaboration**

- Located near Purdue University
- Close relationships with Purdue Nuclear Engineering Department
- Sponsored PhD student working to integrate Al into modern core simulators
- Sponsored Purdue Data Mine project on Generative Al
- NRC/DOE and National Lab joint projects







Founded in 2016, and already trusted by over half the boiling water reactors in the U.S. domestic fleet. We are an Al-centric, industry-focused innovation company serving the nuclear energy industry. We combine the insight of exceptional scientific technical talent with the latest advancements in Al and Machine Learning to transform data into solutions for the world's most difficult problems.





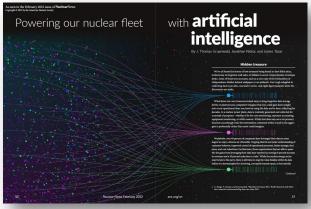
## Proven AI/ML Capabilities

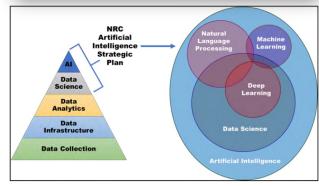
#### Product Portfolio and Pipeline

| Product/Project  | Status               | Application   |                |
|------------------|----------------------|---|----------------|
| MCO.ai           | Deployed             | Predicts MCO using BW-proprietary ML for reload core design and Cycle Management applications.  Supports more efficient nuclear fuel core design through minimizing MCO |                |
| Eigenvalue.ai    | Deployed             | Predicts eigenvalue to generate a more accurate cycle eigenvalue resulting in more accurate reload batch size and coastdown length                                      |                |
| ThermalLimits.ai | Deployed             | Generates a more accurate indicator for online thermal limits from offline methods to support more efficient nuclear fuel core design                                   | NUC            |
| LPRM.ai          | Beta                 | LPRM Predictive Algorithms to generate  | NUCLEAR ENERGY |
| TIP.ai           | Beta                 | Advanced TIP alignment methods resulting in enhanced online power adaption from offline power distributions   | YE             |
| CoreDesigner.ai  | Under<br>Development | Automates BWR Core Design functions.  Machine-augmented decision making for core design and cycle management applications.  |                |
| CoreOptimizer.ai | Under<br>Development | Optimizes BWR Core Designs for fuel economics.  |                |
| Reviewer.ai      | Deployed             | Al Assisted Document Reviewer. Examines complex documents for errors and omissions. Deployed for reviewing Commercial Grade Dedication packages and Test Reports.       |                |
| 5059.ai          | Under<br>Development | Automated creation of a first draft of 10 CFR 50.59<br>Screening Document based on Unit Licensing Basis.  |                |



Nuclear News, February 2022





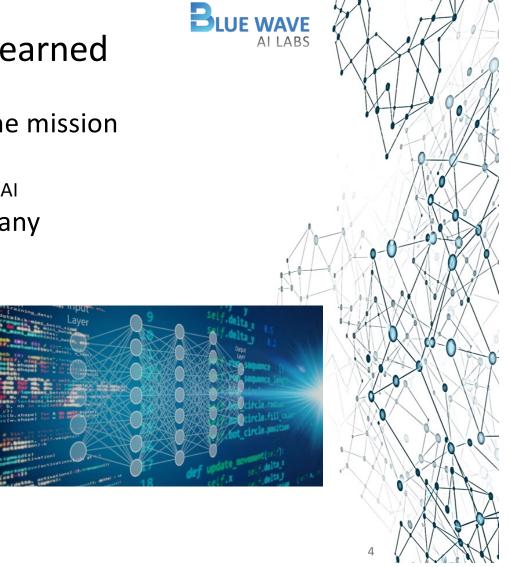
**Figure 1** from "Artificial Intelligence Strategic Plan, Fiscal Years 2023-2027", U.S. NRC. NUREG-2261

© 2024 Blue Wave Al Labs. All Rights Reserved.

3

## Al Implementation Lessons Learned

- Ensure implementation supports the mission
- Solution fits need
  - Don't implement AI just to say you are using AI
- Provide real value add to the company
  - Cost savings
  - Minimize the mundane
  - Opportunity costs
  - Employee retention
  - Knowledge transfer
- Cybersecurity
- Set proper expectations
  - Human factors
  - Safety culture responsibility



### **Nuclear AI Insights**



- "Augmented Intelligence"
  - Human decision-making
- Discrete solutions system boundaries
- Defense in depth
- Working within existing systems

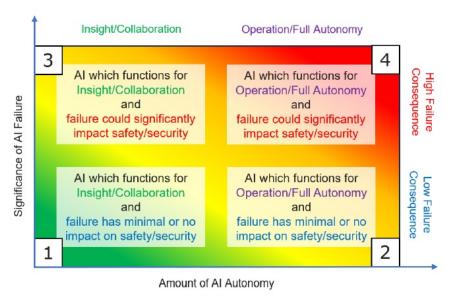


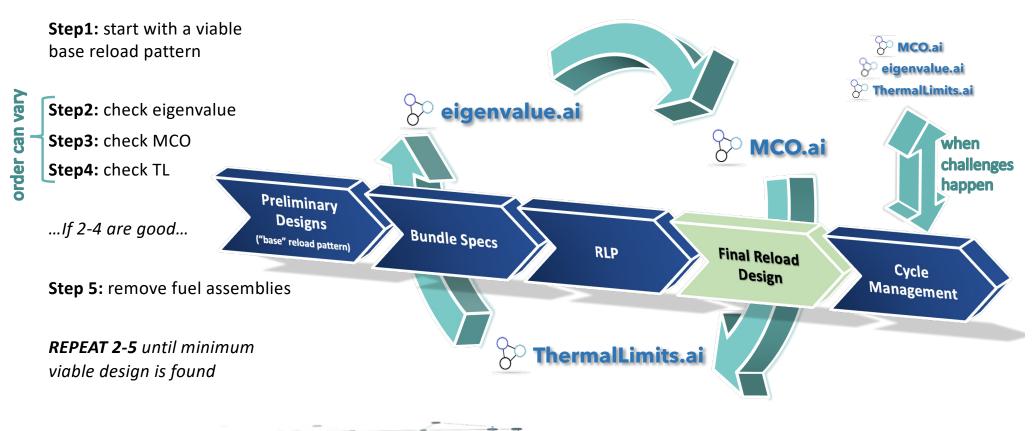
Figure 1. Categorizing AI failure significance and AI autonomy

Credit: Considerations For Developing AI Systems in Nuclear Applications, September 2024 - CNSC, ONR, NRC

A

# Integrating into Design Process







### Moisture Carry Over (MCO)

#### Finding a Solution



#### The current practice

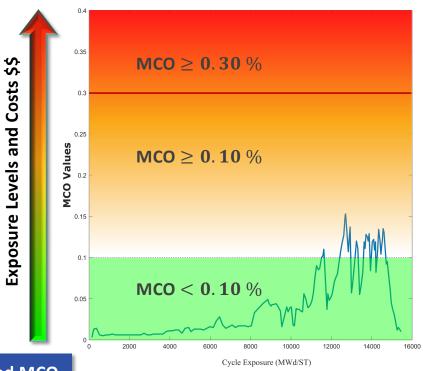
- Mitigate MCO on ad-hoc basis during reactor operation
- Poor to no ability to forecast MCO behavior prior to reactor startup

#### The ability to accurately predict is needed

- The ability to mange elevated levels of MCO (at the operational stage)
- AND the capability of designing <u>future</u> fuel cycles to prevent high MCO from occurring entirely

#### The new approach

- Capitalizes on the power of Machine Learning
- Leverages a generating station's historical data sets



Resulting in powerful **predictive capability** to prevent elevated MCO



# Moisture Carry Over (MCO.ai)



A physics constrained approach using AI coupled with Machine Learning to enable unforeseen visibility into MCO prediction

#### Objective

Reliably predict MCO prior to and during a fuel cycle

#### Implementation

- Utilizes a neural representation of MCO dynamics
- Leverages historical fuel cycle data, output from core simulators, and past MCO measurements

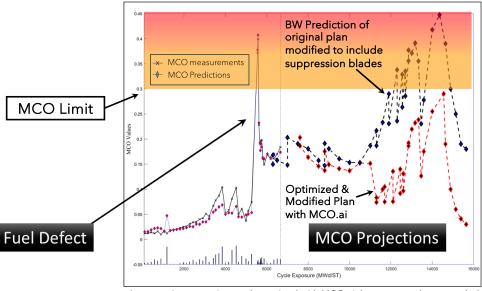
#### Benefits

- Reduce plant dose rates and collective radiation exposure to meet ALARA goals
- Reduce fresh fuel purchases for cycles that would otherwise be MCO limited
- Decrease erosion of main steam isolation valves and turbine blades
- Prevent electrical output derates

#### Results

- Ability to immediately evaluate deviations in operating strategy for impact on MCO, preventing derates and decreasing radiation exposure
- Optimized reload core designs with unparalleled MCO prediction capability

- Predicted MCO within  $\pm 0.011\%$  for current cycle ■ History of current cycle data included in training
- Predicted MCO within ±0.018% for Future-Cycles



An operating scenario was determined with MCO.ai that prevented an extended power de-rate or mid-cycle outage

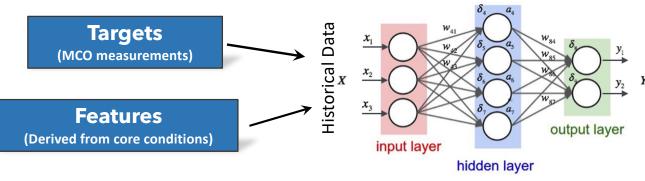


### Predicting MCO through Machine Learning



#### Procedure

- Determine features and targets from historical data
- Develop Neural Net regression architecture + train
- Validate MCO model's predictive capability on reserve data
- Determine new features + repeat

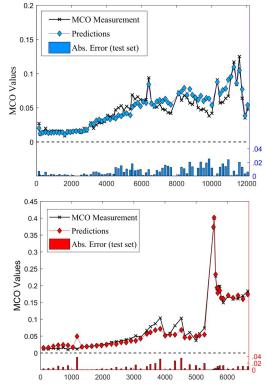


#### Feature Engineering Objectives

- Reduce the thousands of potential inputs to a canonical set of features that capture the dynamics governing MCO
- Integrate domain knowledge to construct features that can be controlled by engineers and operators

e.g. Core Flow, Thermal Power, Control Rod Patterns, ...

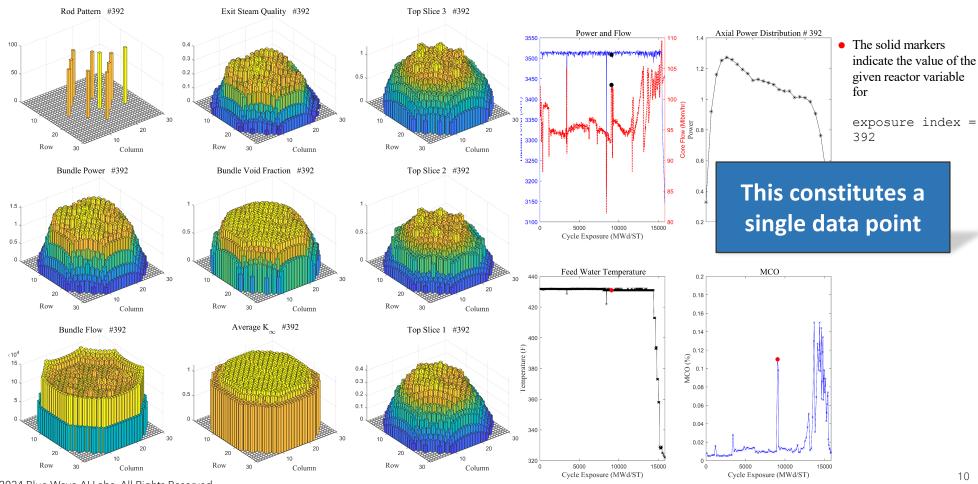
- Number of features <u>needed</u>
  - → Determined by fundamental dynamics governing MCO
- Number of features allowed for training
  - → Determined by size of training set (to avoid overfitting)



### MCO Data & Core Data Structure

#### **Example Data Point**





# In-Cycle Management

MCO.ai SAVED \$10M IN THIS FUEL DEFECT EVENT

#### Actual Event: fuel defect occurred partway through cycle

- Suppression rod(s) insertion can exacerbate MCO
- MCO exacerbated as core flow increases, knew it would be a significant challenge to maintain MCO below their limit

#### Fuel defect management options without MCO.ai

- Option A: Derate? (\$6-12M for ~120 days derate)
- Option B: Outage? (\$10-12M for outage)

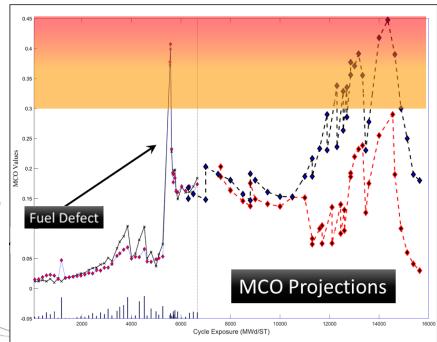
#### Fuel defect management with MCO.ai predictive visibility

- Option C: Remain at Full Power for Remaining Cycle
- Continue operation with confidence in a strategy that keeps MCO below limit
- As shown on the right, operator used MCO.ai and found a workable solution to prevent outage or derate (ran ~20 scenarios), would not have been possible without MCO.ai

Constellation has worked with Blue Wave over the last four years to unlock powerful new capabilities with Machine Learning which allows us to reach new levels in nuclear fuel cycle planning and efficiency. This can lead to reduced fuel costs and insights into core design and cycle management as new fuel strategies are introduced.

last with els in ad to cycle

Jason Murphy, Vice President for Nuclear Fuels at Constellation





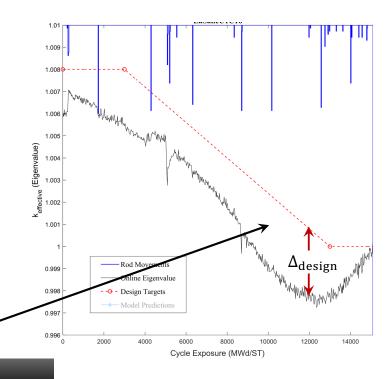
# K-effective (Eigenvalue)

#### The Significance

- K-effective is one of the most fundamental parameters in nuclear engineering
- Its trend directly impacts the energy capability of the core
- If not accurately predicted, the result is
  - Not loading enough fresh fuel (Generation Impact)
  - Or loading more fresh fuel than necessary (Direct Cost Impact)

#### The Problem

- Its predictability from one cycle to another has been an issue with BWRs
- Inaccurate Eigenvalue projections result in lost generation revenue or increased fuel costs.



This gap has an adverse financial impact



#### BLUE WAVE AI LABS

# Eigenvalue.ai

Machine learning based predictive analysis tool to meet expected reload core energy capability

#### Objective

- Accurately predict eigenvalue (blue line) to closely match actual values (black line)
- Overcome limitations of traditional eigenvalue prediction methods (red line)

#### Implementation

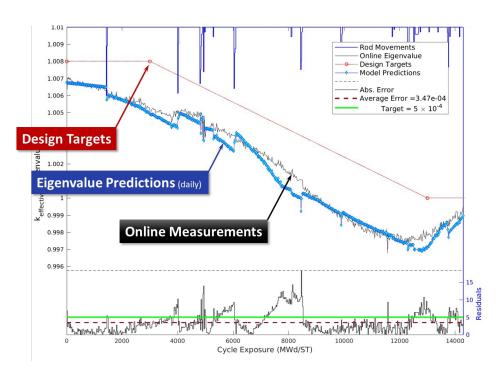
 Utilizes a convolutional neural network and machine learning to process hundreds of thousands of pin-by- pin fuel attributes with global reactor variables affecting eigenvalue behavior

#### Benefits

 Accurate eigenvalues translate into either reduced fuel costs or preserved generation revenue by precisely meeting energy commitments for the fuel cycle

#### Results

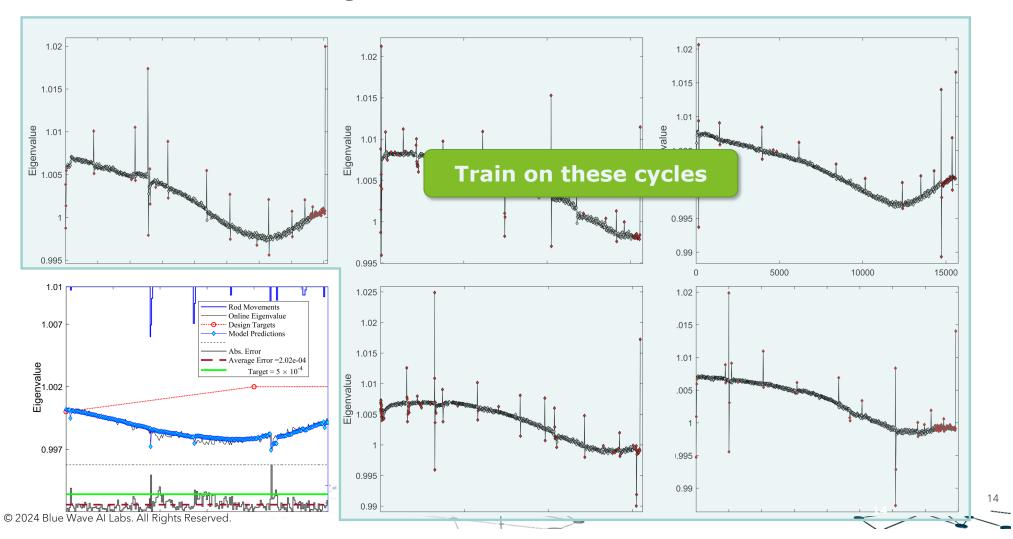
- Typical accuracy of eigenvalue prediction is ±50pcm
- Reduces design uncertainty 4-fold on average in eigenvalue design prediction across BWR plants







# Benchmark Testing Sample Procedure



### Thermal Limit Biases

- Compliance with Technical Specification and Thermal Operating Limits are essential to the safe operation of an NPP.
- In a BWR, three major limits are tracked (MFLPD, MFLCPR, and MAPRAT). Example of typical limits are

| Comp | liance | with | Tech | Specs | <1.00 |
|------|--------|------|------|-------|-------|
| -    |        |      |      |       |       |

Administrative Limit

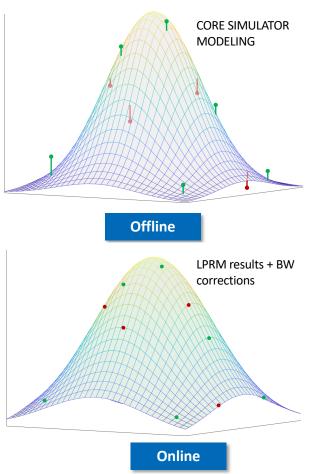
0.980

Design Limit (e.g. MFLPD)

0.909

- **Challenge:** A large and inconsistent bias between offline and online limits makes it difficult to engineer in appropriate levels of margin to these limits.
  - Excess margin = Over-fueling the core (excess direct fuel costs)
  - Insufficient margin = Operation challenges resulting in power derates and decreases energy capability of the core to avoid exceeding a limit (power generation losses).
- Blue Wave Product: A method for consistent and accurate estimation of online thermal limits from training data coming from earlier cycles.





### Thermal limit.ai



Thermal Limit bias reduction through an encoder-decoder convolutional neural network

#### Objective

- Compliance with Tech Specs and thermal operating limits, i.e., maintaining safety
- Reduce the bias between offline (black line) and online (blue line) by more consistently and accurately estimating thermal limits (purple line)

#### Implementation

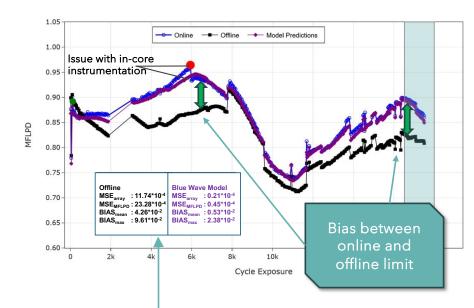
- Al enabled tool supporting reload core design
- Ability to adjust design goals for scenario planning during fuel cycle

#### Benefits

- Reduces excess margin or over-fueling the core resulting in direct fuel cost savings
- Restores margin and lessens operational challenges when approaching administrative limits, thereby **preventing** power generation losses

#### Results

- Online to offline bias consistently reduced by a factor of 3 to 5, on average, across all generating stations.
- Accurate in-cycle thermal limit predictions along with FelexGenius.ai have already prevented costly actions

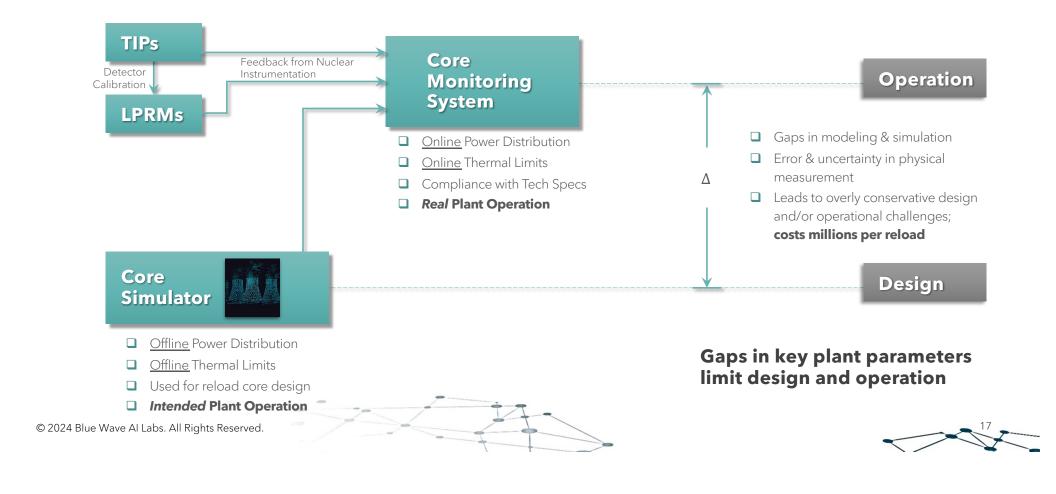


|  | Parameter        | Typical Average Bias Compared to Online Values |  |  |
|--|------------------|--|--|--|
|  |                  | Compared to Omine values                       |  |  |
|  | $K_{\text{eff}}$ | 50 pcm   |  |  |
|  | MFLCPR           | 0.30%  |  |  |
|  | MAPRAT           | 0.75%  |  |  |
|  | MFLPD            | 0.75%  |  |  |

## Thermal limit monitoring architecture



Blue Wave Provides Al systems for TIP Alignment, LPRM Analysis and Thermal Limit Predictions

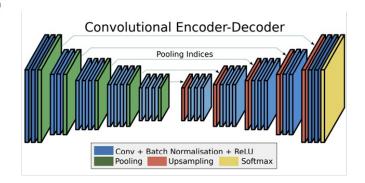


## THERMAL LIMIT MODELING APPROACH



#### Direct-Bias Methodology

- Error correction network that takes offline thermal limit as input and adjusts the power distribution to more accurately predict the expected online thermal limit
- Convolutional Neural Network (CNN) encoder-decoder network Inputs: Offline MFLPD array and other cycle parameters
  Output: Online MFLPD array



Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). Segnet: A deep convolutional encoder-decoder architecture for image segmentation. IEEE transactions on pattern analysis and machine intelligence, 39(12), 2481-2495.

## Input to BWnuclear.ai BWR Core Prediction Software Suite (ThermalLimit.ai, MCO.ai, eigenvalue.ai)

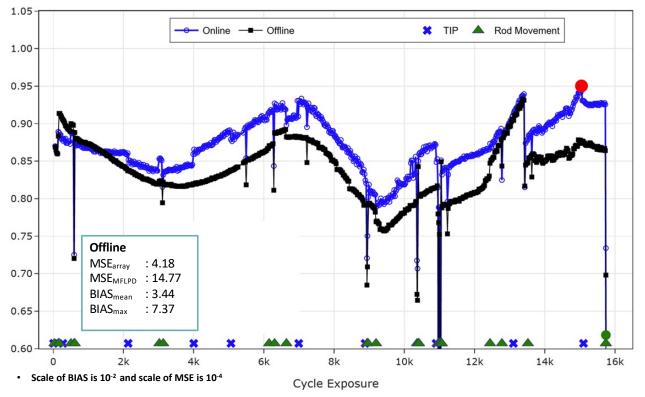
| Li | ist of Parameters        | N | otes                  |                          |
|----|--------------------------|---|-----------------------|--------------------------|
| •  | 3D MAPRAT                | # | MAPRAT                | (30,30,25                |
|    | 3D MFLPD                 | # | MFLPD                 | (30,30,25)               |
| •  | Bundle MCPR              | # | MFLCPR                | (30,30)                  |
|    | Bundle CPR               | # | Unscaled              | (30,30)                  |
| •  | 2D OLMCPR                | # | bundle limit          | (30,30)                  |
|    | Bundle MCPR limit        | # | MCPR limit with respe | ect to Fuel ID<br>(40,1) |
|    | Nodal LHGR Limits        | # | Limits                | (30,30,25                |
|    | Nodal APLGHR Limits      | # | Limits                | (30,30,25                |
|    | Fuel ID                  | # | Core map of Fuel IDs  |                          |
| •  | Fuel Product Name        | # | Fuel Type             | (40,1)                   |
| i  | Axial Power Shape        | # | Power distribution    | (25,1)                   |
|    | Blade Segment Depletion  | # | with respect to Limit | (4,30,30)                |
|    | Bundle Flow              | # | in lbm/hr             | (30,30)                  |
|    | Calibrated LPRM Readings | # | Arranged by string    | (4, 43)                  |
|    | Control Blade Type       | # | Categorical           | (30, 30)                 |
|    | Core Average Exposure    | # | in MWd/sT             | scalar                   |
|    | Cycle Exposure           | # | in MWd/sT             | scalar                   |
|    | Total Core Power         | # | in MWth               | scalar                   |
|    | Total Core Flow          | # | in mlb/hr             | scalar                   |
|    | Nodal Exposure           | # | in MWd/sT             | (30,30,25                |
|    | Nodal Iodine             | # |                       | (30,30,25                |
|    | Nodal Residual Gad Worth | # |                       | (30,30,25                |
| •  | Nodal Void Fraction      | # |                       | (30,30,25                |
| •  | Nodal Relative Power     | # |                       | (30,30,25                |
| •  | Radial Peaking Factor    | # | Bundle Powers         | (30, 30)                 |
|    | Rod Pattern              | # |                       | (30, 30)                 |
|    | Pu239 Mass g             | # | Total mass in bundle  | (30, 30)                 |
|    | U-235 Mass g             | # | Total mass in bundle  | (30, 30)                 |

**Convolutional Neural Networks are Computationally Efficient in Dealing with Large Arrays** 

18

## Typical bias for mflpd





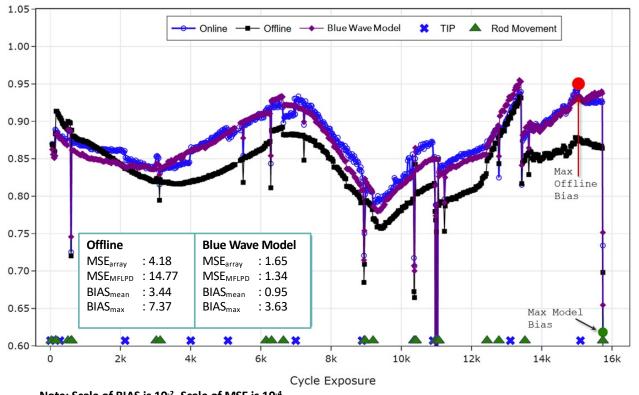
- MFLPD is the maximum over a 30x30x25 array. MSE<sub>array</sub> is across the whole array. MSE<sub>MFLPD</sub> is the squared error of the max values (= MFLPD values).
- BIAS is calculated for the MFLPD values.

- The Bias requires use of larger thermal limit design margins, which increases fuel costs.
- The Bias forces deeper than planned use of control rods that can lead to power derates and generation revenue losses.
- Unplanned rod patterns lead to MFLPD management challenges for Operations and re-work for Nuclear Analysis engineers.
- Online values depend upon accurate LPRM and TIP operation.



## Blue Wave MFLPD model performance





#### **Performance Improvement:**

- Mean bias is reduced by a factor of 3.62
- Max Bias reduced by a factor of 2.03

#### **Potential Savings**

- Revise the design to have fewer fuel bundles
- Avoid generation losses due to derates.

#### **Training Set:**

■ The eight previous Fuel Cycles

Note: Scale of BIAS is 10-2. Scale of MSE is 10-4

MFLPD is the maximum over a 30x30x25 array. MSE<sub>array</sub> is across the whole array. MSE<sub>MFLPD</sub> is the squared error of the max values (= MFLPD values)

BIAS is calculated for the MFLPD values

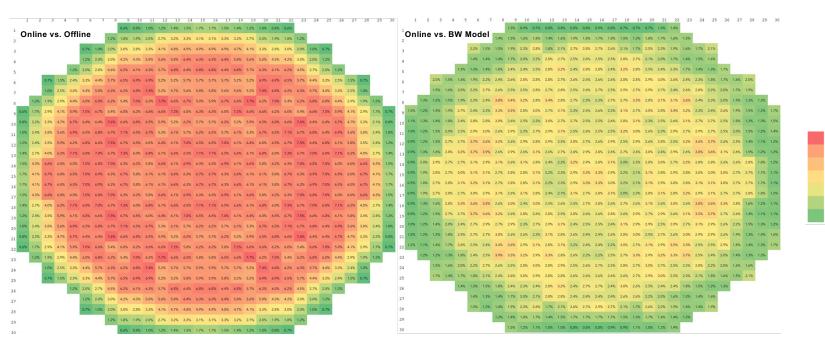
## **MFLPD AXIAL ACCURACY**



The Blue Wave model prediction is significantly more accurate than the physics model prediction

$$\max_{k} |y_{ijk} - y'_{ijk}|$$

$$\max_{k} |y_{ijk} - y_{ijk}^p|$$



Cycle 17

## Flexgenius.ai

#### Local Power Range Monitor (LPRM) Diagnostic Tool

- Display
  - LPRM readings
  - 3D view with anomalous LPRMs highlighted in red
  - List of anomalous LPRMs
  - Rod pattern and power distribution
  - Total power and core flow
- Anomaly Detection Algorithms
  - Drift
  - Frozen LPRM
  - Outlier
  - Change in noise
- Remaining Useful Life (RUL) Algorithm
  - Neural network predicts calibration current based on cycle data
- LPRM Reading Predictions based on ThermalLimit.ai model
- Real Time Data Feed

#### Benefits

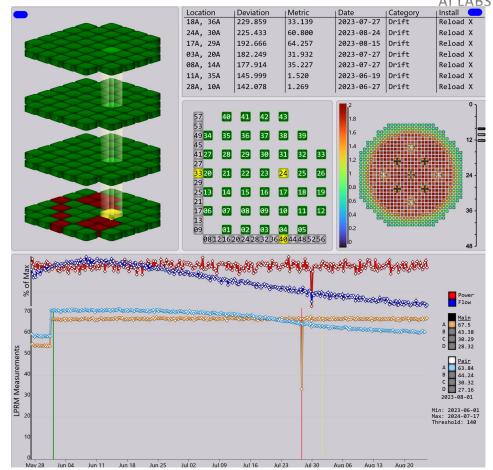
- Significant savings from avoiding unnecessary maintenance actions resulting from incorrect power readings
- Secures reactor safety by helping keep LPRM instrumentation operating and accurately calibrated

#### Future

Support PWR ICI Diagnostics with similar tool





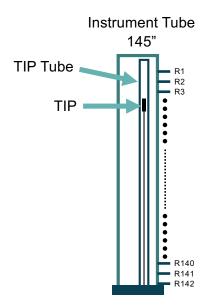


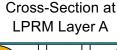


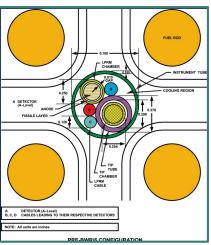
## TIP.ai – TIP TRACE ALIGNMENT



- LPRM Calibration using TIP (Traversing In-Core Probe)
  - The TIP is pushed into the TIP tube from the bottom.
  - The exact vertical location of the TIP at the beginning of the TIP run is unknown.
  - The fuel has 8 spacers at standard distance from each other and the neutron flux at the spacers is lower, which is visible as a dip in the TIP trace. Trace alignment is based on these dips.
- The standard CMS alignment algorithm has proven to be inaccurate.
- Blue Wave has developed a more accurate algorithm
  - A custom filtering kernel with peaks separated by the nominal distance between the spacers in the fuel.
  - The filtering kernel is shifted up and down by varying amounts and convolved with the TIP measurements at each shift value.
  - The location of the minimum of the convolution-vs-shift function corresponds to the shift value that best aligns with the spacers given the nominal separation values.
- Blue Wave TipTrace.ai enables more accurate TIP traces and therefore more accurate LPRM calibration
- Benefits:
  - Accurate LPRM readings are the basis of Thermal Limit Monitoring and avoiding unnecessary generation losses, that can result from incorrect measurements.
  - Incorrect measurements are also a safety issue as Thermal Limits may be underestimated.







Case: Peach Bottom Tip Calibration

| String | Base Method | BW Method | ST_SYM_0 | SYM_BM | SYM_BW | Uncertainty | Coordinates | ∆  Value | Date of Trace       |
|--------|-------------|-----------|----------|--------|--------|-------------|-------------|----------|---------------------|
| 2      | -2          | -2        | 8.68     | 10.39  | 7.86   | 1.2         | (12, 26)    | 0        | 2021-03-16T10:57:54 |
| 3      | 0           | 0         | 3.54     | 7.48   | 3.54   | 0.9         | (16, 26)    | 0        | 2021-03-16T10:57:54 |
| 4      | -2          | -3        | 7.25     | 9.36   | 3.62   | 2.1         | (20, 26)    | 1        | 2021-03-16T10:57:54 |
| 5      | -3          | -2        | 10.73    | 28.78  | 3.4    | 2.3         | (24, 26)    | 1        | 2021-03-16T10:57:54 |



# Case Study

#### Actual Event: Approaching thermal limits limit within a matter of days

- Runaway MFLPD at 0.96 on path to reach procedural limit of 0.98 within days,
   Blue Wave model predicted .918
- Operators thought some LPRMs may need to be placed out of service due to mis-calibration, but couldn't efficiently verify this claim.

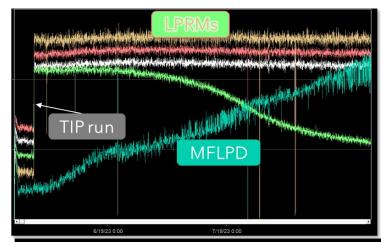
## Without Intervention, a short-term derate would be eventual course of action, followed by insertion of shaper control blades

- MLFPD getting worse, up to 0.975 (model predicted 0.92)
- Operator requested Blue Wave to analyze ALL LPRMs...we did, rank ordering them from most-to-least problematic
- We identified 7 LPRMs with issues, recommending bypass → reducing MFLPD from 0.975 to 0.955

#### Blue Wave predictions proven true and accurate

- Blue Wave recommended performing recalibration with TIP, based on model predictions (still 0.92) ... After TIP the MFLPD went to 0.92!
- Blue Wave tools helped address and closeout three related IRs
- Operator estimates that this support avoided generation losses of approximately \$1.23M
- The Blue Wave Method won NEI Top Innovative Practice award





| LPRM              | Metric  | Notes   |
|-------------------|---------|---|
| 24-49 A - 48-25 A | 229.859 | Noisy   |
| 32-41 A - 40-33 A | 225.433 | Significant Drift, smaller drift in Delta level         |
| 24-41 A - 40-25 A | 192.666 | Bypassed, large offset prior to bypass                  |
| 08-33 A - 32-09 A | 182.249 | Large Drift   |
| 16-25 A - 24-17 A | 177.914 | Fixed offset, not particurlarly drifting apart          |
| 16-49 A - 48-17 A | 145.999 | Large Drift, Present in B and C levels to lessor extent |
| 16-41 A - 40-17 A | 142.078 | Modest Drift  |







## Making a Material Impact on the Nuclear Fleet

with bwnuclear.ai (mco+eigenvalue+thermal limit)



|        |                |                                   | Cycles planned with BWnuclear.ai |                            |                                 |                              |                                |                             |                        |                       |                       |
|--------|----------------|-----------------------------------|----------------------------------|----------------------------|---------------------------------|------------------------------|--------------------------------|-----------------------------|------------------------|-----------------------|-----------------------|
| Unit # | # BW<br>Cycles | Batch Size<br>before<br>BWnuclear | Batch Size<br>2 Cycles<br>Ago    | Savings<br>2 Cycles<br>Ago | Batch Size<br>Previous<br>Cycle | Savings<br>Previous<br>Cycle | Batch Size<br>Current<br>Cycle | Savings<br>Current<br>Cycle | Total Saved<br>Bundles | Fresh Fuel<br>Savings | Spent Fuel<br>Savings |
| 1      | 2              | 272                               |                                  |                            | 268                             | 4                            | 260                            | 12                          | 16                     | \$8M                  | \$1.6M                |
| 2      | 3              | 272                               | 272                              | 0                          | 260                             | 12                           | 264                            | 8                           | 20                     | \$10M                 | \$2M                  |
| 3      | 2              | 272                               |                                  |                            | 268                             | 4                            | 262                            | 10                          | 14                     | \$7M                  | \$1.4M                |
| 4      | 3              | 276                               | 272                              | 4                          | 260                             | 16                           | 268                            | 8                           | 28                     | \$14M                 | \$2.8M                |
| 5      | 2              | 272                               |                                  |                            | 268                             | 4                            | 256                            | 16                          | 20                     | \$10M                 | \$2M                  |
| 6      | 3              | 200                               | 196                              | 4                          | 192                             | 8                            | 188                            | 12                          | 24                     | \$12M                 | \$2.4M                |
| 7      | 2              | 316                               |                                  |                            | 304                             | 12                           | 312                            | 4                           | 16                     | \$8M                  | \$1.6M                |
| 8      | 2              | 308                               |                                  |                            | 300                             | 8                            | 304                            | 4                           | 12                     | \$6M                  | \$1.2M                |
| 9      | 2              | 144                               |                                  |                            | 140                             | 4                            | 140                            | 4                           | 8                      | \$4M                  | \$0.8M                |
| 10     | 1              | 308                               |                                  |                            |                                 |                              | 292                            | 16                          | 16                     | \$8M                  | \$1.6M                |
| 11     | 2              | 208                               |                                  |                            | 204                             | 4                            | 200                            | 8                           | 12                     | \$6M                  | \$1.2M                |
| 12     | 2              | 208                               |                                  |                            | 200                             | 8                            | 200                            | 8                           | 16                     | \$8M                  | \$1.6M                |
| 13     | 1              | 228                               |                                  |                            |                                 |                              | 220                            | 8                           | 8                      | \$4M                  | \$0.8M                |
| Total  | 27             |                                   |                                  |                            |                                 |                              |                                |                             | 210                    | \$105M                | \$21M                 |
|        |                |                                   |                                  |                            |                                 |                              |                                |                             | <b>A</b>               |                       |                       |

\$105M Fresh Fuel Cost Savings\$21M Spent Fuel Cost Savings\$81M Avoided Operational Costs

**\$207M**Total Cost Savings for 13 Units

Average savings per unit per cycle: \$7.7M

Average fresh fuel savings per unit per cycle: 3.8M

27 Cycles Planned with BWnuclear.ai

210 Fuel Bundles Saved

(Fresh Fuel = \$500K/assembly, Spent Fuel = \$100K/assembly)

## BLUE WAVE AI LABS

# Challenges of Documentation in the Nuclear Industry

- Nuclear facilities must maintain comprehensive documentation throughout their lifecycle, from design to decommissioning.
- Nuclear power plants face an average of 173,459
   paperwork hours per facility annually, equivalent to
   about 86 full-time employees dedicated to regulatory
   compliance documentation per plant [1].
- Coupled with the problem of an aging nuclear engineering workforce, this is a problem that could really benefit from AI technologies.



 $\hbox{[1] https://www.americanactionforum.org/research/costs-benefits-nuclear-regulation/}\\$ 

## Blue Wave Nature Language Processing Strategy



## Objectives

- Minimize the Mundane
- Maximize the Staff Hour Savings

## Strategy

 Develop targeted applications relying on Blue Wave Nuclear Corpus and Large Language Models to streamline specific regulatory activities.

## First target

10 CFR 50.59 Screening

#### 50.59 Product Business Case

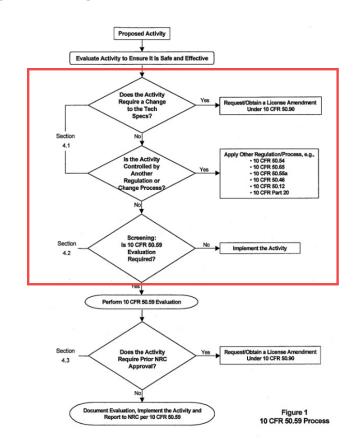
- The effort involved in collecting and reviewing the licensing basis and other documentation required for a 50.59 Screening can take on average 40 hours per screen with hundreds of screenings per year performed at each site.
- All NRC Part 50 licensed reactors maintain a licensing basis and are required to perform 50.59 Screenings against that licensing basis for making changes, and therefore can benefit from an Al tool.
- Added benefits include knowledge transfer, training for inexperienced 50.59 preparers, and investigating the licensing basis for other licensing actions.





## 50.59 Current Workflow

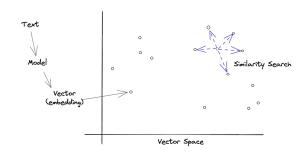
- NEI 96-07, "Guidelines for 10 CFR 50.59 Evaluations," provides detailed guidance on implementing the 50.59 process.
- An engineer performs a keyword-based search over the licensing basis including the Updated Final Safety Analysis Report (UFSAR) and the plant's Technical Specifications.
- Based on their search output and reading of the sections, provide a description of why the section is pertinent.



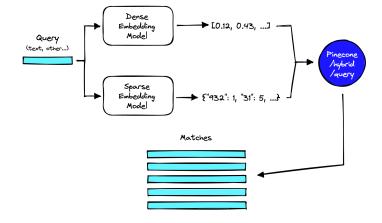


# **Proposed Tool**

- A 50.59 Screener tool would reduce time spent manually searching through the plant's licensing basis
- This is accomplished using a hybrid search solution which uses AI-based semantic search for enhanced context understanding, all the while retaining traditional search functionality
- Conditioned on the input provided and retrieved chunks, a justification is provided as to why the section is relevant.



https://www.tylercrosse.com/ideas/semantic-search



https://www.pinecone.io/learn/hybrid-search-intro/



# Data Pipeline

- The hardest part of any machine learning workflow is the data cleaning and preprocessing. Depending upon the nature of the documents, a rule-based conditional strategy was created.
- We performed several experiments to determine the most optimal chunking strategy eventually settled on a version of the recursive text splitter.
- With the chunks in place, we wrote custom algorithms to tag metadata onto them. For the most part, the metadata is what the end user is interested in.
- Finally, we use an encoder only transformer model to create the vector embeddings.

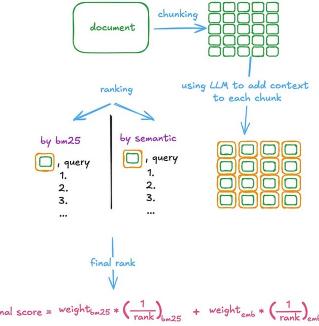




# **Hybrid Search**

- Hybrid search combines traditional keyword search with vector embeddings, capturing both exact matches and the semantic meaning of queries.
- In our approach, we use a Okapi BM25 keyword retriever and a Dense Passage Retriever (DPR) based embedding model as our vector retriever.
- Search results from both retrievers are aggregated, and reranked.
- The retrieved chunks and a few other surrounding ones become the context for generation for the LLM.





https://levelup.gitconnected.com/the-best-rag-technique-yet-anthropics-contextual-retrieval-and-hybrid-search-62320d99004e



# **Summary Generation**

 The retrieved chunks and a few other surrounding ones become the context for generation for the LLM.

Database

Vector store
Index

LLM

Answer

What are sections relevant to Small Break Loss of Coolant Accident (LOCA) analysis and Large Break LOCA analysis?

> Section 14.5.1 Major Reactor Coolant System Pipe Ruptures (Large Break Lossof-Coolant Accident)

This subsection provides detailed information on the analysis of Large Break LOCAs, including the identification of cause and accident description, method of analysis, and results

 $\hbox{@}$  2024 Blue Wave Al Labs. All Rights Reserved.

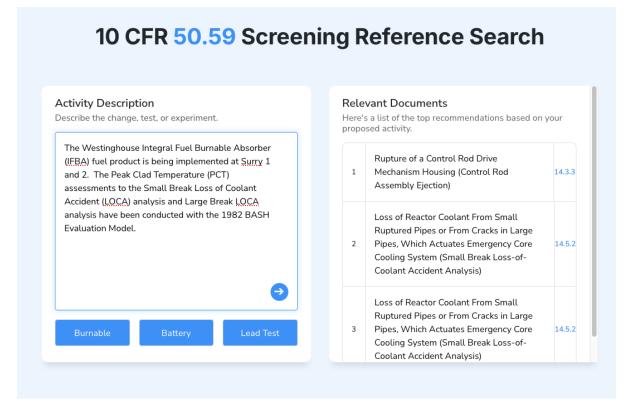
32



## Results and Future Work

- Plenty of excitement regarding the first demo of the screener tool
- Next iterations to fine-tune LLMs to assist with screening decisions
- Lack of domain expertise implies, the LLM cannot "reason" with limited context

Thus, we plan on using historical documents and convert the problem into one of "statistical learning"



© 2024 Blue Wave Al Labs. All Rights Reserved.

