Contrastive Machine Learning and Hyperparameter Optimization for Detecting Nuclear Material Transfers

Jordan Stomps

University of Wisconsin-Madison

ETI Annual Workshop

February 20 – 21, 2024, Golden, CO

How can we build AI models to detect the transfer of nuclear material?

- Traditional methods for tracking nuclear material require the time-consuming manual analysis of measurements
- Use transfer measurements collected around the Multi-Informatics for Nuclear Operating Scenarios (MINOS) tested at Oak Ridge National Laboratory (ORNL)
- Shielded radiological material transfers:
  - Byproducts
    - ^{225}Ac
    - Activated Metals
    - Spent Fuel
  - Nuclear Material
    - Fresh Cm
    - ^{252}Cf
    - Fresh Np
    - Irradiated Np

MINOS Sensor Map

Mission Relevance

- **Goal:** Detect the transportation of shielded radiological material by identifying and characterizing radiation signatures
  - Is a nuclear material transfer occurring?
  - What kind of material is it?
  - How much material is it?
- **Main Objective:** How can Semi-Supervised Machine Learning (SSML) and/or Self-Supervised Learning (SSL) extract information from both unlabeled and labeled data for application in nuclear nonproliferation?
- **Impact:** This research establishes a methodology enabling radiation monitoring in data-rich, label-poor environments

Overview: Contrastive Machine Learning

The contrastive framework has four components:

- **Data augmentation** — transforms samples using a set of augmentation types/rules
- **Base encoder** — encodes representations
- **Projection head** — casts final shape for measuring similarity and calculating loss
- **Contrastive loss function** — maximizes agreement between positive (similar) samples

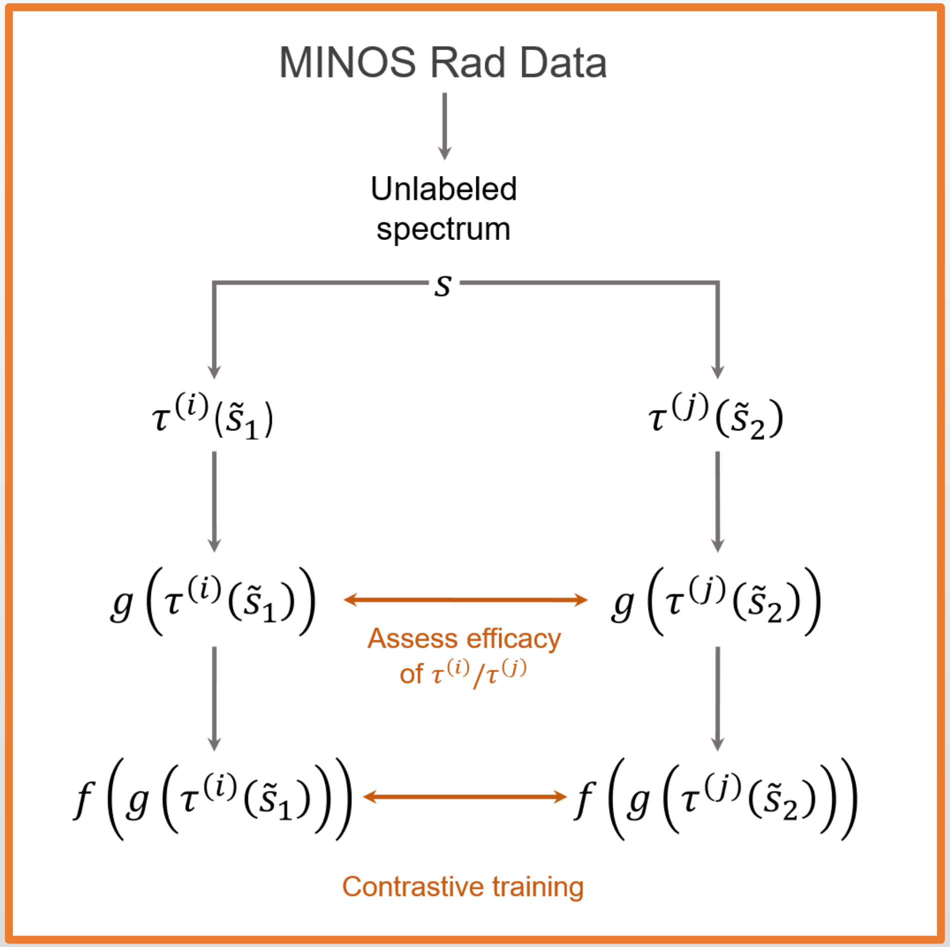

National Nuclear Security Administration

How can gamma spectra be augmented?

- Given an instance s, apply a transform $\tau \in \mathcal{T}$ to get an augmented instance $\tilde{s}$
- Applied augmentations do not obscure labeling information needed for classification (i.e. label-invariant)
- Using baseline estimation and denoising using sparsity (BEADS), each spectrum is decomposed into augmentable components:

$$\boldsymbol{x} = \boldsymbol{s} + \boldsymbol{b} + \boldsymbol{p}$$

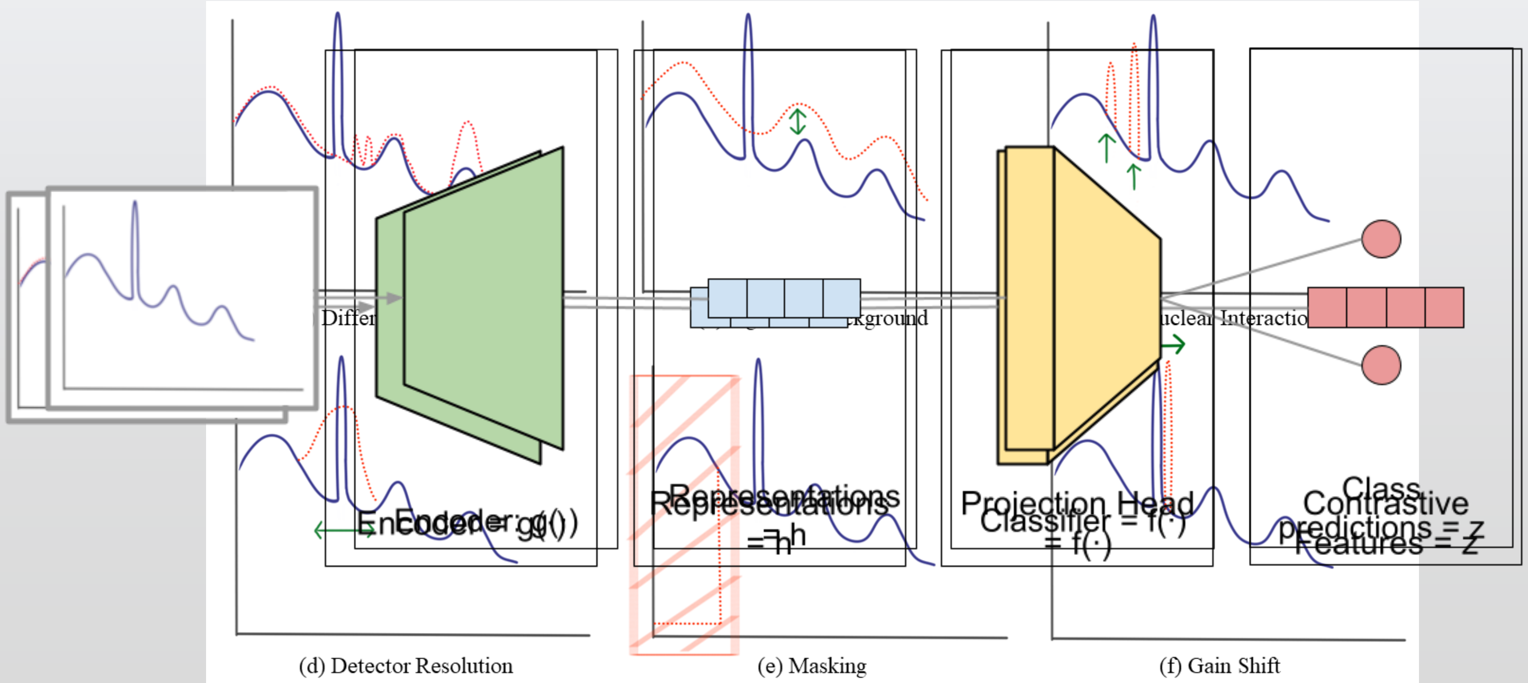

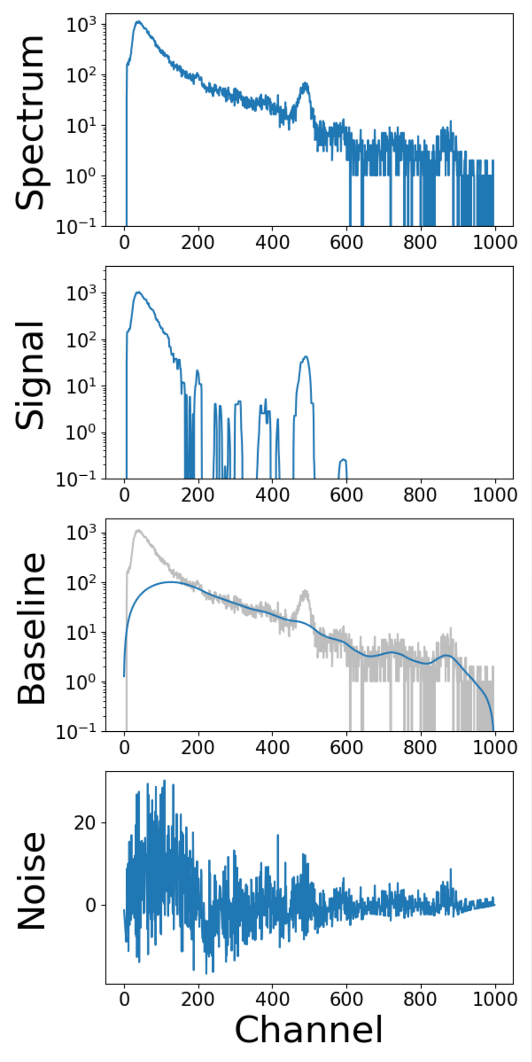

Example BEADS Decomposition

Contrastive models can distinguish between material transfers.

- Encoder is contrastively trained on 68,124 weakly anomalous, unlabeled spectra
- A linear classifier is trained on 52 labeled spectra
- This linear classifier beats Scikit-Learn’s toolbox
  - Supervised models could only train on 52 labeled spectra

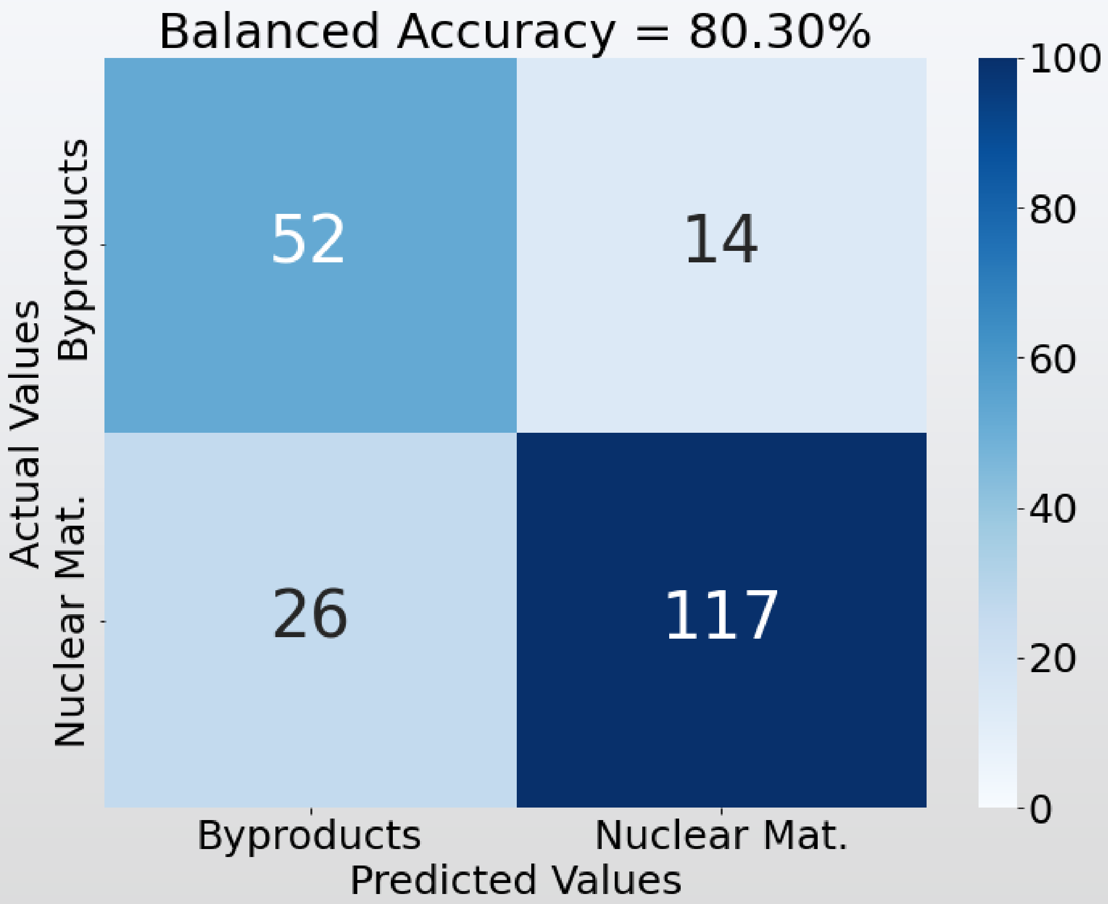

National Nuclear Security Administration

The model detects some transfer types better than others.

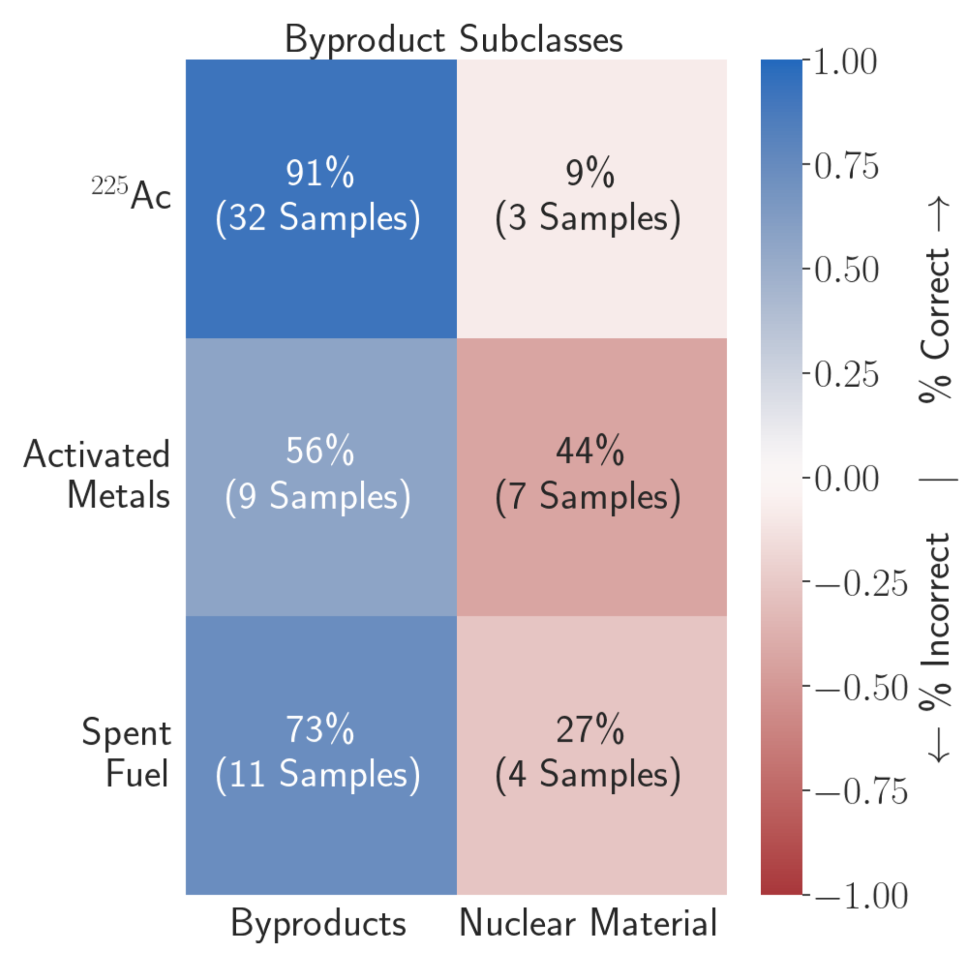

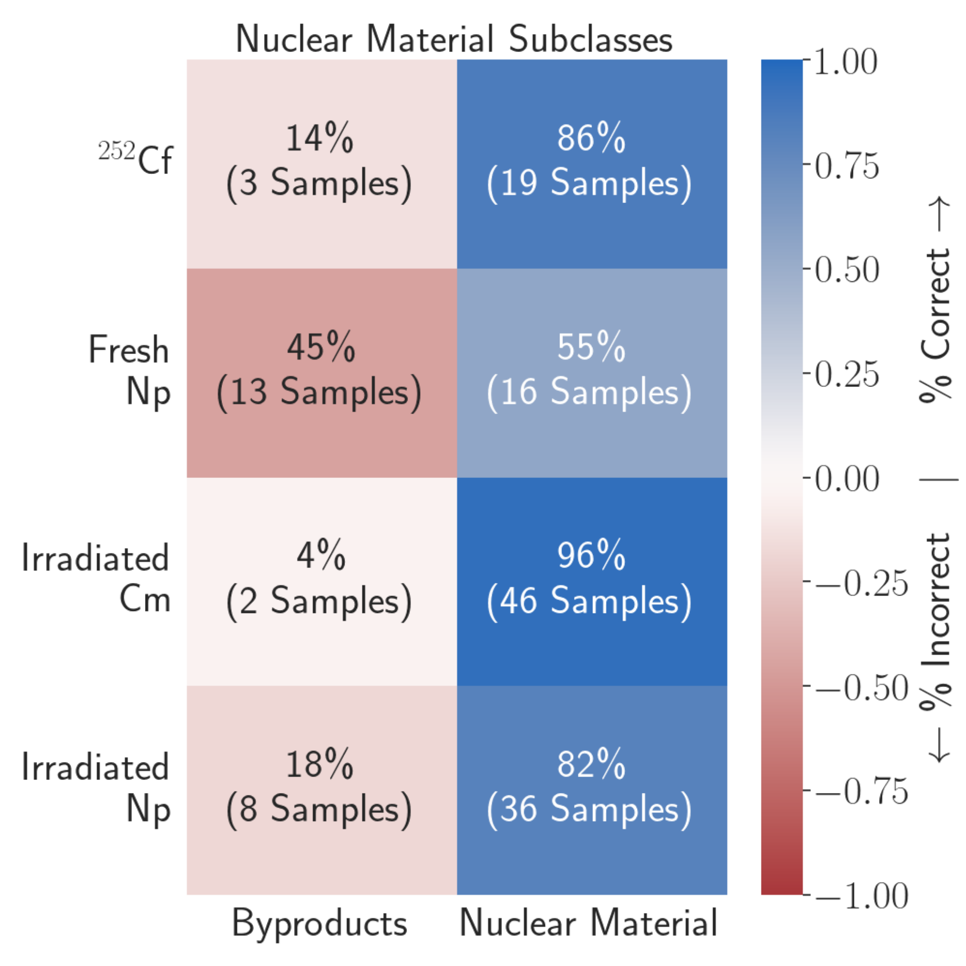

National Nuclear Security Administration

What is the model learning?

- Integrated Gradient can help connect detection to spectral features
- A higher importance means more influence on final classification decision
- Some patterns might be unrealistic, but they mimic regions of interest
- Learned information from unlabeled data is embedded across principal components

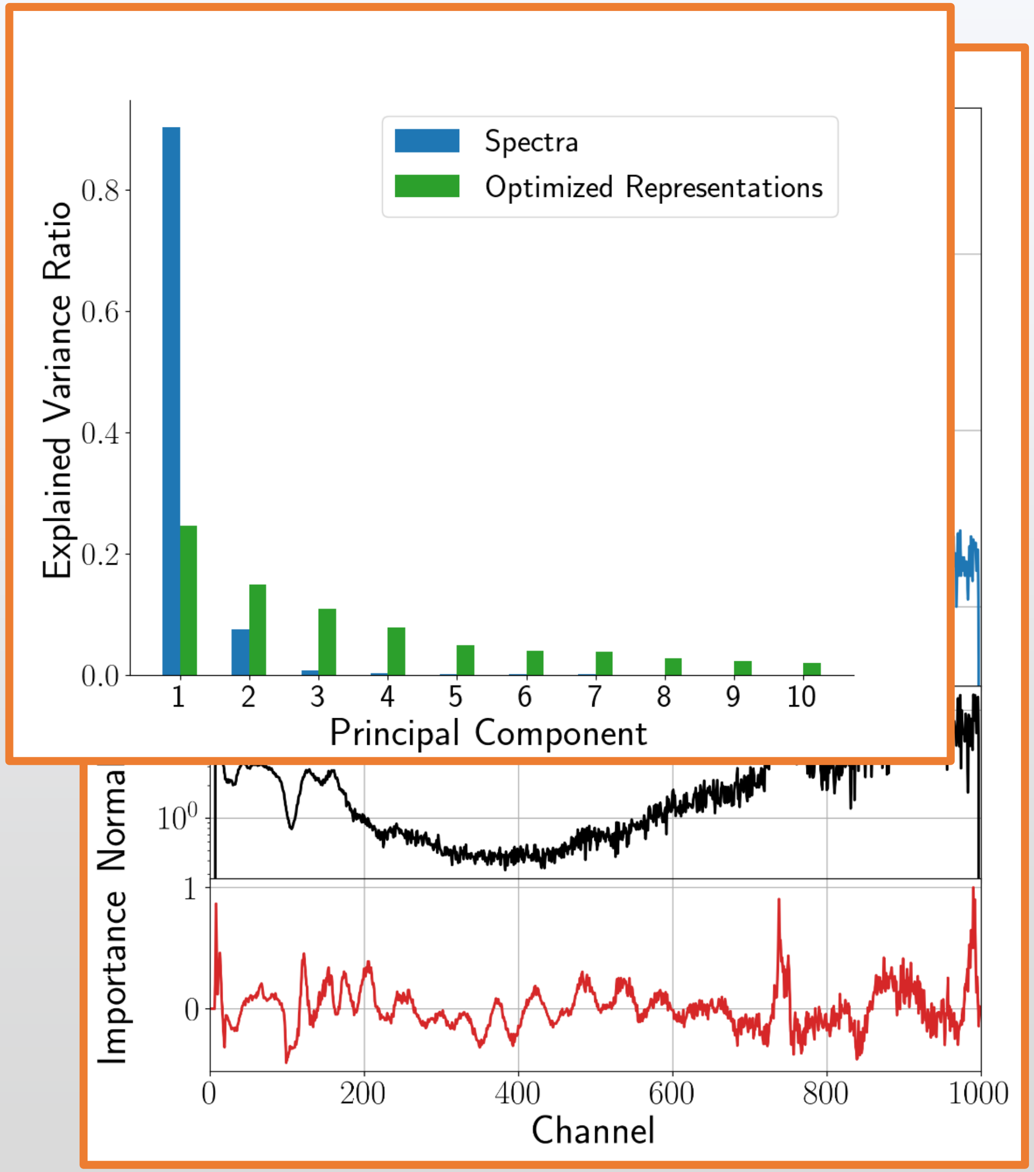

How should scarce labeled data be allocated for training?

- Semi-supervised contrastive learning regularizes representations to the classification tasks
- The classifier needs labeled data for supervised learning
- If labeled data is noisy, more training data does not necessarily mean higher detection accuracy
- If labeled data is limited, training the classifier should be prioritized

Average Accuracy		Representation Model				
		0.00%	10.00%	20.00%	30.00%	40.00%
Classifier	10.00%	75.651	72.827	72.977	76.185	75.496
	20.00%	77.1	76.871	76.795	75.575	-
	30.00%	75.039	77.863	75.42	-	-
	40.00%	78.398	78.932	-	-	-
	50.00%	77.253	-	-	-	-

National Nuclear Security Administration

How long should hyperparameter optimization be run?

- The answer depends on time, computational resources, model complexity, and data quality
- No systematic relationship between hyperparameters and accuracy is observed
- Independent hyperopt sequences quickly converge, but not to the same level

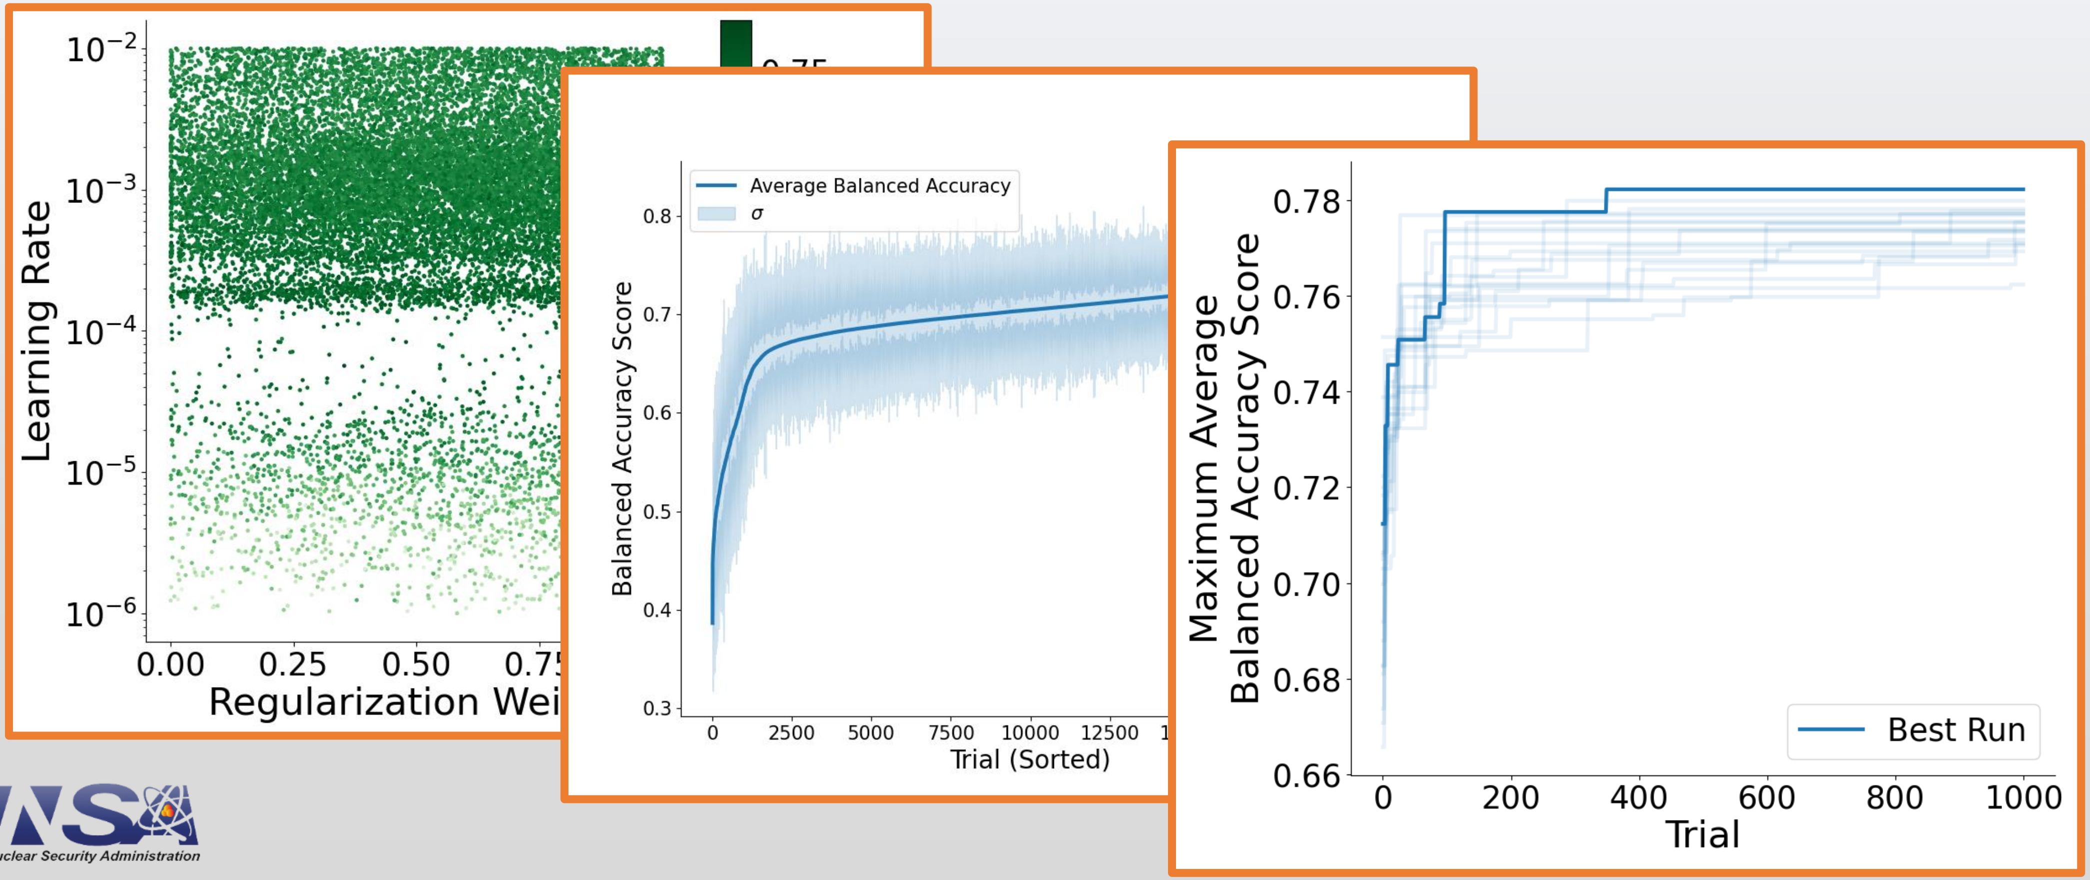

National Nuclear Security Administration

Conclusion

Is there value in using unlabeled data for detection tasks?

- Yes! Information in unlabeled data helps SSML outperform SL techniques.

How can contrastive learning be efficiently used with radiation spectra for material transfer detection?

- By using augmentations specifically designed for radiation detection principles.
- Classifiers built on contrastively trained representations achieve up to 80.30% balanced accuracy.

How can a model be evaluated for explainability?

- PCA and Integrated Gradients suggest relevant pattern recognition for contrastive models.

How can scarce labeled data be effectively allocated to training tasks?

- When labeled data are limited, prioritize classifier training.

When is a contrastive model sufficiently optimized?

- Maximally achieved accuracy quickly converges but must be found empirically.

National Nuclear Security Administration

ETI Impact

- What is the impact of the ETI on your development?
  - Conferences: 2023 INMM & ESARDA Joint Annual Meeting, ETI workshops, University Program Reviews
  - Internship at ORNL
    - NA-22 funded project on robust data analytics for simulated environmental samples
    - Conferences: Conference on Data Analytics (CoDA), ORNL AI Expo
- Personnel transitions: (1/8/2024; Staff Scientist) Nonproliferation Data Scientist at ORNL
- Technology transitions
  - There is an active research interest in leveraging unlabeled or analogous datasets to enhance rare or limited labeled datasets (LDRD on limited data across modalities)
- Thank you to…
  - Advisor: Paul Wilson
  - Mentor: Ken Dayman
  - MINOS Collaboration: Dan Archer, Michael Willis, Andrew Nicholson, and James Ghawaly

ACKNOWLEDGEMENTS

This material is based upon work supported by the Department of Energy / National Nuclear Security Administration under Award Number(s) DE-NA0003921.