Cover sheet

Title: MINOS Analytics Fact Sheets

Authors:
Kary Myers
Garrison Flynn
Tom Reichardt
Ken Dayman
Emily Casleton
Nagi Rao
George Chin
Dave Widemann

# MINOS Fact Sheet: ANALYTICS

## Defining Test Sets in Temporally Correlated Data

## Overview

Analytics models are evaluated using testing data sets that are held out from training data sets. Separation of testing and training is essential for quantifying confidence in a model's predictions. Traditional approaches of selecting testing data at random assumes that 1) the data are independent and identically distributed (iid) and 2) neighboring samples do not inform one another. Nuclear non-proliferation problems are typically driven by event-based classification, presenting an often-ignored challenge where temporal independence is not a valid assumption. Neglecting the effects of temporal correlations with events when defining training sets results in overfitting and risks inflating confidence in future predictions.

## MINOS Context

No two reactor startups (events) are exactly the same. Features may have an observable trend within an event (e.g., strictly increasing), but event-to-event variability can be significantly larger than individual trends. For example, $^{41}Ar$ detected in the effluent stack is an obvious signature of reactor power level within a single event. $^{41}Ar$ is a function of reactor flux as well as the amount of air in the reactor. The feature value of $^{41}Ar$ is strictly increasing as power level increases within a single startup, but becomes confounded with other intermediate power levels when multiple startups are considered. This makes it a good feature for reactor power level, but the influence of air quantity causes indistinguishable event-to-event variability.

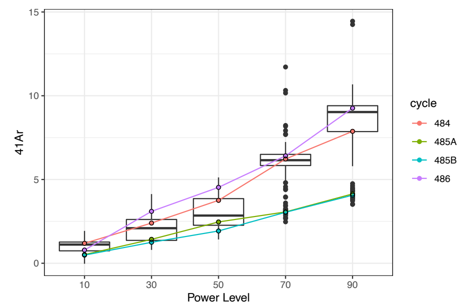


Figure 1. Trends appear across a single startup event but have high variability startup-to-startup.

## Discussion

High autocorrelation of temporally neighboring features within a single event makes predictive modeling a simple interpolation. The trained model is not blinded from any information. However, extrapolative predictions are unlikely to have the same correlations. Therefore, if a model's predictive performance is determined only by interpolative holdout (testing) sets when strong temporal correlation is present then a decision maker is likely to have false confidence in predictions on new data. The solution to avoiding this pitfall in incorrectly applying common machine learning techniques for temporally correlated and non-iid data ***is careful selection of training and testing sets***. The training set should mimic the extrapolation to new data. For the example of reactor start up, an entire startup was held out (Figure 2). This approach will quantify predictive performance realistically and provide more accurate assessments of model robustness.

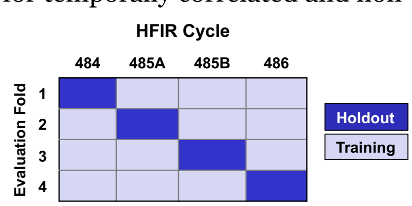


Figure 2. Model cross validation with event-based holdouts

## Further Reading

General introduction to traditional statistical learning can be found in [1]. Discussion of the challenges of correlated and event-based detection in nuclear nonproliferation is presented in [2] and a demonstration of solutions applied to the MINOS problem is found in [3].


[1] James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning*, Vol. 112, New York: Springer.

[2] Flynn G.S. et al. (2019) "Power Level Prediction and Source Localization" LA-UR-19-22790.

[3] Flynn G.S., Parikh N.K, Egid A., Casleton E. (2019) "Predicting Power Levels of Nuclear Reactor by Combining Multiple Modalities," LA-UR-19-26714, *Institute of Nuclear Materials Management*, July 15-18, Palm Desert, CA.

# MINOS Fact Sheet: ANALYTICS

## Unsupervised Learning

## Overview

Unsupervised learning techniques are used to draw inferences from predictor data when response data is limited or unavailable, e.g., by reducing the dimensionality of high-sampling rate datasets and extract physically interpretable features from these datasets. An analyst confronted with many sensors measuring many predictor variables combined with few meaningful observations risks overfitting the data (i.e., small $n$, large $p$). Reducing the dimensionality of the dataset can help identify variables or combinations of variables that best correlate with the system and thereby lower the risk of overfitting (i.e., reducing $p$). Additionally, unsupervised learning techniques allow for more facile data exploration that can help guide the development of supervised models. Finally, a significant opportunity for unsupervised learning is in enabling transferability as labels are not required.

## MINOS Context

Unsupervised learning approaches pursued under MINOS include principal component analysis (PCA) and non-negative matrix- and tensor-factorization (NMF and NTF, respectively). PCA allows the analyst to significantly reduce the dimensionality of the dataset by focusing on the predictors containing the most variance, which in turn aids in data exploration and identification of likely system state predictors. PCA has been applied to a variety of data streams including from seismic, infrasound, current-clamp, and antenna sensors. In the case of infrasound, the first principal component PC1, containing the greatest signal variance, was strongly correlated with reactor power, a discovery which triggered the development of supervised models able to predict reactor power. Similarly, the antenna collections team used PC1 to determine which of their antennas were sensitive to the reactor cooling tower fan motion. In contrast to PCA's extraction of the components encompassing the greatest variance, NMF and NTF provide a parts-based representation of the signal, enabling the extraction of physically interpretable components. NMF of infrasound spectrograms extracted signatures associated with fan motion and water flow, and application of NTF to a ground-probe electromagnetic (EM) dataset – pre-processed via modulation aligned signal projections (MASPs) – extracted features related to the active control of two variable speed cooling fans.

## Discussion

Unsupervised learning provides a way to understand trends in datasets without relying on prior labeled examples beforehand. The understanding they provide can help focus more complex, supervised models in order to avoid overfitting challenging datasets (small $n$, large $p$). But using unsupervised learning techniques alone to predict the system state is challenging because, while able to function without labeled data, it would be difficult to assess how transferable the technique is without access to labeled data. Additionally, changes in the distribution of variance between variables in a dataset can alter the components generated by PCA. This may result in a strongly predictive component becoming less predictive simply because a new source of variance was added to the data.

## Further Reading


- MINOS Fact Sheet on "Small *n*, Large *p*"
- D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," Nature **401**, 788-791 (1999).
- H. Lee et al., "Nonnegative tensor factorization for continuous EEG classification," Int. J. Neur. Syst. **17**, 305-317 (2007).
- S. Wold et al., "Principal component analysis," Chemom. Intell. Lab. Syst. **2**, 37-52 (1987).

# MINOS Fact Sheet: ANALYTICS

## Multiview versus Multimodal Fusion

## Overview

When designing a data fusion system, the types of events or objects of interest to which each sensor is sensitive or blind must be considered as this will determine the types of conclusions the system may draw as well as inform on the optimal ways to perform the data and information fusion. Combining sensors that are sensitive to the same physics yields *multiview* fusion, and the resulting system will return the same types of conclusions as a single sensor, but the system will typically be more robust to changes, fault tolerant, and more precise. If sensors are sensitive to different physics and/or events in the scene, the resulting fusion is *multimodal*, and the system may be draw more complex conclusions about the scene that neither sensing modality alone may discern. The exact properties of the fusion system are also dependent on the architecture of the fusion system. For example, if multiple sensors view the same object or action in the scene, the resulting system is typically more reliable to losses in data and has greater confidence. Further examples of system architectures and the associated ramifications on system performance are discussed in [1].

## MINOS Context

To identify and characterize transfers of radioactive material at the MINOS testbed, the problem was decomposed into an event where a vehicle is observed while radiation is also sensed. Acoustic sensors were used to indicate when a vehicle is present, and radiation detectors were used to indicate when radioactive materials are also sensed. Using either modality alone cannot uniquely identify a transfer of radioactive material. For example, the radiation sensor will false alarm due to $^{41}Ar$ from HFIR or radon washout, but adding the acoustic modality increases the specificity of the system and reduces false alarms. Fusing data collected by seismic sensors to the acoustic data does not add new information but can improve confidence and fault tolerance as seismic sensors can sense vehicles when excess background noise impedes the use of acoustic data. The types of events each combination of sensors is sensitive is shown in Figure 1.

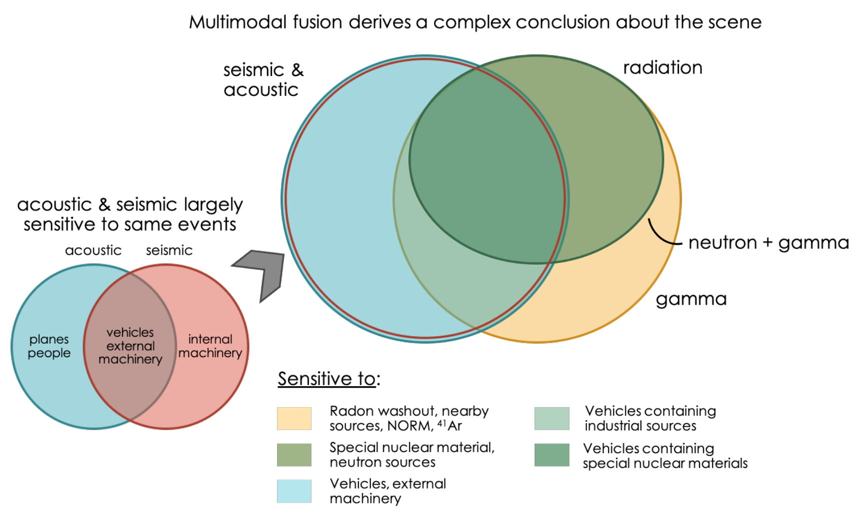


**Figure 1.** Combining multiview and multimodal fusion to identify transfers of radioactive material

## Further Reading

An overview of data fusion system architectures is given in [1], and a report discussing data fusion using artificial intelligence is given in [2]. Further details on analysis of acoustic data for MINOS is given in [3].


[1] W. Elmenreich, "An introduction to sensor fusion," *Vienna University of Technology,* vol. 502, pp. 1-28, 2002.

[2] K. Dayman, B. Phathanapirom, J. Hite, S. Stewart, "Artificial Intelligence—Applications and Implications: Ch. 4: Artificial Intelligence Driven Data Fusion for Security Systems," 2020.

[3] J. Hite and K. Dayman, "Automated Vehicle Detection in a Nuclear Facility Using Low-Frequency Acoustic Sensors," 23rd International Conference on Information Fusion, 2020.

# MINOS Fact Sheet: ANALYTICS

## Denoising Noisy Data

## Overview

Data collection is performed with a specific goal in mind, such as the characterization of a passing vehicle or the identification of nearby nuclear material. However, the data will also contain information about other events, which may or may not be of interest, as well as the environment and random noise.

Denoising the data, through a low-pass filter or more advanced methods, can help separate the signal from the noise, assuming the data is a noisy realization of an underlying signal with some assumed structure.

Data = Desired Signal + Other Trends + Noise

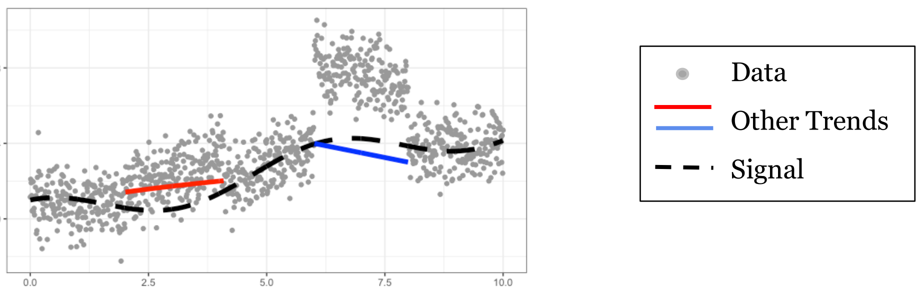


## MINOS Context

- Noisy data tends to exist, although is not exclusive to:
  - Persistently collected data
  - High sampling rates
- Within MINOS, have success denoising data using the 1D fused LASSO [1].
  - Used when coordinates are closely related to neighbors, as in time series data
  - Penalized, flexible, generalized linear regression
  - Results in a piecewise constant fit

Underlying Signal Present

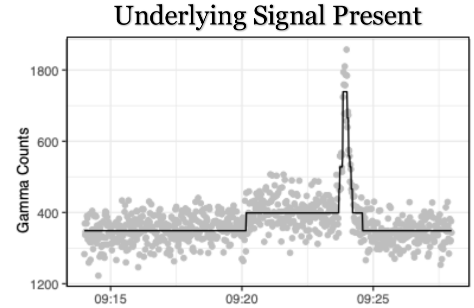


No Signal Present

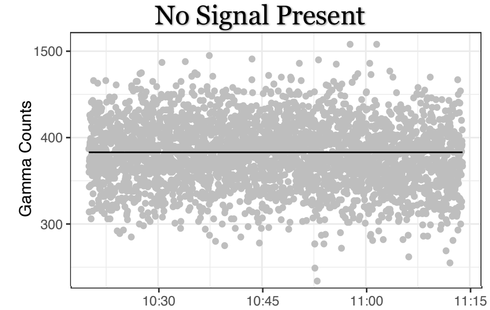


## Discussion

- Denoising data is no substitute for good data collection.
- It is important to define what is "noise". One man's noise is another man's signal.
- Sometimes the data is just noise.
- Within MINOS, the 1D fused LASSO and further processing has allowed for automated denoising and nuclear material transfer path characterization

## Further Reading


1. Tibshirani, R. J., & Taylor, J. (2011). "The solution path of the generalized lasso". *The Annals of Statistics*, 39(3), 1335-1371.
2. Egid A., Osthus D., Woodring J., Weaver B., Klumpp J., Casleton E., Archer D., Nicholson A., Ray W., Garces M., Cardenas E., Chichester D., Reichardt T. (2020) "Characterizing transfers of nuclear material through multi-sensor indicators" *Conference on Data Analysis*, Feb 25-27, Santa Fe, NM.

# MINOS Fact Sheet: ANALYTICS

## Incorporation of Domain Knowledge

## Overview

Domain knowledge from various disciplines, including, physics, chemistry and engineering, has been used to achieve performances superior to purely data-driven approaches for solving analytics problems in complex engineered systems. Incorporation of domain knowledge leads to custom-designed solutions that exploit the underlying laws and system configurations to overcome the challenges of sparse, unstructured and complex datasets. This approach leads to generalizations that transcend the specific scenarios by reducing overfitting and enabling extrapolations beyond the measured datasets.

## MINOS Context

The thermal hydraulics of HFIR's secondary coolant system, which features four pumps and fans, has been exploited to develop a 3-level fuser to estimate the reactor power level by combining thermal, acoustic, and electromagnetic multi-modal features [3]. The fuser's structure reflects the coolant system by aggregating the activities of constant and variable speed fans, and fusing multi-modal features, as shown in Fig. 1, since their combination but not the individual activity or feature adequately reflects the power level. This approach shows successive accuracy improvements of the power level estimate as different sensor modalities are strategically fused in an order that reflects the system knowledge. Classification of irradiated target dissolution events at REDC is addressed by utilizing combinations of isotope decay chains and half-life estimates to select isotope features and time-windows to aggregate measurements [4]. In both scenarios, multiple methods of diverse design are trained on the available data and their predictions are fused based on information fusion theory to achieve performances superior to the individual methods. This information fusion approach and the overall incorporation of system domain knowledge significantly improved the performances over purely data-driven methods.

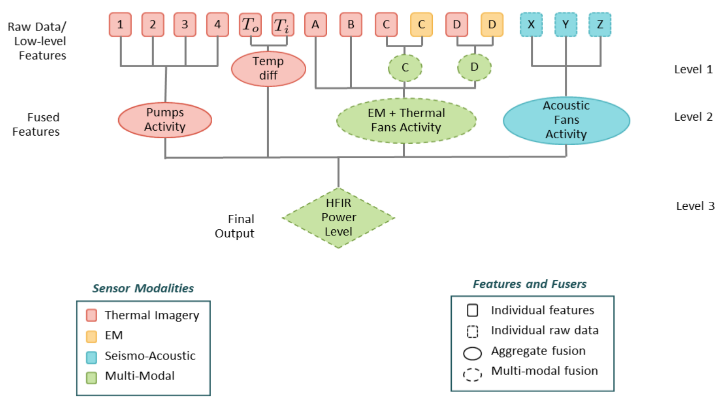


Fig. 1. 3-level fuser for multi-modal power level estimation.

## Discussion

These results represent a starting point of practically effective solutions for nuclear facilities analytics, and also provide foundations for future work needed to develop solutions that transfer to other scenarios by anchoring on invariants derived from the domain knowledge.

## Further Reading

General discussion on data-driven methods for sensor fusion problems may be found in [1-2], and power level estimation and dissolutions classification problems are discussed in [3] and [4], respectively.


1. R. R. Brooks and S. S. Iyengar (editors), Distributed Sensor Networks, 2nd Edition, 2011, Chapman and Hall Publishers. Chapter: N. S. V. Rao, Measurement-based statistical fusion methods for distributed sensor networks.
2. N. S. V. Rao, D. B. Reister, J. Barhen, Information fusion methods based on physical laws, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 1, 2005, pp. 66-77.
3. N. S. V. Rao, C. Greulich, S. Sen, J. Hite, K. J. Dayman, W. Ray, R. Hale, A. D. Nicholson, J. Johnson, R. D. Hunley, M. Maceira, C. Chai, O. Marcillo, T. Karnowski, R. Wetherington, Reactor power level estimation by fusing multi-modal sensor measurements, *International Conference on Information Fusion*, Virtual Conference, 2020.
4. N. S. V. Rao, C. Greulich, S. Sen, J. Hite, K. J. Dayman, A. D. Nicholson, D. E. Archer, M. J. Willis, I. Garishvili, R. Hunley ,J. Johnson, A. J. Rowe, I. R. Stewart, J. M. Ghawaly, Classification of dissolution events using fusion of effluents measurements and classifiers, *Annual Meeting of Institute of Nuclear Materials Management*, 2020.

MINOS Fact Sheet: **ANALYTICS**

## Small *n*, large *p* problem

## Overview

Data-poor (small *n*), feature-rich (large *p*) scenarios are common in real-world applications, and tackling this problem is key. There is currently no agreed upon method to tackle the small *n*, large *p* problem, also referred to as the curse of dimensionality or the short-fat data problem. Statistical approaches require $n>p$, and traditional deep learning practices require a large amount of data from which the network can learn. Furthermore, it is necessary to characterize the error in prediction, which is readily done in statistical approaches but less common in deep learning.

## MINOS Context

The MINOS project has generated a wealth of data spanning sensor modalities. However, the number of unique events is low, and most instances do not have high-resolution labels. This leads to a situation in which we have a small sample size *n* but many features *p*. This is the case in the MINOS project, as few dissolutions are performed over the course of a year and specific materials are transferred at a maximum rate of once per week and a minimum of once, but numerous modalities record the event.

We are interested in the application of deep learning approaches to classify multimodal sensor data, starting with radiological data, and are exploring the application of few-shot learning (see Fig. 1) to address MINOS' material transfer problem. A twin network was trained on pairs of samples from the MUSE radiological dataset that generated a signal upon transfer. The output of this network was a similarity score describing if the two samples are of the same or different classes. Classification was performed via k-testing, in which each sample was compared against a reference sample and the reference that gives the highest score was taken as the class. We also performed uncertainty quantification using the dropout as Bayesian approximation approach. Through this approach, we were able to quantify the uncertainty surrounding a single sample.

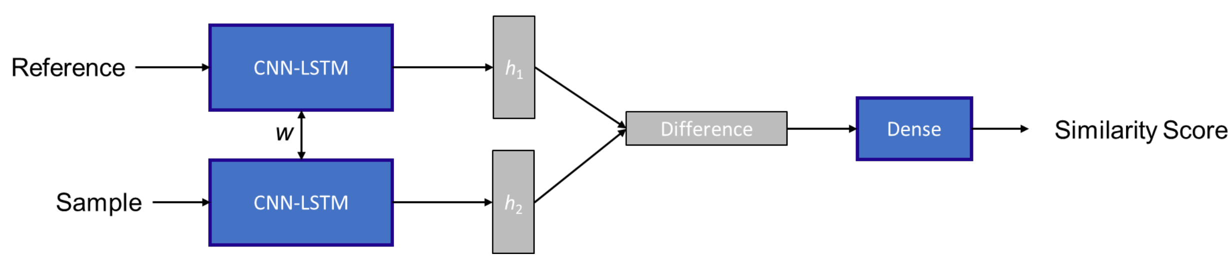


Figure 1: Schematic of the twin network in the one-shot learning method applied to the material transfer use case. CNN-LSTM represents the twin feature extraction layers (convolutional neural network (CNN) + long-short term memory (LSTM)) in which weights $w$ are shared during training. The difference in the tensors $h$ output from the feature extraction layers are passed to a dense layer from which a similarity score is generated.

## Further Reading


- MINOS End of Year Report FY20
- G. Koch, R. Zemel, and R. Salakhutdinov, Siamese neural networks for one-shot image recognition, International Conference on Machine Learning (ICML) Deep Learning Workshop, Vol. 2, (2015).
- O. Vinyals, C. Blundell, T. Lillicrap, and D. Wierstra, Matching networks for one shot learning, Advances in Neural Information Processing Systems, (2016), pp. 3630-3638.
- W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, J.-B. Huang, A closer look at few-shot classification, International Conference on Learning Representations (ICLR), (2019).

# MINOS Fact Sheet: ANALYTICS

## Data Analysis Pipeline: Detector before Classifier

### Overview

In industrial signal processing applications, it can be common to have a complex data analysis pipeline that requires multiple steps. This pipeline is even more relevant when building large parameter neural network-based classifiers for sparse datasets. Below, we describe why and how our detector-to-classifier pipeline is used for MINOS.

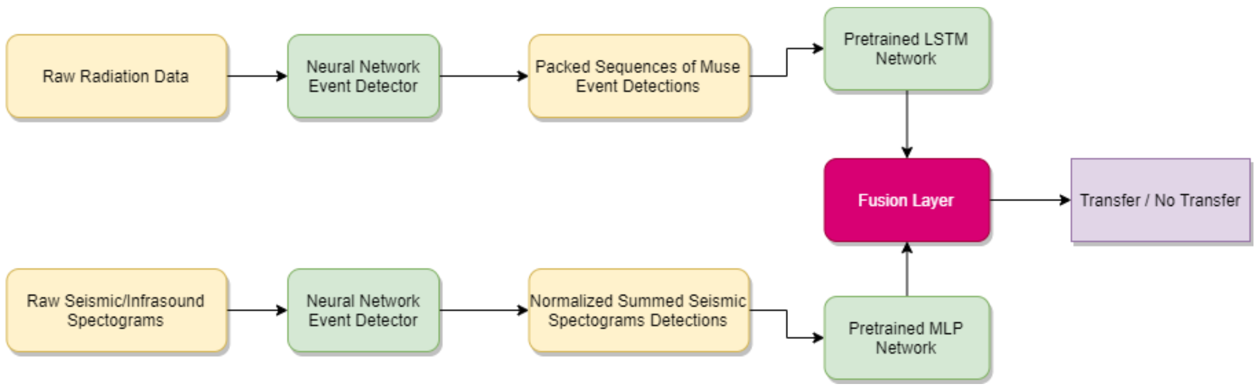


Figure 1: Detector to classifier pipeline for multimodal MINOS datasets.

### MINOS Context

This discussion is for MINOS seismic, infrasound and radiological sensors data. The classification algorithm must ingest a day's worth of data and then make a binary prediction, determining if a transfer occurred on that specific day.

### Discussion

Our solution is to build a signal detector for each data modality. The purpose of the detector is to rapidly scan a day's worth of data and extract only signals that are above the background noise. Thus, a day's worth of data is compressed, by a factor of ~99%, into a list of detections that are relevant to prediction of material transfer. Each detection contains the signal's spectrum, sensor location and timestamp. A material transfer classifier is then trained on the list of detections.

There are two risks associated with the detector to classifier pipeline described above. First, the performance of the classifier is directly related to the accuracy of the detector. If detector misses too many signals, then the classifier will not have the data to make the correct prediction. To mitigate this risk, we suggest lowering the detection threshold, so as not to have any false negatives.

The second risk associated with this approach is that the classifier model must be able to handle variable length input data. For each day, the number of detections can vary by a large amount, e.g. 2 to 128 detections. To mitigate this risk, we built a recurrent neural network, using Long Short-term Memory gates. This network is capable of handling variable length sequences and is designed to find temporal patterns amongst the detections.

### Further Reading

- MINOS End of Year Report, FY2020
- Saini and Widemann LLNL MINOS Q3 slide deck